Mathematical Modelling of Social Factors in Decision Making Processes at the Individual and Population Levels

by

John Lang

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Applied Mathematics

Waterloo, Ontario, Canada, 2016

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of contributions

- Chapters 2-3:
  - The research that resulted in Chapters 2-3 was performed by John Cameron Lang under the supervision and guidance of Dr. Hans De Sterck.
  - Chapters 2-3 are based on manuscripts published in *Mathematical Social Sciences* [109] and the *Journal of Complex Networks* [110], respectively. Both manuscripts were primarily written by John Cameron Lang, with suggested revisions and amendments provided by Dr. Hans De Sterck.
- Chapter 4-5:
  - The research that resulted in Chapters 4-5 was performed by John Cameron Lang under the supervision and guidance of Drs. Daniel M. Abrams and Hans De Sterck.
  - Chapter 4 is based on a manuscript published in *BioMed Central (BMC) Public Health* [108] that was primarily written by John Cameron Lang. Suggested revisions and amendments were provided by Drs. Daniel M. Abrams and Hans De Sterck.
  - Chapter 5 was primarily written by John Cameron Lang. Suggested revisions and amendments were provided by Drs. Daniel M. Abrams and Hans De Sterck.

Abstract

In this thesis we apply mathematical modelling techniques to investigate the implications of social influence on decision making processes in two related contexts.

The first problem concerns the mathematical modelling of civil unrest. We consider the collective action problem facing individuals who are deciding whether or not to join a political revolution or protest in a dictatorial regime that employs censorship and repression. In studying this problem we develop both a population-level model and a network-based individual-level (or agent-based) model. The population-level model establishes a conceptual framework that can be used to understand the role that new communication technologies (e.g. the Internet, satellite television, Short Message Service (SMS) text messaging, social media, etc.) may have played in facilitating the political revolutions of the Arab Spring. We establish the consistency between the individual-level model and the population-level model, and show methodologically how these two modelling strategies can be applied to complement one another, establishing a hierarchy of differential equation models that explicitly take the structure of the social network into account. Finally, using proxy network data for network structure pre- and post-adoption of new communication technologies, we perform small-scale computational simulations of our individual-level model in order to establish quantitative evidence that the political revolutions of the Arab Spring may have been facilitated by new communication technologies.

The second problem concerns the spread of smoking and obesity in populations. We consider two conformity problems that individuals face when deciding whether to join one population sub-group over another (or possibly over many others) in the context of two non-communicable diseases. We begin by studying the smoking epidemic over the past century, where individuals are given the choice to smoke or not to smoke. We establish a new data set for smoking prevalence over the past century in seven developed countries and use it to calibrate a population-level mathematical model for the dynamics of smoking prevalence. We compare our model's predictions to an independently established measure of individualism/collectivism, i.e. Hofstede's Individualism versus Collectivism (IDV) measure, and find evidence that a society's culture can have a quantitative effect on the spread of a contagion. Finally, we study the dynamics of individuals' body mass index (BMI - defined as weight divided by height squared). We establish an individual-level model that also has implications at the population level. At the population level our model fits empirical BMI distributions better than the log-normal and skew-normal distribution functions, i.e. two distributions commonly used to fit right-skewed data, and provides a mechanistic explanation for the right-skewness observed in empirical BMI distributions. At the individual level our model is able to reproduce the average and standard deviation

in individuals' year-over-year change in BMI. At both the individual and population levels our model finds evidence in support of the hypothesis that social factors play a role in the dynamics of individuals' BMI.

Acknowledgements

A special thanks to Drs. Hans De Sterck and Daniel M. Abrams for guiding me throughout my thesis.

Thank you to the thesis examination committee for their time and attention.

Thank you to the entire Applied Mathematics Department at the University of Waterloo for helping to accommodate my special circumstances.

Thank you to Drs. Thomas Homer-Dixon, James Fowler, and Matthew Scott for the insightful conversations.

Dedication

To the friends and family that have made my life so rich.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

Human beings are undoubtedly social creatures. Consciously or unconsciously, conspicuously or inconspicuously, the decisions we make and the actions we take are being influenced by the innumerable social interactions that are an unavoidable part of everyday life. The primary focus of this thesis is to apply mathematical modelling techniques to investigate the implications of social influence on decision making processes. Broadly speaking, we investigate two different ways in which social influence manifests itself. First, in Chapters 2-3 we investigate collective action problems, i.e. tasks whose successful completion requires the participation of a minimum fraction of the population. Second, in Chapters 4-5 we investigate problems involving conformity, i.e. situations where individuals choose to conform their behaviour to one social group over another (or possibly over many others). The methods used in this thesis vary depending on the problem under consideration and fall into two broad categories. At one extreme, we use agent-based models (ABMs) [20, 117] that model the behaviour of each individual in a population. At the other extreme, we use population-level models that model the aggregate behaviour of large sub-groups in a population [99, 101, 191]. We now provide additional context for collective action and conformity problems and compare and contrast the methods used in this thesis in Sections 1.1-1.2, respectively, before summarizing the main results and contributions of this thesis in Section 1.3.

1.1 Collective Action and Conformity Problems

1.1.1 The Stag Hunt Game

Of all the animal species on Earth, human beings are among the most, if not the most, cooperative. Even when compared against our nearest neighbours (e.g. chimpanzees), humans are especially cooperative beings [172, 173]. Although exactly how the cooperative nature of human beings evolved is not known, and indeed can likely never be known for sure, it seems likely that human beings' cooperative nature evolved in response to selection pressure that can be conceptually identified with "The Stag Hunt" game [28], a symmetric two player game that is played as follows. Two players each choose between cooperating (c), i.e. hunting a stag collaboratively, or defecting (d), i.e. hunting a rabbit individually. Capturing a stag requires that both players cooperate and results in a large payoff, whereas capturing a rabbit is a purely individual effort and results in a lower payoff. If both players cooperate, then they succeed in capturing a stag and they each receive the highest possible payout (p_3). In contrast if they both defect, i.e. if they each choose to hunt rabbits individually, then they each receive a small payout (p_1). If one cooperates and the other defects then the defecting player gets a medium payout (p_2)[1], i.e. they successfully hunt a rabbit, and the cooperating player gets a small penalty (p_0), i.e. they receive neither the payoff from the capture of a rabbit or from a stag. The magnitudes of the payoffs are $p_3 > p_2 \geq p_1 > p_0$. This game is illustrated in Fig. 1.1.

The Stag Hunt game is useful as a conceptual tool for understanding the most simplified form of a collective action problem, i.e. a task whose successful completion requires a certain minimum collective participation. The Stag Hunt game has two Nash Equilibria[2]: one where both players cooperate (c, c) and one where both players defect (d, d). Clearly the desired equilibrium is the cooperative one, however, maintaining this equilibrium is non-trivial. Intuitively, we expect that maintaining the cooperative equilibrium (c, c) requires

[1] We note that the payoff from hunting a rabbit depends on whether your opponent also chooses to hunt a rabbit, i.e. $p_2 \geq p_1$ in order to allow for the possibility that capturing a rabbit is easiest when you are not in competition with your opponent for rabbits.

[2] In a multiplayer game a Nash Equilibrium is a set of strategies for each player such that (a) each player knows the strategies of the other players, and (b) given the strategies of the other players no player can improve their outcome by changing strategy. For example, in a two player game a pair of actions (x, y), i.e. where Player 1 chooses x and Player 2 chooses y, is a Nash Equilibrium if (a) given that Player 2 chooses y, Player 1 achieves their highest possible payoff from action x, and (b) given that Player 1 chooses x, Player 2 achieves their highest possible payoff from action y. In the Stag Hunt game (c, c) is a Nash Equilibrium because, given that Player 1 chooses to cooperate, Player 2 achieves their highest payoff from also choosing to cooperate, and vice versa.

Player 1\2	c	d
c	p_3, p_3	p_0, p_2
d	p_2, p_0	p_1, p_1

Figure 1.1: The Stag Hunt game. Rows represent possible actions of Player 1. Columns represent possible actions of Player 2. Payoffs are given as ordered pairs for Players 1 and 2, respectively, where $p_3 > p_2 \geq p_1 > p_0$. For example, if Player 1 defects (d) and Player 2 cooperates (c) then Player 1 receives payoff p_2 and Player 2 receives payoff p_0.

that players be able to communicate with and also to trust each other.

Using the Stag Hunt game as a point of reference, one possible narrative for the evolution of the especially cooperative nature of human beings, as posited by evolutionary anthropologists [172, 173], proceeds in two stages. In the first, humans are exposed to evolutionary pressure in the form of small-scale collective action problems analogous to the Stag Hunt game. Successful adaptation in the first stage resulted in an increasing population, which in turn led to competition between groups of individuals. In this second stage, collaboration between individuals (and groups of individuals) is drastically scaled up in order to defend against other groups of humans, i.e. invaders. The necessity of intergroup collaboration in order to defend from invaders necessitated organization of small groups, i.e. tribes, into larger groups. Communication, coordination, trust, and ultimately collaboration, between these small tribes laid the foundation for the development of culture. Individuals of different, but related, tribes became able to identify potential partners for collaboration through shared language, foods, and social norms. A corollary of the development of culture and group identity, therefore, is the desire of and for individuals to conform to social norms.

Although the narrative presented above is only one of many, it is nevertheless useful in that it provides a conceptual framework for the two broad themes that will be discussed in this thesis: collective action problems (see Chapters 2-3) and conformity problems (see Chapters 4-5).

1.1.2 Collective Action Problems

An early mathematical model for collective action problems in the context of social spreading phenomena was provided by Granovetter in 1978 [75]. This work considers a binary choice collective action problem, i.e. individuals are given a choice between one of two alternatives, where the payoffs depend on the choices made by the population in aggregate. Examples of phenomena where this model might apply include, but are by no means limited to, the enforcement of unpopular norms [36], social influence in marketing [98], the spread of health behaviours [27, 106], rumour spreading [153], and even in engineering applications (e.g. the spread of failures in electrical grids) [187]. Most of these situations can be summarized conceptually as follows. Individuals are given the choice between inaction and action. The payoff for inaction, i.e. the payoff of their default state, is zero. The payoff for action depends on the number of other individuals that also chose action. Individuals who choose action are choosing to contribute to a collaborative endeavour that will only succeed if sufficiently many individuals participate. If only very few individuals choose action then the collaboration fails and their payoff is negative. However, if a sufficiently large number of individuals choose action then the collaboration is successful and their payoff is positive. Individuals will therefore only choose action if they expect a sufficiently large fraction of the population to also choose action. This situation can be modelled mathematically by a *linear threshold process* where each individual i is given a threshold θ_i and will only choose action if the fraction of the population that has chosen action meets or exceeds θ_i. We note that, given this characterization, these processes can be considered as straightforward generalizations of the Stag Hunt game.

The concept of the linear threshold model was revived in the context of political revolution after the 1989 Eastern European revolutions [107]. A linear threshold model was applied in this context in order to explain the unanticipated nature of the 1989 revolutions as well as the incredible speed with which they spread. The argument presented is that, while much of the population may have relatively high thresholds for action θ_i, a cascade of activations may nevertheless be possible so long as a critical fraction of the population has developed sufficiently low thresholds. This provides a satisfactory explanation for the speed with which the revolution occurred, since the sequence of activations of low to high threshold individuals acts as a chain reaction. It also provides a satisfactory explanation for the unforeseen nature of the revolution, since only a relatively small minority of the population is required to have low thresholds (i.e. only a relatively few people need to be biased towards action).

Although the distribution of thresholds is of critical importance in a linear threshold model, so too is the structure of the underlying communication network. Indeed, if you are

isolated and unable to communicate or coordinate your actions with other individuals then you will likely never choose to act, regardless of the value of your threshold for action θ_i. In 2010-2011 another wave of massive protests culminated in revolutions in several Arab countries [19, 92, 152]. As with the 1989 revolutions in Eastern Europe, these political protests and revolutions, which came to be collectively known as the Arab Spring, were largely unforeseen and spread with an incredible speed across a huge geographical area. What was notably different from the 1989 revolutions, however, was the degree to which the Arab Spring was broadcast to the world. New communication technologies allowed real time communication and coordination between protesters and the rest of the world. Indeed, the dominant narrative that was developed by the media was how the Internet and social media played a critical role in facilitating the Arab Spring [5, 89, 92, 100, 116, 144, 162, 165, 192]. In Chapters 2-3 we apply a linear threshold model to investigate this hypothesis and to describe a mechanism through which new media technologies may have facilitated the protests and revolutions of the Arab Spring.

1.1.3 Conformity Problems

In the binary choice collective action problems studied in the previous section, individuals are given a choice of action or inaction. If they choose action then they are choosing to join a collaborative effort, the outcome of which depends on the number of collaborators they have. In contrast, in the binary choice conformity problem individuals are given the choice to belong to one of two population sub-groups: A or B, or equivalently A or *not* A. Instead of choosing whether or not to join a collective endeavour, individuals are choosing whether or not to join a collective identity. But, suppose that we identify the choice to belong to group A with action and the choice to belong to group B with inaction. Then, using the terminology developed for collective action problems in the previous section, the basic premise of the binary choice conformity problem can be described as follows. Individuals are given a choice of either action (i.e. joining group A) or inaction (i.e. joining group B). Individuals who choose inaction receive a composite payoff made up of the *individual utility* they derive from inaction (i.e. the utility they derive purely from conforming their behaviour to the standards of group B) and the *social utility* they derive from social interactions with other inactive individuals (i.e. the utility they derive from social interactions with other group B members). Similarly, individuals who choose action receive a composite payoff made up of the *individual utility* they derive from action (i.e. the utility they derive purely from conforming their behaviour to the standards of group A) and the *social utility* they derive from social interactions with other active individuals (i.e. the utility they derive from social interactions with other group A members). While

at first glance the binary choice collective action and conformity problems appear to be distinct, they are in fact very closely related - at least from the mathematical modelling perspective adopted in this thesis. We remark that this observation is consistent with the narrative developed in Section 1.1.1, which identifies both collective action and conformity problems as potentially having a common evolutionary origin.

In Chapter 4 we apply the concept of a binary choice conformity problem to study the dynamics of smoking prevalence. Specifically, individuals are given a choice between smoking (i.e. belonging to group A) and not smoking (i.e. belonging to group B). They are assigned a composite payoff made up of individual utility (i.e. smokers receive utility derived purely from the act of smoking and non-smokers receive utility derived purely from the act of not smoking) and social utility (i.e. smokers receive utility from interactions with other smokers, and non-smokers receive utility from interactions with other non-smokers). Using this conceptualization to frame the problem of modelling the dynamics of smoking prevalence, Chapter 4 applies a mathematical model that is closely related to an antecedent model of social group competition applied to both the dynamics of language death [1] and the decline of religious affiliation [2]. The first case divides the population into a group A that speaks a particular minority language (e.g. Scottish Gaelic in Scotland, or Quechua in Peru) and a second group B that does not speak this minority language. Over the period of time that the model is applied (roughly a century) the groups speaking minority languages decline in size monotonically. The second case divides the population into a group A that is religiously unaffiliated and a second group B that is religiously affiliated. Over the period of time that the model is applied (again, roughly a century) the groups of religiously unaffiliated increase in size monotonically. We note that in both of these previous cases the dynamics that are being modelled are monotonic. In contrast, the dynamics of smoking prevalence over the past century are not monotonic, and therefore, the model presented in Chapter 4 must be modified accordingly, see Section 1.3 or Chapter 4 for additional details.

A natural generalization of the binary choice conformity problem is the multiple choice conformity problem, i.e. the situation where individuals are given the choice to belong to one of many different population sub-groups. This is the case studied in Chapter 5, which generalizes the mathematical model presented in Chapter 4 (and also [1, 2]) for the dynamics of the body mass index (BMI). Since BMI can take any value in a continuous range, this model presents individuals not with a choice between membership in one of two groups, but between membership in one of an infinite number of groups.

In both Chapters 4 and 5 we develop our experimental design in order to investigate the hypothesis that social interactions play an important, and indeed quantifiable, role in contagion processes. Specifically, we find that for two epidemics of non-communicable

diseases (i.e. the smoking and obesity epidemics) we are able to provide evidence in support of this hypothesis by investigating the importance of the contribution from *social utility* with respect to *individual utility*, see Section 1.3 and Chapters 4-5 for additional details.

1.2 Agent-based versus Population-level Modelling

Population-level mathematical modelling is a modelling strategy that aggregates individuals according to some characteristic and then formulates a mathematical model in terms of individuals' average behaviours. The classic example of a population-level model is the Kermack-McKendrick Susceptible-Infected-Removed (SIR) model for the evolution of an epidemic [99][3], which divides individuals into three compartments according to their status: Susceptible (S) individuals may become infected through contact with an Infected individual (I), and Infected individuals become Removed (R) through recovery or death. Once compartments have been established, the model is formulated in terms of individuals' average behaviour by specifying an ordinary differential equation (ODE) for the evolution of each compartment. For the Kermack-McKendrick SIR model, with infection (β) and removal (γ) rates that are assumed to be constant over the lifetime of an Infected individual, the system of ordinary differential equations is

$$\begin{aligned} \dot{S} &= -\beta SI, \\ \dot{I} &= \beta SI - \gamma I, \text{ and} \\ \dot{R} &= \gamma I, \end{aligned}$$

where $\dot{X}$ denotes the time derivative of X. Because the Kermack-McKendrick SIR model divides the population into three distinct compartments, this population-level model is commonly known as a compartmental model. Whereas the SIR model divides the population into distinct compartments, however, there exist applications where individuals can take on a continuum of states, e.g. see Chapter 5 which models individuals' body mass index (BMI). In this case, the population-level model is given by a partial differential equation (PDE) that describes the dynamics of the distribution of individuals over the continuum of possible states. We note that this partial differential equation can be thought of as analogous to the compartmental model in the limit where the number of compartments goes to infinity.

[3]For more contemporary examples of simple compartmental models and their applications we suggest [101] and [191].

At the opposite extreme of population-level models are individual-level models, i.e. agent-based models (ABMs) in which the behaviour of each individual is specified separately. Recently, despite difficulties in analyzing ABMs, they have become increasingly popular due to advances in scientific computing and techniques in numerical and computational analysis that have allowed for the efficient computational analysis of ABMs via simulation. As a result, there is a growing body of research on the strengths and weaknesses of ABMs relative to compartmental models [114, 136, 149].

As an undergraduate student I was taught that when developing a mathematical model there are two important factors that need to be balanced. First, it is important that the mathematical model be sophisticated enough to be able to capture the essence of the phenomenon being studied. Second, it is important that the mathematical model be kept as simple as possible. The first point is self-evident. Obviously, a mathematical model that is incapable of capturing the phenomenon being modelled is useless. The second point, although less obvious than the first, is really just a rephrasing of Occam's razor in the context of mathematical modelling. Alternatively, the second point, i.e. our preference for simplicity over complexity, can be justified by considering the tradeoffs inherent when comparing a simple model to a more complex model. Specifically, although more complex models (e.g. ABMs) are better able to capture nuanced behaviour than simple models (e.g. population-level models), it is important to note that they may also (a) require additional assumptions in order to specify behaivour of individuals in greater detail, (b) become less analytically tractable and more computationally expensive to simulate, (c) require the calibration of many more additional parameters, and (d) result in overfitting and otherwise spurious results. In this thesis we attempt to apply ABMs and simple population-level models in parallel in order to take advantage of their complementary strengths and weakness. For example, although simple population-level models typically admit a complete mathematical analysis and straightforward interpretation, they are limited in their ability to model fine scale structure or nuanced phenomena. Conversely, although ABMs are able to model fine scale structure and nuanced phenomena, they generally do not admit a complete mathematical analysis and are susceptible to overfitting. Using both ABMs and simple population-level models together allows one modelling technique to be used to complement or enhance the other. Specifically, we note that the works of Chapters 3 and 5 both take advantage of this important relationship, see Section 1.3 for additional details.

1.3 Thesis Summary

Chapter 2 begins by investigating how changing patterns of social interaction brought about by new media, e.g. social networks, satellite television, Short Message Service (SMS) text messaging, etc..., might have facilitated the surprising and dramatic events of the political protests and revolutions of the 2010-2011 Arab Spring. Specifically, using both conceptual and network based justifications, we develop a simple one-dimensional ordinary differential equation model, i.e. the step visibility function (SVF) model, describing the dynamics of a political protest in a dictatorial regime that employs censorship and repression. This situation represents a collective action problem, since individuals who choose to protest unilaterally face significant reprisal from the regime, whereas individuals who protest in sufficient numbers are immune to retaliation from the regime due to the regime's finite resource constraint. We discuss the model interpretation and implications in the context of the Arab Spring revolutions, among others. In this context it is important that we select a simple mathematical model that admits a complete mathematical analysis. This allows us to completely understand all possible dynamical behaviours exhibited by the model, and hence, allows us to develop a straightforward interpretation in the context of political revolutions in dictatorial regimes. The major contribution of our simple model is that it provides a conceptual framework for understanding potential outcomes from political protests/revolution in dictatorial regimes, and how these potential outcomes are influenced by the structure of the underlying network of social interactions.

The work of Chapter 3 follows the mathematical model presented in Chapter 2. As we discussed above, the model presented in Chapter 2 is a simple compartmental model that is useful as a conceptual tool for understanding the role of new communication technology in political protests/revolutions in dictatorial regimes. However, in order to arrive at this simple mathematical model, many significant assumptions were necessary. Specifically, instead of explicitly accounting for the communication network (and how this network changes over time), the model presented in Chapter 2 makes a number of assumptions that allows us to implicitly model the communication network. Chapter 3 relaxes these assumptions by explicitly incorporating the communication network into a linear threshold agent-based model (ABM) for the dynamics of political protest in dictatorial regimes. In so doing, the ABM allows us to directly investigate the role of new communication technologies, and also to justify the assumptions made in our simple compartmental model *a posteriori*. Moreover, although the ABM of Chapter 3 does not admit a complete mathematical analysis, we note that because we carefully specify the ABM to be consistent with the simple compartmental model previously developed in Chapter 2, we are able to use the analysis provided in Chapter 2 as a guide for the computational analysis performed for our

ABM. Finally, we use insights gained from the ABM to refine the original simple model of Chapter 2 to enhance its ability to approximate the aggregate behaviour of the ABM. Small scale numerical simulations of the ABM on a physical contact network, and on a subset of the Facebook network, support the thesis that new communication technologies facilitated the protests and revolutions of the Arab Spring. The major contributions of this section are that it (a) confirms *a posteriori* the assumptions made to develop the simple model of Chapter 2, (b) provides evidence for the hypothesis that the Arab Spring was facilitated by new communication technologies, and (c) outlines methodologically how simple low dimensional compartmental models and high dimensional ABMs can be developed in a consistent fashion and how they can be used together to complement each others' weaknesses.

Whereas Chapters 2-3 investigate social effects in the context of political movements, Chapters 4-5 investigate social effects in epidemiological contexts, i.e. in the context of non-communicable diseases. Chapter 4 models the dynamics of smoking prevalence as a binary choice conformity problem where individuals choose to belong to one of two social groups (i.e. smokers or non-smokers). Specifically, Chapter 4 uses a simple compartmental model for the dynamics of smoking prevalence over the time span of approximately one century. As with Chapter 2, the simple compartmental model that we develop admits a complete mathematical analysis, as well as a straightforward interpretation. More importantly, however, given the relative sparsity of the data on smoking prevalence, the selection of a simple mathematical model also allows us to minimize the likelihood that our model suffers from problems of overfitting. The main contributions of this chapter are threefold. First, since antecedent mathematical models for binary choice conformity problems [1, 2] are only capable of modelling monotonic changes in prevalence, and since smoking prevalence peaked in the mid-twentieth century, we modify these previous approaches by including a term modelling the known health risks of smoking. This quantity, which increases over the time frame of our model and which decreases the utility that individuals receive directly from the act of smoking, allows our model to capture the rise and fall of smoking prevalence over the past hundred years. Second, we use data on tobacco consumption to estimate smoking prevalence in seven advanced economies over the past century. In each of these countries, therefore, we have established the longest reported time series for smoking prevalence. Finally, by calibrating our model to estimated smoking prevalence data, and comparing our fitted parameters to a previously established measure of individualism (Hofstede's IDV), we provide evidence in support of the hypothesis that a society's culture (in this case its level of individualism) can have a quantifiable influence on the dynamics of a social spreading process.

Whereas Chapters 2-4 investigate binary choice problems, i.e. situations where in-

dividuals choose between one of two alternatives (e.g. to protest or not to protest, to smoke or not to smoke), Chapter 5 modifies our approach for the case where individuals choose between a continuum of choices. Specifically, Chapter 5 develops a mathematical model for the dynamics of the body mass index (BMI)[4]. As with Chapter 3, Chapter 5 proposes models at both the individual and population-levels. More precisely, we specify an agent-based model for the dynamics of BMIs at the individual level. This agent-based model, which consists of a system of coupled Langevin equations, has known implications at the population-level. Specifically, the system of Langevin equations that we develop is related to a Fokker-Planck partial differential equation that describes the evolution of the probability density function for BMIs, i.e. the evolution of the distribution of BMIs in the population. Using three independently collected data sets we show that our model is capable of reproducing empirical BMI data at both the individual and population levels. We make two main contributions at both the individual and population-levels. At the population-level our model provides a mechanistic explanation for the right-skewness of the BMI distribution, and is also better at reproducing the distribution of BMIs in the population than two other probability distribution functions (i.e. the skew-normal and log-normal distributions) commonly used to describe right-skewed data. At the individual level, our model is able to reproduce empirical data for the average and standard deviation of year-over-year change in individuals' BMI, and also provides evidence in support of the hypothesis that social influences play an important role in the dynamics of individuals' BMI. Finally, we note that one of the three BMI data sets mentioned above, which is comprised of anonymized medical records from the Northwestern Medical network of hospitals and clinics, was compiled specifically as part of this research project. The compilation of this data set represents a major contribution, since it represents the largest known collection of BMI panel data, i.e. individuals' year-over-year change in BMI can be calculated for 329,543 distinct individuals resulting in over 1,017,518 data points spanning a 17 year period.

Publications resulting from the work performed in this thesis are summarized in the following section. Following Chapters 2-5, the conclusion of this thesis is given in Chapter 6.

[4]The body mass index is defined as the ratio of body mass (in kilograms) to height (in meters) squared and is a commonly used indicator for body composition, i.e. the higher an individual's BMI the greater the percentage of fat in overall body composition.

1.3.1 Publications Related to this Thesis

- J.C. Lang and H. De Sterck. The Arab Spring: A simple compartmental model for the dynamics of a revolution. *Mathematical Social Sciences*, 69:12-21, 2014.
  - Chapter 2 is based on this paper [109].
- J.C. Lang and H. De Sterck. A hierarchy of linear threshold models for the spread of political revolutions on social networks. *Journal of Complex Networks*, 2016. DOI: 10.1093/comnet/cnv030.
  - Chapter 3 is based on this paper [110].
- J.C. Lang, D.M. Abrams, and H. De Sterck. The influence of societal individualism on a century of tobacco use: modelling the prevalence of smoking. *BMC Public Health*, 15(1):1-13, 2015.
  - Chapter 4 is based on this paper [108].

Chapter 2

The Arab Spring: A Simple Compartmental Model for the Dynamics of a Revolution

This chapter is based on research published in *Mathematical Social Sciences* [109].

2.1 Introduction

> "After decades of political stagnation... new winds of hope were felt in the Middle East, accompanied by a new catchword making the rounds in the American media, 'Arab Spring'... The age of the old patriarchs, it appeared, was nearing its end. And the new media - satellite television, mobile phones, the Internet - were often regarded as having precipitated this development by undermining governments' hegemonic control over the flow of information."

When A. Hofheinz wrote these words about the Arab Spring he was referring to modest advancements being made in democracy and political liberalization in a handful of Middle Eastern countries in 2005 [85]. He did not foresee the events sparked by Mohamed Bouazizi's self-immolation in a small Tunisian city on December 17, 2010 that ultimately led to the 2010-2011 Arab Spring revolutions[1]. Nevertheless, his analysis of new media and their impact on Arab society is eerily prescient, especially considering that in 2005 social media

[1]Henceforth, when we use the term Arab Spring when referring to the 2010-2011 Arab Spring.

was either in its infancy or completely non-existent. Indeed, Facebook was launched in 2004 and was still an invitation-only service in 2005, YouTube was founded in early 2005, and Twitter was not founded until the spring of 2006 [21].

The predominant (but not universal [7]) view today is that the Internet and social media played a critical role in the Arab Spring [5, 89, 92, 100, 116, 144, 162, 165, 192]. Although some rigorous work has been done attempting to determine a link between social media and protests using, for example, Twitter data [89, 116], most of these opinions are based on anecdotal evidence and *ad hoc* reasoning. Thus, the goal of this chapter is twofold. First, we introduce a simple compartmental model for the dynamics of a revolution in dictatorial regimes that employ censorship and police repression. A defining property of the model is the use of visibility and policing terms that feature rapid transitions as a function of the size of the revolution, for which we provide conceptual and network-based mathematical justifications. The complete mathematical classification of the dynamical behaviour of the model leads to a division in parameter space that is interpreted naturally in terms of stability of the regime (stable police state, meta-stable police state, unstable police state, and failed state). We show that these dynamical properties of the model are generic for a broad class of visibility and policing functions that feature rapid transitions. Second, we investigate how the model can be applied to the Arab Spring revolutions in Tunisia and Egypt, taking into account the influence of the Internet and new media on the visibility of the revolution and the ensuing reduced efficacy of censorship. Within the model this leads to significant, discontinuous changes in regime stability, which greatly increase the probability of realized revolutions. These properties of the model inform possible answers to questions on causes and timing of the Arab Spring revolutions, and the role of the Internet and new media. The broader relevance of the model classification is also investigated by applying it to the current political situation in some other countries with regimes that employ censorship and police repression.

It should be noted that models of opinion/norm formation [36], conflict [9, 43, 105, 151], and revolution [71, 107] already exist. However, these models either do not apply specifically to the peaceful political revolutions of the Arab Spring or are highly complex. Although complex models may in principle be able to offer a more complete description, they also have limitations such as requiring the calibration of a large number of parameters. This makes complex models analytically intractable, difficult to interpret, and computationally expensive to simulate. In Epstein's categorization of reasons to model [51], ours is to illustrate the core dynamics of the balance between, on one hand, the growth of a revolutionary movement and the influence of censorship and information flow on this growth, and on the other hand, the suppression of the revolution by police force. As a simple conceptual model, our model follows the tradition of well-known differential equation models

in conflict analysis like the Richardson arms race model [151]. Thus, we emphasize that the main goal of this chapter is to develop a simple model that is nevertheless able to capture essential features of political revolutions in dictatorial regimes that employ censorship and police repression.

The remainder of this chapter is divided into five additional sections. We specify and justify our model in Section 2.2, and we provide a mathematical analysis of our model in Section 2.3. Section 2.4 specifies and analyzes an extension to our model that relaxes some of the simplifying assumptions made in the basic form of the model, confirming that the dynamical properties of the basic model are generic for a broader class of models. We investigate how our model can be applied to the Arab Spring revolutions in Tunisia and Egypt, and the current situation in some other countries with regimes that employ censorship and police repression, in Section 2.5. Finally, conclusions are given in Section 2.6.

2.2 Model Specification

2.2.1 Basic Model

We begin by specifying a simple model describing the process by which citizens engage in revolution in regimes that employ censorship and police repression. Let $r = r(t)$ be the fraction of *protesters* or *revolutionaries* in the population at time t. The model is given by a single differential equation for $r(t)$,

$$\dot{r}(t) = \frac{dr}{dt} = \underbrace{c_1 \; v(r;\alpha) \; (1-r)}_{g(r)} - \underbrace{c_2 \; p(r;\beta) \; r}_{d(r)}, \tag{2.1}$$

where $\alpha, \beta \in (0,1)$ and $c_1, c_2 > 0$ are parameters, and where the functions g, d: $[0,1] \to \mathbb{R}^+$ are called the *growth* and *decay* terms, respectively, since they model the growth and decay of the fraction of protesters.

In our model the fraction of protesters can only grow when the protest movement is sufficiently large to be visible to the general population, defying attempts at censorship by the regime. We call the proportionality constant c_1 and the parameter α the *enthusiasm* and *visibility* of the protesters, respectively. The *visibility term* $v(r;\alpha)$ is modelled as a step function which shuts off the growth term when the fraction of protesters is below the *visibility threshold* $1-\alpha$, i.e.

$$v(r;\alpha) = \mathbb{I}_{\{r>1-\alpha\}},$$

where $\mathbb{I}_{\{.\}}$ is the indicator function, i.e.

$$\mathbb{I}_{\{X\}} = \begin{cases} 1 & \text{if } X \text{ is true} \\ 0 & \text{if } X \text{ is false} \end{cases}.$$

We illustrate the visibility term $v(r;\alpha)$ and the growth term $g(r)$ in Fig. 2.1a. The *policing term* $p(r;\beta)$, is also modelled as a step function, which shuts down the decay term when the fraction of protesters is above the *police capacity* β, i.e.

$$p(r;\beta) = \mathbb{I}_{\{r<\beta\}}.$$

We illustrate the police term $p(r;\beta)$ and decay term $d(r)$ in Fig. 2.1b. We call the proportionality constant c_2 the *policing efficiency*. Finally, we observe from Eq. (2.1) that if $r = 0$ or $r = 1$ then $\dot{r} = 0$. Thus, we call $r = 0$ and $r = 1$ the equilibria of *total state control* and of the *realized revolution*, respectively.

We provide a conceptual justification for the model developed above, as well as a network-based mathematical justification for the step function, in the following section.

2.2.2 Conceptual Justification of Basic Model (Eq. (2.1))

In order to arrive at the model given in Eq. (2.1) some simplifying assumptions were made. First, our model is developed for describing rapid revolutionary transitions on a short time scale (on the order of months), and neglects demographic and other long-term effects. Second, we assume that the regime is very unpopular and that all individuals would privately like to see the regime changed. While this does not apply to all cases of political revolution (e.g., the 2011 revolutions in Libya or the current civil war in Syria), it is nevertheless a reasonable modelling approximation in many cases (e.g. the Arab Spring revolutions in Tunisia or Egypt). These two assumptions imply a constant population that can be divided into two compartments: the population participating ($r(t)$) and not participating ($1 - r(t)$) in the revolution. We note that, by the second assumption, the fraction of the population potentially willing to join the revolution at time t is $1 - r(t)$.

Dictatorial regimes are known to keep tight control on the flow of political information through state control of the media and through censorship [36, 47, 85, 92, 100, 107], for obvious reasons: if political protests are kept hidden from the general population, protest movements have little chance of growing. We model this effect through the visibility term $v(r;\alpha)$, which we assume undergoes a rapid transition from 0 to 1 that, for simplicity, we model as a step function. As soon as the fraction of protesters reaches the visibility

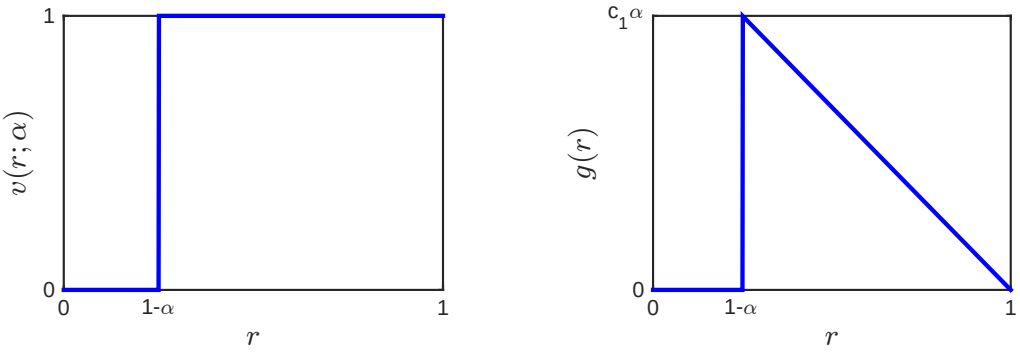

(a) Visibility and growth terms. (Left) Visibility term $v(r;\alpha)$. (Right) Growth term $g(r)$.

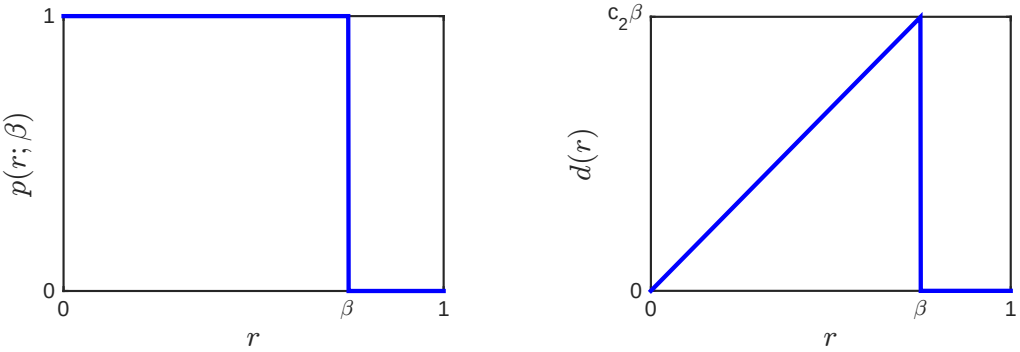

(b) Policing and decay terms. (Left) Policing term $p(r;\beta)$. (Right) Decay term $d(r)$.

Figure 2.1: Visibility $v(r;\alpha)$, growth $g(r)$, policing $p(r;\beta)$, and decay $d(r)$ terms for the basic model (Eq. (2.1)).

threshold $1-\alpha$ and is large enough to be visible to the general population, the revolution is assumed to grow proportional to $1-r$, where the time scale of growth is determined by the protesters' enthusiasm c_1.

As further motivation for the rapid transition from 0 to 1 in $v(r;\alpha)$ consider that, given the policing limitations of the regime, the decision of individuals whether or not to act is a collective action problem [107]. If individuals protest individually then the state is capable of severe retaliation, however, if individuals protest in sufficient numbers then the state loses its ability to punish. Thus, the case can be made that the most important factor for individuals deciding to join a revolution is the *perceived size* of the revolution. A network-based mathematical justification of how this would lead to rapid transitions from 0 to 1 in the visibility term of our model is provided in the following section.

We assume that the regime is capable of arresting/dispersing protesters at a rate proportional to the size of the revolution $r(t)$, provided that the number of protesters does not exceed the regime's finite police capacity β. Provided that no new protesters join the revolution ($v = 0$) and that the number of protesters does not exceed the regime's police capacity ($p = 1$), this corresponds to exponential decay in the number of protesters with the time scale determined by the policing efficiency c_2. We make the further simplifying assumption that the regime loses all ability to punish once the number of protesters exceeds the regime's police capacity. Thus, $p(r;\beta)$ takes the form of a step function in our basic model.

We note that the growth term in our model (with parameters c_1 and α) can be related to aspects of *grievance* [43], the utility of protest, and the overall emotional state of the population [138]. In contrast, parameters c_2 and β can be related to aspects of *political opportunity* [43] and *state capacity* [82].

In Section 2.4 we relax the step function transitions of $v(r;\alpha)$ and $p(r;\beta)$ to broader classes of sigmoid-type rapid transitions between 0 and 1, and show that the resulting more general models are capable of exhibiting the same dynamics as our basic model with step functions. This shows that the discontinuities in the step function transitions between 0 and 1 are not essential for obtaining the dynamics of the model and do not introduce artifacts in the dynamics, and that the dynamics are indeed generic for a broader class of models with rapid transitions between 0 and 1 in visibility and policing terms. We choose step functions in our basic model because they require the smallest number of parameters in describing rapid transitions between 0 and 1, and yet maintain the essential dynamics of more complicated models with more generic rapid transitions between 0 and 1.

Finally, due to the simplicity of our model, it is unable to capture singular events such as the self-immolation of Mohamed Bouazizi on December 17, 2010. Although this type of event could be modelled stochastically, to keep our model as simple as possible we introduce the concept of *shocks*: events external to our model which affect the fraction of revolutionaries $r(t)$ either directly, or indirectly via a change in the parameters α, β, c_1, or c_2.

2.2.3 Network-Based Justification of Basic Model (Eq. (2.1))

Section 2.2.2 introduced the decision of an individual to participate in a revolution as a collective action problem. If individuals protest individually then the state is capable of severe retaliation, however, if individuals protest in sufficient numbers then the state loses its ability to punish. It is a reasonable assumption, therefore, that each individual i in a

network has a threshold $\theta_i \in (0,1)$ which determines his or her participation. In particular, an individual is *considering joining the revolution* only if the fraction of the individual's neighbors participating in the revolution is at least θ_i [107]. We use this basic framework to justify the choice of $v(r;\alpha)$ as a rapid transition between 0 and 1 (modelled by a step function in our basic model) by considering the following derivation.

Let individuals in a population be represented by nodes V in a network $G = G(V,E)$, where edges E linking individuals represent social interactions. We define the average degree and threshold of individuals V to be ϕ and θ, respectively. Suppose that in the time interval $(t, t+\Delta t)$ an individual $i \in V$ who is considering joining the revolution joins with probability $c_1 \Delta t$. Equivalently, taking $\Delta t \to 0$, a node that is considering joining the revolution will join at the first arrival time of a Poisson process with rate c_1. In order to estimate the fraction of the average individual's neighbours that are participating in the revolution, without having to specify additional information about the structure of the underlying network, it is necessary to make a simplifying assumption. Specifically, following [104, 128, 137] we assume that the states of an individual's neighbors are uncorrelated. It now follows that the probability that an average individual considers joining the revolution is

$$\begin{aligned}
&\sum_{j=\lceil\theta\phi\rceil}^{\lfloor\phi\rfloor} \binom{\lfloor\phi\rfloor}{j} r^j (1-r)^{\lfloor\phi\rfloor - j} \\
&= \sum_{j=0}^{\lfloor\phi\rfloor - \lceil\theta\phi\rceil} \binom{\lfloor\phi\rfloor}{j} (1-r)^j r^{\lfloor\phi\rfloor - j} \\
&= \mathrm{BinCDF}(\lfloor\phi\rfloor - \lceil\theta\phi\rceil; \lfloor\phi\rfloor, 1-r), \qquad (2.2)
\end{aligned}$$

where $\mathrm{BinCDF}(x;n,p)$ is the cumulative distribution function for the binomial distribution with n trials and probability of success p evaluated at x. An approximation for the number of nodes considering joining the revolution is thus

$$(1-r)\ \mathrm{BinCDF}(\lfloor\phi\rfloor - \lceil\theta\phi\rceil; \lfloor\phi\rfloor, 1-r).$$

It follows that

$$\Delta r \approx c_1\ \Delta t\ (1-r)\ \mathrm{BinCDF}(\lfloor\phi\rfloor - \lceil\theta\phi\rceil; \lfloor\phi\rfloor, 1-r).$$

Dividing by Δt and taking the limit $\Delta t \to 0$ yields

$$\frac{dr}{dt} \approx c_1\ (1-r)\ \mathrm{BinCDF}(\lfloor\phi\rfloor - \lceil\theta\phi\rceil; \lfloor\phi\rfloor, 1-r),$$

where $\text{BinCDF}(\lfloor\phi\rfloor - \lceil\theta\phi\rceil; \lfloor\phi\rfloor, 1-r)$ corresponds to the visibility term $v(r;\alpha)$ of model (2.1).

Individuals only consider joining the revolution if it has already grown to a certain extent, but an individual's ability to determine the extent of the revolution is constrained by the number of neighbors the individual has. Specifically, the more neighbors one has the more certain one can be about the true extent of the revolution. We emphasize that determining the true extent of the revolution before deciding to act is of critical importance because if one overestimates the support for a revolution then one risks acting unilaterally. We would, therefore, expect the average threshold θ and average degree ϕ to be negatively correlated (increased ϕ corresponds to decreased θ). We illustrate the dependence of Eq. (2.2) on the parameters θ and ϕ in Fig. 2.2. All panels of Fig. 2.2 show that Eq. (2.2) is a function of r with a steep transition from 0 to 1, with the transition becoming steeper as ϕ increases, which lends justification to our choice of a rapid transition from 0 to 1 for $v(r;\alpha)$. Although we do not know the exact relationship between θ and ϕ, Fig. 2.2 illustrates how when θ and ϕ are negatively correlated Eq. (2.2) can be approximated by a steep transition from 0 to 1 that has one parameter determining where the transition occurs (corresponding to parameter α in the visibility term $v(r;\alpha)$). In our basic model we model this steep transition from 0 to 1 as a step function, while more general sigmoid-type transitions are considered in Section 2.4. Note that, in a macroscopic view, the average number of effective neighbors ϕ that individuals can interact with is increased significantly by new media, which corroborates our interpretation of $v(r;\alpha)$ as a *visibility term*.

We conclude this section with a few remarks regarding the assumption made above that the states of an individual's neighbors are uncorrelated. Given that the process we are studying spreads over social networks, which tend to have high clustering coefficients, this assumption would benefit from additional justification. The primary justification we offer for this assumption is empirical. First, we note that this assumption is applied in many models of spreading processes, for example rumour spreading processes [128] and epidemiological processes [104, 137]. These models have been shown to be good first order approximations of higher order "effective degree" or Agent Based models that do not employ assumptions about the correlations of the states of an individual's neighbors [114]. Second, we have recorded preliminary results confirming that the visibility function in a simple model that includes effects of the network (by statistically sampling real social network data sets without assuming uncorrelated neighbors) takes the form of a rapid transition of the type illustrated in Fig. 2.2 also for these real networks. Further details are given in Chapter 3.

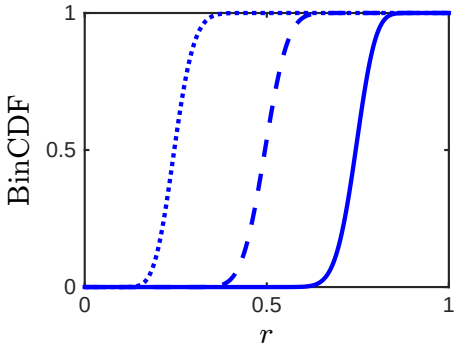

(a) Constant $\phi = 100$ with (solid line) $\theta = 0.75$, (dashed line) $\theta = 0.5$, and (dotted line) $\theta = 0.25$.

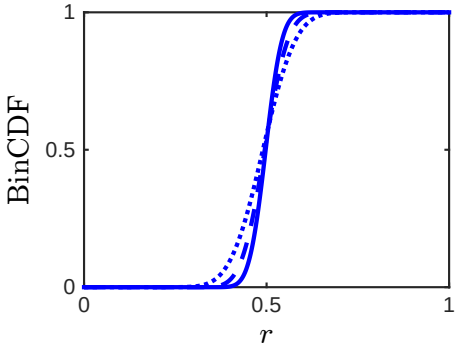

(b) Constant $\theta = 0.5$ with (solid line) $\phi = 200$, (dashed line) $\phi = 100$, and (dotted line) $\phi = 50$.

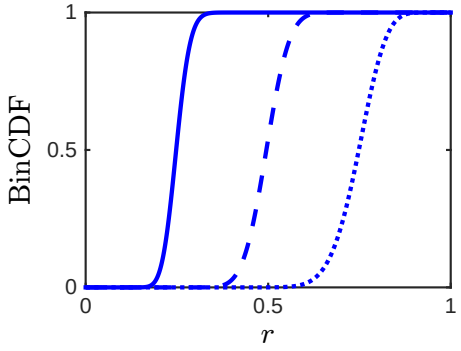

(c) Negatively correlated θ and ϕ. (Solid line) $(\theta, \phi) = (0.75, 50)$, (dashed line) $(\theta = 0.50, \phi = 100)$, and (dotted line) $(\theta, \phi) = (0.75, 50)$.

Figure 2.2: Dependence of $\text{BinCDF}(\lfloor\phi\rfloor - \lceil\theta\phi\rceil; \lfloor\phi\rfloor, 1-r)$, i.e. of the probability that an average individual considers joining the revolution (Eq. (2.2)), on average threshold θ and average degree ϕ.

2.3 Model Analysis

2.3.1 Classification of Parameter Regimes

The mathematical classification of the different types of dynamical behaviour that may occur in model (2.1) proceeds case-wise by considering parameter regions $\alpha + \beta = 1$, $\alpha + \beta < 1$, and $\alpha + \beta > 1$, which we call Regions I, II, and III, respectively, see Fig. 2.3a. See Fig. 2.3b for the phase portraits of the different regions.

Before considering Regions I, II, and III separately we begin by considering the equilibria $r = 0$ and $r = 1$. When

$$r < \min\{1 - \alpha, \beta\}$$

we have $v(r; \alpha) = 0$ and $p(r; \beta) = 1$, which implies

$$\dot{r} \leq 0 \text{ and } \dot{r} = 0 \iff r = 0.$$

Similarly, when

$$r > \max\{1 - \alpha, \beta\}$$

we have $v(r; \alpha) = 1$ and $p(r; \beta) = 0$, which implies

$$\dot{r} \geq 0 \text{ and } \dot{r} = 0 \iff r = 1.$$

It follows that $r = 0$ and $r = 1$ are locally asymptotically stable equilibria whose basins of attraction contain the intervals $[0, \min\{1 - \alpha, \beta\})$ and $(\max\{1 - \alpha, \beta\}, 1]$, respectively.

Region I: $\alpha + \beta = 1$

When $r = \beta = 1 - \alpha$ we have $v(r; \alpha) = p(r; \beta) = 0$, which implies $\dot{r} = 0$. It follows that $r = \beta = 1 - \alpha$ is an unstable equilibrium, and that $r = 0$ and $r = 1$ are locally asymptotically stable equilibria with basins of attraction $[0, \beta)$ and $(1 - \alpha, 1]$, respectively. We note that Region I, because it is one-dimensional, is unlikely to manifest itself and so we mostly disregard it in what follows.

Region II: $\alpha + \beta < 1$

As in the previous section, $r = 0$ and $r = 1$ are locally asymptotically stable equilibria with basins of attraction $[0, \beta)$ and $(1 - \alpha, 1]$, respectively. When $r \in [\beta, 1 - \alpha]$ we have $v(r; \alpha) = p(r; \beta) = 0$, which implies $\dot{r} = 0$. It follows that all $r \in (\beta, 1 - \alpha)$ are locally stable equilibria and $r \in \{\beta, 1 - \alpha\}$ are unstable equilibria.

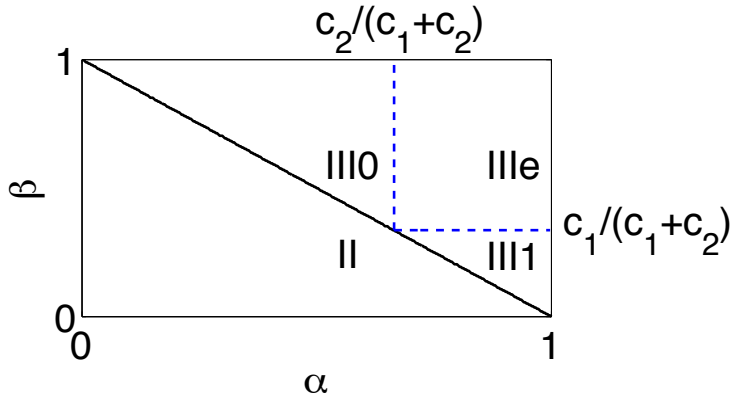

(a) Division of $\alpha - \beta$ parameter space of model (2.1) into Regions I, II, IIIe, III0, and III1. Region I is given by the solid line $\alpha + \beta = 1$.

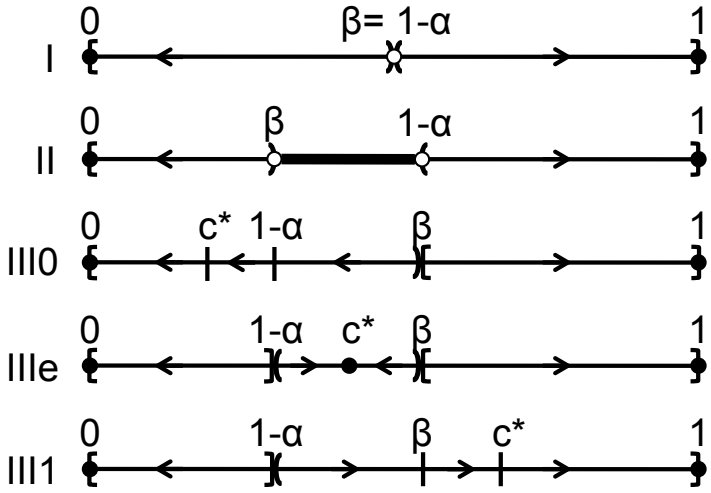

(b) Equilibria, stability and basins of attraction on the r-axis ($r \in [0, 1]$) for parameters α, β, c_1 and c_2 of model (2.1) lying in Regions I, II, III0, IIIe and III1. Closed (open) circles represent locally asymptotically stable (unstable) equilibria. Left (right) arrows indicate regions where $\dot{r} < 0$ ($\dot{r} > 0$). The thick horizontal line indicates locally stable equilibria. Square ("[" and "]") or curved ("(" and ")") brackets indicate boundaries of basins of attraction that contain or do not contain their boundary points, respectively. Short vertical lines indicate special values of r that are not equilibria or boundaries of basins of attraction.

Figure 2.3: Parameter regime and phase diagram for Regions I, II, III0, IIIe, and III1 of the basic model (Eq. (2.1)).

Region III: $\alpha + \beta > 1$

When $r = 1 - \alpha$ we have $v(r; \alpha) = 0$ and $p(r; \beta) = 1$, which implies $\dot{r} < 0$. Analogously, when $r = \beta$ we have $v(r; \alpha) = 1$ and $p(r; \beta) = 0$, which implies $\dot{r} > 0$. Thus, the locally asymptotically stable equilibria $r = 0$ and $r = 1$ have basins of attraction containing $[0, 1 - \alpha]$ and $[\beta, 1]$, respectively. Restricting our attention to the interval $r \in (1 - \alpha, \beta)$ and solving the algebraic equation $\dot{r} = 0$ gives $r = c^*$, where

$$c^* = \frac{c_1}{c_1 + c_2}.$$

We now observe that our analysis breaks down into a further three sub-cases:

- Region III0 ($c^* \leq 1 - \alpha < \beta$): If $c^* \in [0, 1 - \alpha]$ then

  $$\forall r < \beta : \dot{r} \leq c_1(1 - r) - c_2 r < 0,$$

  which implies that the interval $(1 - \alpha, \beta)$ lies in the basin of attraction of $r = 0$.

- Region III1 ($1 - \alpha < \beta \leq c^*$): If $c^* \in [\beta, 1]$ then

  $$\forall r > 1 - \alpha : \dot{r} \geq c_1(1 - r) - c_2 r > 0,$$

  which implies that the region $(1 - \alpha, \beta)$ lies in the basin of attraction of $r = 1$.

- Region IIIe ($1 - \alpha < c^* < \beta$): If $c^* \in (1 - \alpha, \beta)$ then

  $$\dot{r} = c_1(1 - r) - c_2 r \begin{cases} > 0 & \text{if } r \in (1 - \alpha, c^*) \\ = 0 & \text{if } r = c^* \\ < 0 & \text{if } r \in (c^*, \beta) \end{cases},$$

  which implies that there exists a third locally asymptotically stable equilibrium $r = c^*$ with basin of attraction $(1 - \alpha, \beta)$.

2.3.2 Analytic Solution to the Basic Model (Eq. (2.1))

Although the analysis of the previous section is more useful in the interpretation and classification of the parameter regimes of model (2.1), it is also possible to solve model (2.1)

explicitly using the technique of separation of variables. Doing so for the initial condition $r(t_0) = r_0$ yields the following. In Region I ($1-\alpha = \beta$) and Region II ($1-\alpha > \beta$) we have

$$r(t) = \begin{cases} r_0 e^{-c_2(t-t_0)} & \text{if } r_0 < \beta \\ r_0 & \text{if } r_0 \in [\beta, 1-\alpha] \\ 1-(1-r_0)e^{-c_1(t-t_0)} & \text{if } r_0 > 1-\alpha \end{cases} .$$

In Regions III0 ($c^* \leq 1-\alpha < \beta$), IIIe ($1-\alpha < c^* < \beta$), and III1 ($1-\alpha < \beta \leq c^*$) we have

$$r(t) = \begin{cases} r_0 e^{-c_2(t-t_0)} & \text{if } r_0 < 1-\alpha \\ \left[c^* + (r_0 - c^*)e^{-(c_1+c_2)(t-t_0)}\right] \mathbb{I}_{\{t<t_\alpha\}} & \\ \quad +(1-\alpha)e^{-c_2(t-t_\alpha)}\mathbb{I}_{\{t\geq t_\alpha\}} & \text{if } r_0 \in [1-\alpha, \beta] \text{ and } c^* \leq 1-\alpha < \beta \\ c^* + (r_0 - c^*)e^{-(c_1+c_2)(t-t_0)} & \text{if } r_0 \in [1-\alpha, \beta] \text{ and } 1-\alpha < c^* < \beta \\ \left[c^* + (r_0 - c^*)e^{-(c_1+c_2)(t-t_0)}\right] \mathbb{I}_{\{t<t_\beta\}} & \\ \quad + \left[1-(1-\beta)e^{-c_1(t-t_\beta)}\right] \mathbb{I}_{\{t\geq t_\beta\}} & \text{if } r_0 \in [1-\alpha, \beta] \text{ and } 1-\alpha < \beta \leq c^* \\ 1-(1-r_0)e^{-c_1(t-t_0)} & \text{if } r_0 > \beta \end{cases} ,$$

where

$$t_\alpha = t_0 - \frac{1}{c_1+c_2}\log\left(\frac{1-\alpha-c^*}{r_0-c^*}\right), \text{ and}$$
$$t_\beta = t_0 - \frac{1}{c_1+c_2}\log\left(\frac{\beta-c^*}{r_0-c^*}\right).$$

2.4 Model Extension: Sigmoidal visibility $v(r)$ and policing $p(r)$ terms

The assumption that the visibility $v(r;\alpha)$ and policing $p(r;\beta)$ terms are step functions is a strong assumption that results in discontinuities in the vector field of Eq. (2.1) at $r = 1-\alpha$ and $r = \beta$. Specifically, choosing $v(1-\alpha;\alpha) = p(\beta;\beta) = 0$ results in $r \in \{1-\alpha, \beta\}$ (a) being unstable equilibria in Regions I and II, and (b) being in the basin of attraction of $r = 0$ and $r = 1$, respectively, in Regions III0, IIIe, and III1. We acknowledge that this choice is arbitrary in nature, and that defining $v(1-\alpha;\alpha)$ and $p(\beta;\beta)$ differently would change

the vector field of Eq. (2.1) at the points $r \in \{1-\alpha, \beta\}$. Nevertheless, we conclude this section by showing that essentially the same dynamics can occur when the step functions in the visibility and policing terms are relaxed to generic sigmoid-type transitions between 0 and 1. This more general model does admit additional dynamical regimes that are not present in model (2.1). However, as we will show in Chapter 3, the dynamics of model (2.1) are sufficient to characterize the dynamics of the more general model when the general model is implemented with sigmoidal visibility functions that have been computed from two examples of real world empirical communication networks. Thus, we confirm that approximating the rapid transitions in the visibility and policing terms by step functions is justified by greatly simplifying Eq. (2.1) and its analysis, without fundamentally altering its dynamics or interpretation.

2.4.1 Specification of Sigmoidal visibility $v(r)$ and policing $p(r)$ terms

Consider the equation

$$\frac{dr}{dt} = \dot{r} = \underbrace{c_1\ (1-r)\ \nu(r)}_{\gamma(r)} - \underbrace{c_2\ r\ \rho(r)}_{\delta(r)}, \tag{2.3}$$

where $\nu(r)$ and $\rho(r)$ are the visibility and policing terms, and where $\gamma(r)$ and $\delta(r)$ are the growth and decay terms, respectively. We choose $\nu(r)$ from the family of sigmoidal functions $\mathbb{S}$ comprised of (twice piecewise differentiable) functions $s : [0,1] \to [0,1]$ satisfying

(i) $\exists s_0, s_1 \in [0,1]$ such that

- $\forall r \in (s_0, s_1) : s(r)$ is twice differentiable,
- $\forall r \leq s_0 : s(r) = 0$, and
- $\forall r \geq s_1 : s(r) = 1$,

(ii) $\forall r \in (s_0, s_1) : s'(r) > 0$,

(iii) $\exists \xi \in (s_0, s_1)$ such that

- $\forall r \in [s_0, \xi] : s''(r) \geq 0$ and
- $\forall r \in [\xi, s_1] : s''(r) \leq 0$,

(iv) $\forall r \in (s_0, s_1) : s''(r)s(r) \leq [s'(r)]^2$.

Analogously, we choose $\rho(r)$ so that $\rho(1-r) \in \mathbb{S}$. The functions $s(r)$ in $\mathbb{S}$ follow a sigmoidal profile of the type illustrated in Fig. 2.4. Specifically, for $r \in (s_0, s_1)$ $s(r)$ is twice continuously differentiable, strictly increasing, and changes concavity at most once, as specified by conditions (i)-(iii). Condition (iv) is imposed to guarantee that the growth and decay terms of models (2.1) and (2.3) are qualitatively similar in the sense that they are all single peaked[2], see the right panels of Figs. 2.1a and 2.1b, and Fig. 2.5. In the following section, Proposition 2.4.1 shows that the growth term $\gamma(r)$ and decay term $\delta(r)$ of model (2.3) are single-peaked if conditions (i)-(iv) are satisfied. $\mathbb{S}$ describes a fairly broad class of functions $s(r)$ that transition monotonously from 0 to 1 in a way that $\gamma(r) = c_1\,(1-r)\,s(r)$ is single-peaked. These functions can be continuous, e.g. the linear transition from $(s_0, 0)$ to $(s_1, 1)$

$$s(r) = \begin{cases} 0 & \text{if } x \leq s_0 \\ \frac{r-s_0}{s_1-s_0} & \text{if } x \in (s_0, s_1) \\ 1 & \text{otherwise} \end{cases},$$

they can have continuous first derivatives, e.g. the truncated, scaled, and translated $\sin(\cdot)$ function

$$s(r) = \begin{cases} 0 & \text{if } x \leq s_0 \\ \frac{1}{2}\sin\left[\frac{\pi}{2}\left(\frac{2r-s_0-s_1}{s_1-s_0}\right)\right] + \frac{1}{2} & \text{if } x \in (s_0, s_1) \\ 1 & \text{otherwise} \end{cases},$$

or they can have arbitrarily many derivatives exist, e.g. the scaled and translated $\tanh(\cdot)$ function

$$s(r) = \begin{cases} 0 & \text{if } x \leq s_0 \\ \frac{1}{2}\tanh\left[\log(r-s_0) - \log(s_1-r)\right] + \frac{1}{2} & \text{if } x \in (s_0, s_1) \\ 1 & \text{otherwise} \end{cases}.$$

By construction, each of these examples satisfy conditions (i)-(iii). For detailed calculations confirming that these functions satisfy condition (iv) see Appendix B.1.

[2] A function is single peaked if it has a unique global maximum and no other local maxima.

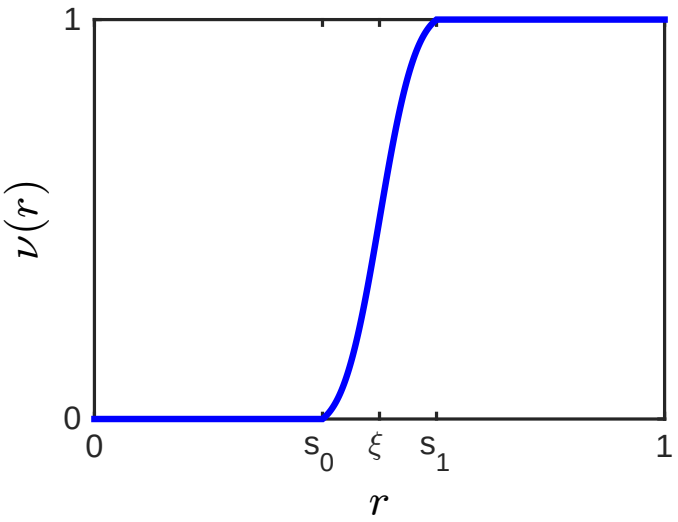

Figure 2.4: Sigmoidal visibility function $\nu(r) \in \mathbb{S}$ with a fast transition from 0 to 1 for model extension (Eq. (2.3)).

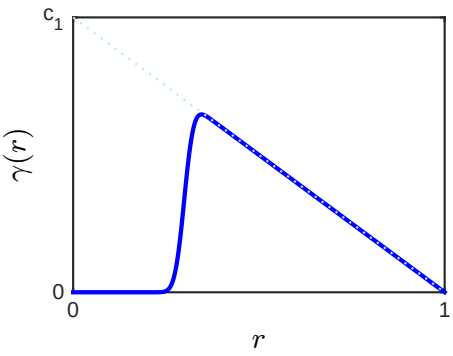

(a) (Solid line) Growth term $\gamma(r)$. (Dotted line) $\gamma(r) = c_1(1-r) + o(1)$, as $r \to 1^-$.

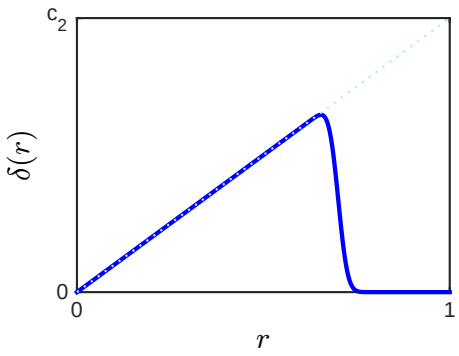

(b) (Solid line) Decay term $\delta(r)$. (Dotted line) $\delta(r) = c_2 r + o(1)$, as $r \to 0^+$.

Figure 2.5: Growth ($\gamma(r)$) and decay ($\delta(r)$) terms for model extension (Eq. (2.3)).

2.4.2 Analysis of Model Extension (Eq. (2.3))

We begin by considering the following proposition.

Proposition 2.4.1 *Let $s \in \mathbb{S}$. Then the function $\gamma(r) = c_1(1-r)s(r)$ is single peaked.*

Proof: Let $V = (s_0, s_1)$. Observe that $\forall r \in V : \gamma'(r) \geq 0$ if and only if

$$\frac{s'(r)}{s(r)} \geq \frac{1}{1-r}.$$

Now, observe that by property (iv), $\forall r \in V$:

$$\frac{d}{dr}\left[\frac{s'(r)}{s(r)}\right] = \frac{s''(r)s(r) - [s'(r)]^2}{s^2(r)} \leq 0.$$

We now have that for $r \in V$ the function $s'(r)/s(r)$ is monotonically decreasing and the function $(1-r)^{-1}$ is strictly monotonically increasing. Observe that

$$\lim_{r \to s_0^+} \log(s(r_0)) = -\infty \implies \lim_{r \to s_0^+} \frac{d}{dr}[\log(s(r))] = +\infty.$$

Since $\frac{d}{dr}[\log(s(r))] = \frac{s'(r)}{s(r)}$, it follows that

$$\lim_{r \to s_0^+} \frac{s'(r)}{s(r)} = \infty > \lim_{r \to s_0^+} \frac{1}{1-r} = \frac{1}{1-s_0}.$$

If

$$\lim_{r \to s_1^-} \frac{1}{1-r} > \lim_{r \to s_1^-} \frac{s'(r)}{s(r)}$$

then there exists a unique point $r^* \in V$ where curves $s'(r)/s(r)$ and $(1-r)^{-1}$ intersect such that $\gamma'(r^*) = 0$, at which $\gamma(r)$ achieves its global maximum. Otherwise, $\forall r \in V : \gamma'(r) > 0$ and $\gamma(r)$ achieves its global maximum at $r = s_1$.

□

Proposition 2.4.1 allows us to roughly sketch the growth and decay terms, see Fig. 2.5. The number of intersections between the growth and decay functions determines the number of equilibria in model (2.3). In what follows, we explain how the dynamics of model (2.3) with generic sigmoidal visibility and policing functions closely follows the corresponding dynamics of model (2.1) in Region II (an open interval of equilibria $\subset (0,1)$), Regions III0 and III1 (one equilibrium $\in (0,1)$), and Region IIIe (three equilibria $\in (0,1)$), thus establishing the equivalence in terms of dynamic behaviour of models (2.1) (with step functions) and (2.3) (with sigmoidal functions) for the specific parameter Regimes II, III0, IIIe, and III1. We summarize the phase portraits of Regions II, III0, IIIe, and III1 for model (2.3) in Fig. 2.6. In the following chapter we compute the sigmoidal visibility function for two empirical communication networks and confirm that our analysis of Regions II, III0, IIIe, and III1 is sufficient to describe the possible dynamical behaviours that we observe. However, we note that for arbitrary visibility and policing functions the dynamical behaviours of model (2.3) are not necessarily limited to Regions II, III0, IIIe, and III1.

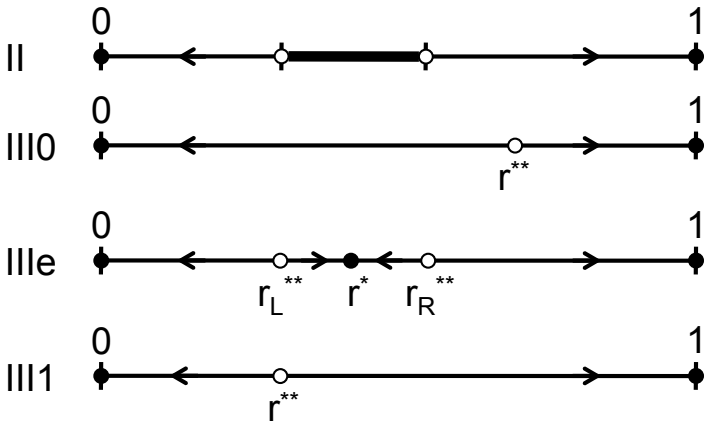

Figure 2.6: Equilibria, stability and basins of attraction on the r-axis ($r \in [0,1]$) for model extension (Eq. (2.3)). Closed (open) circles represent locally asymptotically stable (unstable) equilibria. Left (right) arrows indicate regions where $\dot{r} < 0$ ($\dot{r} > 0$). We note that although these behaviours are sufficient to describe the possible dynamics observed when sigmoidal visibility functions are computed from two empirical communication networks (see Chapter 3), the dynamics of model (2.3) are not necessarily limited to these behaviours.

Region II

Suppose that the support of $\gamma(r)$ and $\delta(r)$ have no points in common. In this situation we have (a) $r = 0$ and $r = 1$ are locally asymptotically stable equilibria, (b) there exists an open interval $I \subset (0, 1)$ of locally stable equilibria, and (c) the infimum and supremum of I are unstable equilibria. This situation is analogous to Region II of model (2.1).

Regions III0 and III1

Consider the situations depicted in Fig. 2.7. In both panels, $r = 0$ and $r = 1$ are locally asymptotically stable equilibria whose basins of attraction are separated by an unstable equilibrium point $r = r^{**}$. Figure 2.7a depicts the situation where $\max_r \gamma(r) < \max_r \delta(r)$, $\gamma'(r^{**}) < 0$, and $\delta'(r^{**}) < 0$. Because the decay term is dominant, we consider this situation to be analogous to Region III0 of model (2.1). Similarly, we consider the Fig. 2.7b to depict a situation analogous to Region III1 of model (2.1).

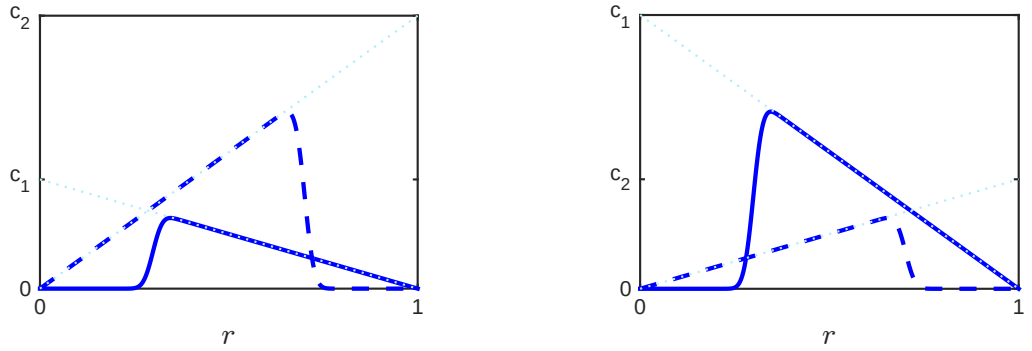

(a) Region III0: $\max_r \gamma(r) < \max_r \delta(r)$, $\gamma'(r^{**}) < 0$, and $\delta'(r^{**}) < 0$.

(b) Region III1: $\max_r \gamma(r) > \max_r \delta(r)$, $\gamma'(r^{**}) > 0$, and $\delta'(r^{**}) > 0$.

Figure 2.7: Growth ($\gamma(r)$ - solid line) and decay ($\delta(r)$ - dashed line) terms for model extension (Eq. (2.3)) with one interior intersection (Regions III0 and III1). (Dotted line) $\gamma(r) = c_1(1 - r) + o(1)$, as $r \to 1^-$ and $\delta(r) = c_2 r + o(1)$, as $r \to 0^+$.

Region IIIe

Consider the situation depicted in Fig. 2.8 in which there are three locally asymptotically stable equilibria $r = 0$, $r = r^*$ and $r = 1$ $(0 < r^* < 1)$ whose basins of attraction are

separated by two unstable equilibria $r = r_L^{**}$ and $r = r_R^{**}$ ($r_L^{**} < r_R^{**}$). This situation is analogous to Region IIIe of model (2.1).

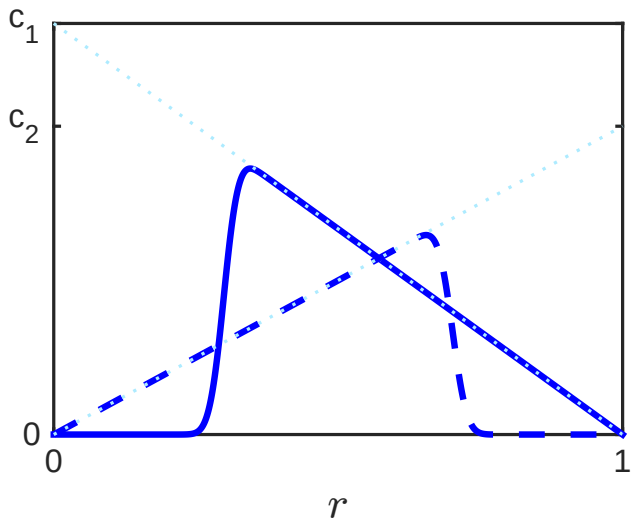

Figure 2.8: Growth ($\gamma(r)$ - solid line) and decay ($\delta(r)$ - dashed line) terms for model extension (Eq. (2.3)) with three interior intersections (Region IIIe). (Dotted line) $\gamma(r) = c_1(1-r) + o(1)$, as $r \to 1^-$ and $\delta(r) = c_2 r + o(1)$, as $r \to 0^+$.

Note that, corresponding to Region I of model (2.1), there are also special limiting cases of parameter choices for model (2.3) in which, for example, some of the equilibria in Fig. 2.6 may coincide. We choose to neglect these cases in our analysis due to the unlikelihood of them being manifested (generic small perturbations of $\gamma(r)$, $\delta(r)$, c_1, or c_2 would take us away from such a special case). Alternativelty, model (2.3) also admits additional behaviours that are not relevant in the context of political revolutions. For example, when the initial slope of the policing function is less than the initial slope of the visibility function then the equilibrium $r = 0$ becomes unstable. From a practical perspective it would not be possible to observe a state in this parameter regime at the equilibrium of total state control ($r = 0$), since any arbitrarily small shock would cause the fraction of revolutionaries $r(t)$ to grow away from $r = 0$. This dynamical regime is, therefore, clearly not relevant in the context of the Arab Spring, since each of the political revolutions and protests of the Arab Spring occured in regimes that were initially in the equilibrium of totatl state control ($r = 0$). These additional behaviours are also neglected from our analysis.

2.4.3 Remarks: Relationship between basic model (Eq. (2.1)) and model extension (Eq. (2.3))

Figures 2.3b and 2.6 show that models (2.1) and (2.3) have equivalent phase portraits for Regions II, III0, IIIe, and III1. Moreover, as we shall see in Chapter 3, the dynamical regimes identified in model (2.1), and that we have shown to be present in model (2.3), are sufficient to characterize the behaviours observed when we implement model (2.3) with sigmoidal visibility functions computed from two real world empirical communication networks. This means that the simplifying assumption of model (2.1) representing the steep increase in the visibility and policing terms by a step function (which is desirable because it gives a model with fewer parameters) is not limiting in the sense that it describes the same dynamical behaviour that is observed in a more complicated model in which the steep increases are modelled by generic sigmoidal functions, at least when the sigmoidal visibility functions are computed from empirical network data as in Chapter 3. This justifies the choice of step functions in model (2.1), since it leads to the simplest model that captures the relevant dynamics. In addition, it indicates that the dynamics we identified for model (2.1) will also occur in more complicated models for the dynamics of revolutions of type (2.3) that feature visibility and policing terms that change rapidly between 0 and 1 in a sigmoidal fashion.

2.5 Model Interpretation

In this section we first provide an interpretation of the classification of parameter regions, see Fig. 2.3a, in terms of political regime types and their stability and potential for revolutionary events, see Fig. 2.3b. We then investigate the application of the model to the Arab Spring revolutions, discussing the Arab Spring context and events, and societal factors relevant for the Arab Spring that have been identified in the political science literature. Finally, we discuss applying the classification resulting from the model to the current political situation in some other countries with regimes that employ censorship and police repression.

2.5.1 Interpretation and Classification of Parameter Regimes

We first provide an interpretation of the parameter regions of the model, previously summarized in Fig. 2.3, in the context of dictatorial regimes and their stability.

States with parameters in Region II have uncountably many stable equilibria between β and $1-\alpha$, which occur because the police capacity β of the regime is too low to clear the protesters and the visibility α is too low to attract new protesters, thus preventing any one group from taking control, see Region II in Fig. 2.3b. We therefore interpret Region II to correspond to *failed states*. In the context of our model, we investigate the application of the failed state parameter region to the case of Somalia, in Section 2.5.3.

Regions III0, IIIe, and III1 differ only in the stability of the interval $(1-\alpha, \beta)$. For Region III0 the interval $(1-\alpha, \beta)$ lies in the basin of attraction of total state control $r = 0$. Because of the contribution of $(1-\alpha, \beta)$ to the stability of the regime, we refer to Region III0 as a *stable police state*. Examples of states that, in the context of our model, are consistent with the parameter regime of Region III0 may include Tunisia and Egypt prior to 2010, see Section 2.5.2, and Iran in 2009, see Section 2.5.3. Analogous to the case of Region III0, we refer to RegionIII1 as an *unstable police state*, since $(1-\alpha, \beta)$ lies in the basin of attraction of the realized revolution $r = 1$. Region IIIe introduces an intermediate locally asymptotically stable equilibrium $r = c^*$, which lies between the equilibria of total state control $r = 0$ and of the realized revolution $r = 1$, and which has the interval $(1-\alpha, \beta)$ as its basin of attraction. We refer to $r = c^*$ as the equilibrium of *civil unrest* and to Region IIIe as a *meta-stable police state*. We hypothesize that Egypt and Tunisia transitioned to Region IIIe, and possibly to Region III1, in Section 2.5.2. Section 2.5.3 provides an additional example of a country, China, that in the context of our model shows characteristics of countries that would be classified in Region IIIe.

The examples given above of countries that could potentially be classified according to the above interpretation are discussed in detail in Sections 2.5.2 and 2.5.3, and are summarized conceptually in Fig. 2.9. We note that as the situation in a particular country evolves the regime may pass from one parameter region to another. In the context of the Arab Spring in Tunisia and Egypt, we discuss how adoption of new media may affect the parameter region of a country and how moving from one parameter region to another affects the likelihood of a realized revolution in Section 2.5.2.

2.5.2 Application: Arab Spring in Tunisia and Egypt

We now investigate the application of the model to the Arab Spring revolutions in Tunisia and Egypt. We first provide a timeline of the main events in the revolutions, followed by discussions on the effects of new media on the model parameters, and the effects of the model parameters on model stability. We then investigate how the model can be related to the events of the Arab Spring revolutions in Tunisia and Egypt.

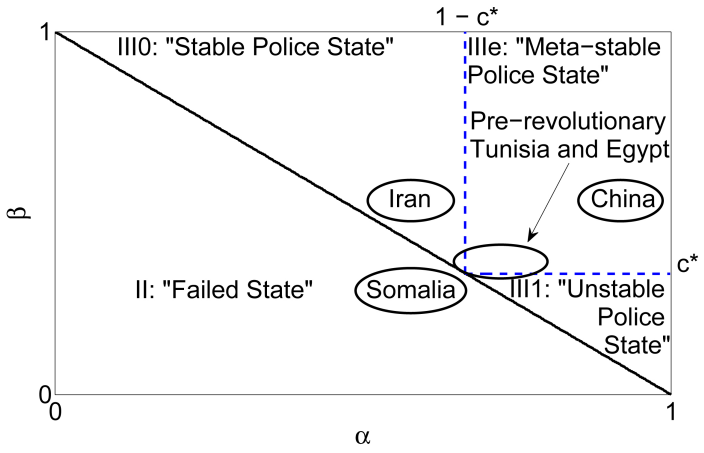

Figure 2.9: Division of $\alpha - \beta$ parameter space of basic model (Eq. (2.1)) into Regions I, II, IIIe, III0, and III1, and conceptual summary of the case studies of Tunisia, Egypt, Iran, China, and Somalia.

Timeline of Arab Spring in Tunisia and Egypt

To aid in the interpretation of our model in the cases of the Tunisian and Egyptian Arab Spring we provide the following rough timeline of major events [19, 47, 89, 92, 144, 152, 157].

- Dec. 17, 2010: Mohamed Bouazizi self-immolates in the Tunisian city of Sidi Bouzid.
- Dec. 18, 2010: Protests erupt in Sidi Bouzid. Protesters begin recording and uploading videos of the protests and police response to the Internet.
- Dec. 27 - 28, 2010: Protests break out in the Tunisian capital, Tunis. Tunisian President Ben Ali denounces protests in a televised address.
- Jan. 5, 2011: Mohamed Bouazizi dies from burn injuries.
- Jan. 14 - 15, 2011: Ben Ali resigns and flees to Saudi Arabia. An interim government is formed.
- Jan. 25, 2011: The "Day of Protest" in Tahrir Square, Cairo, is the first major Arab Spring protest in Egypt.
- Jan. 26, 2011: Egyptian police clear Tahrir Square.

- Jan. 28, 2011: Protesters occupy Tahrir Square, Egyptian President Mubarak addresses the nation, and major Internet disruptions begin.
- Feb. 1, 2011: US President Obama withdraws support for the Mubarak regime, the army refuses to act against protesters, and major Internet disruptions end.
- Feb. 2, 2011: State vandals and thugs attack protesters in Tahrir Square. Army officers intervene on behalf of protesters.
- Feb. 11, 2011: President Mubarak resigns.

Effects of New Media on Protesters' Enthusiasm c_1 and Visibility α

As outlined in Section 2.1, our development of model (2.1) was motivated by our interest in providing a basic dynamic model for Arab Spring revolutions and studying the effect of new media on censorship and the stability of dictatorial regimes as in the Arab Spring. We propose that the Internet, social media, satellite TV, and cell phone communication technologies may have empowered protesters by enhancing their

(1) capacity for organization and coordination [15, 92, 144],

(2) ability to assess the current public support for the revolution [5, 92, 100, 144, 192], and

(3) awareness of the nature and severity of government repression [92, 159].

(1) The decision of whether or not to protest is a coordination problem [107], the realization of which led activists to use the Internet to coordinate protests in Tunisia and Egypt [92, 100, 144, 192]. For example, technologies such as SMS and Twitter messaging were used between co-revolutionaries to communicate which streets were the most/least obstructed by security forces [92]. This enhanced the speed with which revolutionaries mobilized, and in the context of our model this corresponds to an increase in protesters' enthusiasm c_1. Social media and the Internet also contributed to the relatively leaderless way in which the Arab Spring revolutions developed. Compared to revolutions with a more hierarchical leadership structure, a leaderless revolution is difficult if not impossible to disrupt by targeting only a handful of individuals [15, 92]. This increased resilience also corresponds to an increase in protesters' enthusiasm c_1 in our model.

(2) Dictatorial regimes attempt to control protests through censorship by lowering visibility α in order to ensure that protests remain virtually invisible to the general population.

The Internet, social media, satellite TV, and cell phones all work towards increasing visibility α by disrupting the regime's monopoly on the distribution of information. For example, in Tunisia the Internet and social media created a virtual space where Tunisians could express their true opinions with minimal censorial oversight or fear of reprisal [92, 100, 144]. In the microscopic network model of Section 2.2.3, an increase in the number of links (average degree in the graph) along with a decrease in the threshold for action result in a shift of the sharp increase in participation toward smaller fractions of the population, see Fig. 2.2 (bottom panel). Cell phones and social media sites vastly increased the speed with which information travelled, allowing Tunisians - and the entire world - to follow the revolution with unprecedented detail and speed [92, 116, 192]. Satellite TV further enhanced visibility α by corroborating and then rebroadcasting stories relating to the size of the revolution and the regime's brutal response [5].

(3) Awareness of the brutality and severity of the government's reaction may increase both protesters' enthusiasm c_1 and visibility α. An increase in protesters' enthusiasm c_1 may be a direct result of increasing resentment of the regime. In contrast, the increase of visibility α is likely to be indirect. Specifically, otherwise apolitical individuals may be induced to join the revolution [92, 159], presumably by lowering their personal thresholds for participation.

Sensitivity of Model (2.1) to Protesters' Enthusiasm c_1 and Visibility α

The Internet and social media had at best modest penetration in countries of the Arab Spring prior to the self-immolation of Mohamed Bouazizi. Indeed, approximately 25% of Tunisians and 10% of Egyptians had used the Internet at least once prior to the Arab Spring [89]. Given this fact, it is reasonable to ask how a modest level of adoption of new technologies might have a significant impact on the outcome of a revolution. Recall that in Section 2.2.2 we argued that keeping our model simple requires introducing the concept of external shocks: events external to our model which affect the fraction of revolutionaries $r(t)$ either directly, or indirectly via a change in the parameters α, β, c_1, or c_2. An analysis of how changes in parameters may affect the outcome of a revolution must take into account these shocks.

A small increase in visibility α or protesters' enthusiasm c_1 can move parameters from the stable (Region III0) to the meta-stable (Region IIIe) police state, see Fig. 2.9. For parameters in Region III0 a shock $\Delta r > \beta$ is required to pass from total state control $r = 0$ to the basin of attraction of the realized revolution $r = 1$, see Fig. 2.10 (top line). In contrast, for parameters in Region IIIe passing from total state control $r = 0$ to the basin of attraction of the realized revolution $r = 1$ can result from two smaller shocks $\Delta r_1 > 1 - \alpha$

and $\Delta r_2 > \beta - c^*$, where $\Delta r_1 + \Delta r_2$ may be significantly smaller than β, see Fig. 2.10 (bottom line). Note that for parameters in Region IIIe the lower bounds for sufficiently strong Δr_1 and Δr_2 decrease with increasing visibility α and protesters' enthusiasm c_1, respectively. If shocks occur distributed according to some probability distribution, then it is reasonable to assume that shocks of sufficient magnitude to mobilize large fractions of the population lie in the tail of this distribution. For many reasonable probability distributions, halving the size of shock necessary to trigger a revolution more (and potentially much more) than doubles the likelihood of a revolution occurring in any given amount of time. Thus, small increases in visibility α or protesters' enthusiasm c_1 resulting from modest Internet penetration or social media usage can have a large impact on the expected amount of time one has to wait until a revolution is triggered.

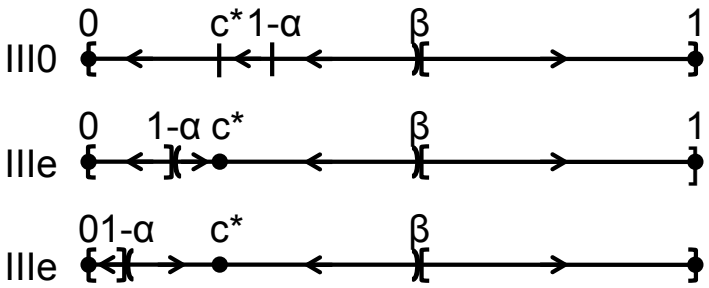

Figure 2.10: Sensitivity of basic model (Eq. (2.1)) to visibility parameter α. Closed circles represent locally asymptotically stable equilibria. Left (right) arrows indicate regions where $\dot{r} < 0$ ($\dot{r} > 0$). Square ("[" and "]") or curved ("(" and ")") brackets indicate boundaries of basins of attraction that contain or do not contain their boundary points, respectively. Short vertical lines indicate special values of r that are not equilibria or boundaries of basins of attraction.

The sensitivity of the final outcome of a revolution to the protesters' enthusiasm c_1 and visibility α also provides us with a potential explanation for why the revolutions of the Arab Spring, and indeed revolutions in general, come as such a surprise to so many regimes and political observers [5, 89, 92, 107]. Increasing either protesters' enthusiasm c_1 or visibility α eventually decreases the size of the basin of attraction of total state control $r = 0$, and does so in a discontinuous fashion. (See, e.g., Fig. 2.10: when α increases such that, at first, c^* is just below $1 - \alpha$, and then becomes slightly larger than $1 - \alpha$, the basin of attraction of $r = 0$ changes from $[0, \beta)$ to $[0, 1 - \alpha)$ in a discontinuous fashion.) This undermines the regime by decreasing the size of shock necessary to trigger a revolution. However, since (a) $r = 0$ always remains locally asymptotically stable, (b) large shocks

are exceedingly rare, and (c) determining the precise extent to which a regime controls the society through media control and policing (corresponding to determining accurate values for the parameters in a model like ours) is very difficult [82], the exact size of shock necessary to trigger a revolution is impossible to determine until such a shock occurs. It follows that for someone observing a regime before and after the adoption of social media there would be few, if any, outward signs of instability: the regime appears stable until it isn't.

Dynamics of Arab Spring Revolutions in Tunisia and Egypt

Guided by the discussion in the previous two sections, it is interesting to consider hypotheses for the Tunisian and Egyptian Arab Spring in the context of our model. It is a reasonable hypothesis that Tunisia and Egypt were, in the context of our model, in parameter Region III0 (stable police state) for a long time in the years before 2010 (Ben Ali was in power in Tunisia for 24 years (1987-2011) and Mubarak was in power in Egypt for 30 years (1981-2011)). However, the realized revolutions of 2010-2011 appear to indicate that Tunisia and Egypt had evolved to significantly less stable regimes (Region IIIe or III1) by 2010, see Fig. 2.9. Once in Region IIIe or III0, we observe that there are many potential candidates for shocks, including but not limited to those listed at the beginning of this section in the Timeline of Arab Spring in Tunisia and Egypt, that may, within the context of our model, have moved the regime to a state of civil unrest or realized revolution.

While it is clear that, in the context of our Arab Spring revolution model, reliable time series measurements of quantities that would correspond to the fraction of protesters in the revolution are not available, and while it is thus not realistic to consider fitting the model and its parameters to observed time series, it is interesting to note that there are recent efforts that attempt to gather quantitative data that can be used for social science research using, for example, Blogs and online social media platforms like Twitter. See, e.g., [89, 116] for the case of the Arab Spring revolutions. Unfortunately, with the currently available data this type of comparison can only be done at a superficial level, but it is an intriguing prospect that this kind of approach may offer new ways to quantitatively test models in social science in the future when more and higher-quality quantitative data of this type may become available.

2.5.3 Application: Non-Arab Spring countries

Although our model was developed with the specific goal of modelling revolutions in dictatorial regimes that employ censorship and police repression initially applied to the Arab

Spring, it is interesting to consider its application to various situations in other countries. In this section we consider the cases of the failed 2009 "Green Revolution" in Iran, and of present-day China and Somalia.

2009 Iranian "Green Revolution"

The protests following Iran's 2009 election, dubbed the "Green Revolution", were ultimately put down by the regime despite widespread use of social media technology. The large amount of resources that were available to the Iranian regime and the relative novelty of applying social media for use in protest [23, 162] may be consistent with a regime with parameters, in the context of our model, in Region III0 (stable police state). This, in turn, would be consistent with the failure of the Green Revolution.

China

While the current regime in China differs from the pre-revolutionary regimes in Tunisia and Egypt in many aspects, it is interesting to consider how our model may apply to China in terms of the influence of state control on the media and the Internet, and police control of dissident opinion. The number of "mass group incidents" reported annually in China has been rising consistently for at least two decades [188]. Constant low levels of protest may correspond to the civil unrest equilibrium ($r = c^*$) in the meta-stable police state (Region IIIe). In our model rising levels of protest would correspond to rising c^*. We note that, in the context of our model, a continued rise of c^* through increased Internet and social media exposure may eventually result in increasing the chance of a realized revolution, as argued in Section 2.5.2.

Somalia

The failed state region (Region II) features low visibility α (weak media) and low police capacity β (weak government), which prevents individuals from joining any popular movements and prevents the government from reigning in existing movements, respectively. This results mathematically in an uncountable number of equilibria contained in $(\beta, 1-\alpha)$. This appears consistent with the slow and erratic rise and fall of local militia and the succession of weak central governments seen in Somalia from 1991 [122]. Our model predicts that a successful national state could arise from either (a) improving police capacity of the transitional government (increasing β), or (b) increasing social cohesion and the capacity of the

media in Somalia (increasing α). Interestingly, Somalia has developed a sophisticated and affordable telecommunications sector [53], which may mean that an increased visibility α is not unrealistic.

2.6 Conclusion

We have introduced a simple compartmental model for the dynamics of a revolution in dictatorial regimes that employ censorship and police repression. The model features visibility and policing terms that describe rapid transitions between 0 and 1 as a function of the size of the revolution, for which we have provided conceptual and network-based mathematical justifications. The dynamical behaviour of the model was classified, leading to a division in parameter space that is interpreted naturally in terms of stability of the regime (stable police state, meta-stable police state, unstable police state, and failed state). We show in Section 2.4 that these dynamical properties of the model are generic for a broad class of visibility and police capacity functions that feature rapid transitions between 0 and 1. We investigated how the model can be applied to the Arab Spring revolutions in Tunisia and Egypt, taking into account the influence of the Internet and new media on the visibility of the revolution and the ensuing reduced effectiveness of censorship. Within the model this leads to significant, discontinuous changes in regime stability, which greatly increases the probability of realized revolutions. These properties of the model inform possible answers to questions on causes and timing of the Arab Spring revolutions, and the role of the Internet and new media. The broader relevance of the model classification was also investigated by applying it to the current political situations in Iran, China and Somalia. We note here that our model is general enough to potentially capture a wide range of social change phenomena in democratic regimes as well. Consider, for example, social norms such as the recognition of gay marriage [10] or the practice of cremation versus burial [42]. Both of these issues have recently gained substantial support over a relatively short period, despite significant resistance. We emphasize, however, that in these cases the opposition to these issues comes not from the policing capacity of a government but from elements of society that are reluctant and/or resistant to change, and may fiercely oppose the change until the case is deemed lost. We note that both of these social issues can also be considered to be binary choice conformity problems. This further illustrates the close relationship between collective action problems and conformity problems that we previously discussed in Chapter 1. Although we do not develop models for either of these two issues specifically, a detailed model of a binary choice conformity problem is provided in Chapter 4 which models the prevalence of smoking as a binary choice conformity problem where individuals

choose either to become smokers or non-smokers.

Simple models like ours have the advantage of relying on just a few basic assumptions about individual and communal behaviour. More elaborate models can easily be imagined, but adding detail to a model comes at the expense of its tractability. Indeed, the very simplicity of our model is what admits a complete and rigorous mathematical classification of its dynamical behaviour, as well as an interpretation that offers interesting hypotheses about how Arab Spring-type revolutions may unfold. Of course, simple models like ours are also subject to many limitations. For example, our model is not capable of describing the Arab Spring revolution in Libya and Syria, or the counter-revolution in present-day Egypt, because they do not correspond to our basic assumption that the population is uniform in its dislike for the current regime. In particular, tribalism in Libya [135], religion in Syria [87], and religion, secular democracy, and the vested interests of the military in present-day Egypt [103] divide the population into factions that cannot easily be accounted for in a one-dimensional mathematical model. Moreover, events in these countries are further complicated by significant external interference and interventions. These complications, among others, must be addressed by more sophisticated models and constitute a significant avenue of future research.

Chapter 3

A Hierarchy of Linear Threshold Models for the Spread of Political Revolutions on Social Networks

This chapter is based on research published in the *Journal of Complex Networks* [110].

3.1 Introduction and Motivation

In this chapter, we propose a linear threshold agent-based model (ABM) for the spread of political revolutions on social networks in dictatorial regimes. In our ABM, nodes of the network can be in two states: active, i.e., participating in the revolution, or inactive. Transitions from the inactive to the active state are governed by a growth process that uses a traditional linear threshold mechanism: an individual v may change from the inactive to the active state if the fraction of neighbours in the active state exceeds the linear threshold θ_v. Linear threshold models have been studied before in the context of influence maximization [98], enforcement of unpopular norms [36], and general "complex contagion" processes [34, 35, 69, 72, 75, 107, 187]. The development and analysis of linear threshold processes has been highly influenced by more established epidemiology [6, 12, 84] and rumour spreading [113, 191] models, with one key difference: whereas a contact with a single "infected" individual is enough to spread a contagion in most epidemiology or rumour spreading applications, i.e. in "simple contagion" processes, multiple contacts with infected individuals are generally necessary to spread a contagion in a linear threshold

process. Indeed, epidemiological and rumour spreading models can be considered to be a special case of linear threshold models where the threshold parameter is chosen sufficiently small such that one infected neighbour enables propagation. Most of the linear threshold models that have been studied in the context of social spreading or contagion processes [34, 35, 36, 69, 75, 98, 187] do not incorporate a mechanism by which active nodes become inactive (which is relevant to the context of our application), evolve in discrete time, and are primarily interested in the final fraction of active nodes. In contrast, our ABM specifies a mechanism by which active nodes become inactive, evolves in continuous time, and is concerned with modelling the temporal dynamics of the number of active nodes.

The main justification for our linear threshold modelling approach in the context of political revolutions is that the decision to join a political revolution can be assumed to be a collective action problem [107]: if individuals act unilaterally they are subject to retaliation by the regime, whereas if they act collectively then the regime loses the ability to punish due to a lack of resources. The linear threshold modelling approach is consistent with the collective action principle, since an individual transitions from an inactive to an active state only after the number of neighbours in the active state has been observed to reach the critical fraction θ_v, and the individual deems the revolution of sufficient size to consider participation. In order to incorporate communication network structure explicitly, we represent interactions between individuals $v \in V$ by edges $e \in E$ in an undirected graph $G = G(V, E)$. This contrasts with previous attempts to model political revolutions using linear threshold models that assume homogeneous mixing [75, 107], i.e. that assume that $G = G(V, E)$ is the complete graph.

The goals and contributions of this chapter are as follows. We start by presenting the linear threshold ABM for the spread of political revolutions on social networks, and then derive a hierarchy of simplified ordinary differential-equation (ODE) models of varying degrees of sophistication that characterize the solutions of the linear threshold ABM, see Table 3.1 and Fig. 3.1. Simplified aggregate or population-level dynamical models can provide more cost-effective dynamical simulations, and can give insight in the qualitative dynamics of the ABM and its parameter regimes through dynamical analysis of the simplified model. It is a significant challenge to incorporate actual network structure in simple ODE models, and in this chapter we present two new effective ways to incorporate network structure into the one-compartmental ODE that approximates the dynamical evolution of the expected fraction of the population that participates in the revolution in the linear threshold ABM model. These approaches make use of the degree distribution of the graph or samplings of the network[1], and result in the *binomial visibility function* (BVF) and the *empirical visibility function* (EVF) population-level ODE models, which we demonstrate

[1]We use the terms graph and network interchangeably.

to be equivalent in the limit of large network and sample size. The performance of these ODE models in approximating the ABM is compared to the performance of the *degree approximation* (DA) model, a more expensive higher-order model that we modify from [128].

Table 3.1: Summary of hierarchy of models from highest (network-level) to lowest (population-level) complexity.

Model	Abbreviation	Section	Equation(s)
Agent-Based Linear Threshold	ABM	3.3	(3.1) and (3.2)
Degree Approximation*	DA	3.5	(3.15)
Binomial Visibility Function	BVF	3.4.1	(3.7) and (3.10)
Empirical Visibility Function	EVF	3.4.1	(3.7) and (3.11)
Step Visibility Function	SVF	2.2.1	(2.1) and (3.8)

*The Degree Approximation (DA) model is modified from the approach of [128].

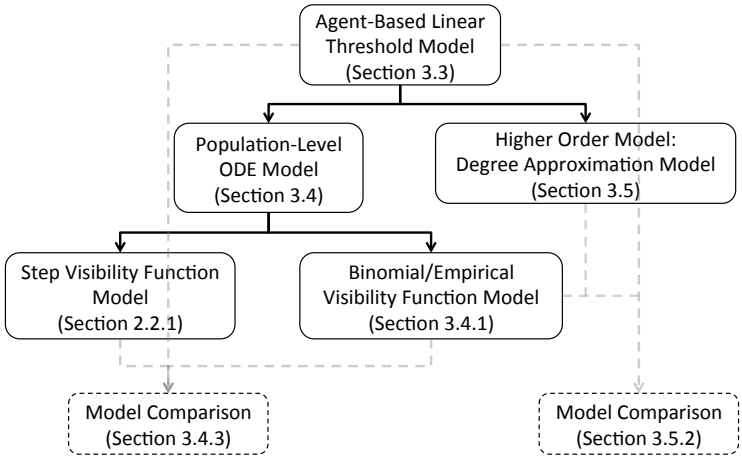

Figure 3.1: Summary of hierarchy of models from highest (network-level) to lowest (population-level) complexity.

We also relate the ABM to the population-level one-compartmental ODE model for the spread of political revolutions that was proposed in Eq. (2.1) of Section 2.2.1 (see also [109]). Because model (2.1) indirectly takes network structure into account through a step function, whose single parameter (i.e. the protester visibility parameter α)[2] determines

[2]We recall that interpretation of the visibility parameter α is that people will only join the revolution when the revolution is of sufficiently large size to be visible to the population, while repressive regimes attempt to make unrest invisible through censorship.

the fraction of participants at which the growth term (for the size of the revolution) is switched on discontinuously, we call model (2.1) the *step-visibility function* (SVF) model. The SVF model was justified in Section 2.2.3 (see also [109]) by a simplified network model with strong assumptions (e.g. we assume that the states of an individual's neighbours are uncorrelated), and as such it is a precursor of the BVF/EVF models introduced in this chapter. In this chapter, we show that our ABM is mathematically consistent with the SVF model and we test the assumptions from Section 2.2.3 with the BVF/EVF models and the ABM model on real networks, corroborating a posteriori the assumptions from Section 2.2.3, and showing that model (2.1) (i.e. the SVF model) reproduces the qualitative behaviour of the ABM and the spread of a revolution under a linear threshold model. This provides further justification for the SVF model, and at the same time provides inexpensive ways to predict the parameter regime behaviour of the ABM presented in this chapter using simple compartmental models like the SVF and the new BVF/EVF models.

A further contribution of this chapter is that we propose an extension of the concept of basic reproduction number (R_0) from epidemiological modelling [84] to characterize the potential of networks to spread the revolution and show that this quantity is proportional to the initial slope of the BVF/EVF we introduce in this chapter. The basic reproduction number for our ABM with a linear threshold process is easy to compute, and we show that it gives useful predictions for the extent of the spreading of the revolution. In small-scale numerical tests we investigate experimentally the differences in spreading behaviour that occur under the linear threshold ABM model when applied to some empirical (modern) online social networks versus (traditional) offline social networks, searching for quantitative evidence that political revolutions may be facilitated by the modern online social networks of social media. We emphasize that the broad applicability of the linear threshold process implies that the framework, methods, and results presented in this chapter are relevant to the study of many related social spreading processes.

The rest of this chapter is structured as follows. Section 3.2 describes the Facebook and physical contact empirical networks on which we test our ABM and approximations. Sections 3.3-3.4 describe the linear threshold ABM and how population-level ODE models can be derived from it, respectively. This includes the SVF model from Section 2.2.1, and the new BVF and EVF models which explicitly incorporate the structure of the network. Section 3.4 also explores the equivalence of the BVF and EVF and provides a detailed numerical comparison between the ABM, the SVF, and the BVF/EVF models. Section 3.5 confirms the usefulness of the BVF and EVF models by showing that they perform no worse than the higher order degree approximation model of [128], while being much less expensive to evaluate. Section 3.6 extends the concept of basic reproduction number (R_0) from epidemiology [84] to the linear threshold process specified in Section 3.3 and illustrates

how it can be interpreted in terms of the BVF/EVF. Finally, Section 3.7 explores how our methods can be used to study the differences in spreading behaviour of political revolutions on online and offline social networks under the linear threshold mechanism, and Section 3.8 concludes.

3.2 Network Data

In this section we describe the two empirical social networks that will be used in this chapter to validate and compare the ABM, DA, BVF/EVF and SVF models (Table 3.1 and Fig. 3.1) for the spread of political revolutions on social networks under the linear threshold model.

The first network we consider is the Facebook subnetwork presented in [120]. It was constructed by combining the *ego-networks*[3] of ten individuals and then taking the largest connected component of the resulting network. We refer to this network as the *Facebook social network* (G_F). The second network we consider is the physical contact network presented in [158]. It was constructed by distributing wireless sensors to students, teachers, and staff at a U.S. high school during a one day period. When two wireless sensors are in proximity of one another, i.e. when they are less than approximately 3m apart, they register an interaction with a temporal resolution of 20s. Therefore, the communication network we extract from this data is referred to as the *physical contact network* (G_P). For details on the network extraction protocol for G_P we refer to Appendix C.1. To facilitate the comparison of experimental results, the minimal contact duration to register an edge in G_P is chosen such that G_P and G_F have approximately the same average degree, see Appendix C.1.

The Facebook network G_F has $N = 3,963$ nodes with sparsity[4] $S = 0.01$, and is visualized in Fig. 3.2a. The physical contact network G_P has $N = 776$ nodes with sparsity $S = 0.06$, and is visualized in Fig. 3.2b. The cumulative degree distributions[5] are displayed in Fig. 3.3. The distributions of local clustering coefficients are displayed in Fig. 3.4.

[3]For a network $G = G(V, E)$ the *ego-network* of an individual $v \in V$ is a subnetwork of G composed of the individual v (called the *ego-node*), the neighbours of individual v, and all of the connections between these individuals.

[4]The *sparsity* of a network is defined as the fraction of possible edges that are present in the network.

[5]If ρ_k is the fraction of nodes with degree k (i.e., the degree distribution) of the graph, then the cumulative degree distribution is the function $F(k) = \sum_{j=1}^{k} \rho_j$. The cumulative degree distribution is displayed in place of the degree distribution because the degree distribution is subject to significantly more noise than the cumulative degree distribution. This is a common approach for empirical networks.

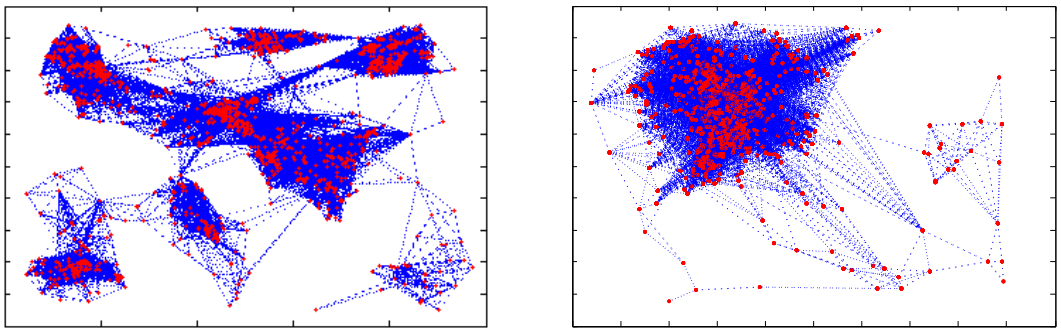

(a) Facebook social network (G_F) (b) Physical contact network (G_P)

Figure 3.2: Visualization of Facebook (G_F) social and physical contact (G_P) networks using the Gursoy-Atun algorithm [78] as implemented in the Matlab Boost Graph Library package [70].

It is interesting to consider how the physical contact and Facebook networks differ on some common network measures, see Table 3.2. Comparing the physical contact and Facebook social networks along these admittedly limited measures nevertheless highlights substantial structural differences between these networks. The physical contact network appears more homogeneous than the Facebook social network, in the sense that the Facebook social network can be visually grouped into distinct communities, whereas the physical contact network cannot (Fig. 3.2). The network measures presented in Figs. 3.3-3.4 and Table 3.2 are generally supportive of this observation. For example, the cumulative degree distributions displayed in Fig. 3.3 show that the physical contact network has an approximately normal distribution with a relatively small standard deviation, and hence, a thin tail. In contrast, the Facebook social network has a much broader distribution (possibly scale-free or exponential, although this cannot be determined conclusively with such a small network size) with a relatively large standard deviation, and hence, a fat tail. Similarly, it is known that online social networks often have larger clustering coefficients than offline social networks. In Section 3.7 we use the Facebook and physical contact networks as representatives of online and offline social networks, respectively, to investigate differences in propagation properties that may arise between online and offline social networks within the linear threshold propagation model. We also use these networks to test and compare the propagation models we derive in this chapter.

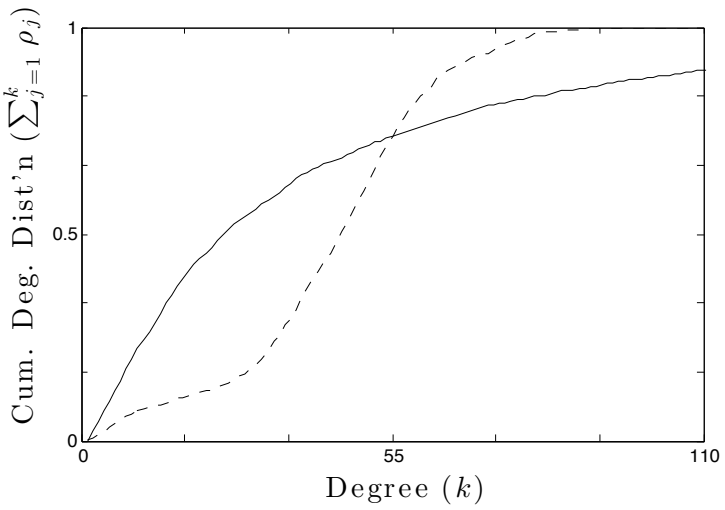

Figure 3.3: Cumulative degree distribution of Facebook social (G_F - solid) and physical contact (G_P - dashed) networks.

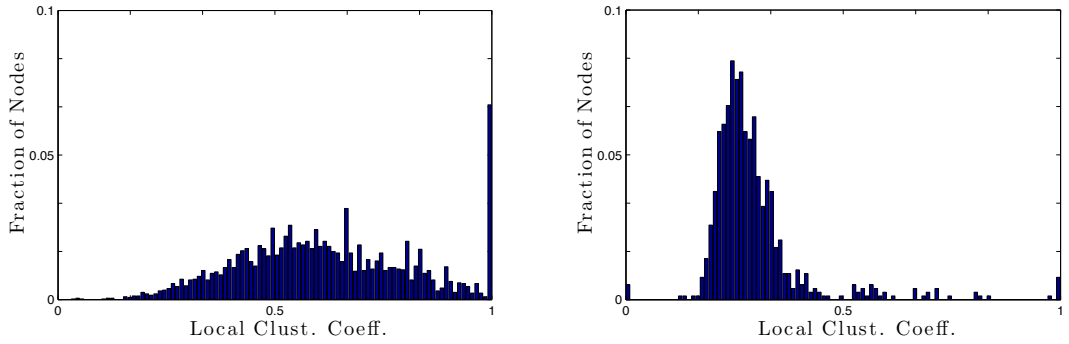

(a) Facebook social network (G_F)

(b) Physical contact network (G_P)

Figure 3.4: Distribution of local clustering coefficient for Facebook social (G_F) and physical contact (G_P) networks.

Table 3.2: Summary of common network measures for Facebook social (G_F) and physical contact (G_P) networks.

Quantity	Network	
	Physical Contact (G_P)	Facebook (G_F)
Number of Nodes (N)	776	3,963
Average Degree ($\mu_k \pm \sigma_k$)	44.8 ± 18.2	44.5 ± 52.4
Minimum/Maximum Degree ($k_{\min}/k_{\max}$)	1/109	2/1034
Sparsity (S)	0.06	0.01
Diameter (D)	8	8
Average Path Length($\mu_l \pm \sigma_l$)	2.37 ± 0.82	3.77 ± 1.29
Average Local Clust. Coeff. ($\mu_c \pm \sigma_c$)	0.29 ± 0.12	0.62 ± 0.20

3.3 Specification of the Linear Threshold Agent-Based Model (ABM)

Consider a population of individuals who are represented by nodes $V = \{v_i\}_{i=1}^N$, and whose interactions are represented by edges $E = \{e_i\}_{i=1}^M$, of the graph $G = G(V, E)$ with degree distribution ρ_k. An individual $v \in V$ at time t can be in one of two states, i.e. $s_v(t) = 0$ (*inactive* in the revolution) or $s_v(t) = 1$ (*active* in the revolution). We assume that the network G is static so that the dynamics of the ABM can be fully specified by providing rules for the transition of individuals from an inactive state to an active state and vice versa, together with parameter values and an initial condition. For the moment, we set aside the issue of the choice of parameters and initial conditions and restrict the remainder of this section to the specification and justification of the transition rules.

3.3.1 Growth Process: Inactive ($s_v = 0$) to Active ($s_v = 1$)

To specify the growth process of the ABM, we use the standard linear threshold model that has been used before in the context of social spreading or contagion processes such as opinion formation, technology adoption, marketing, rioting, and political movements [34, 35, 36, 75, 98, 187]. As mentioned in Section 3.1, the choice of individuals to join a revolution is a collective action problem [107]: individuals are averse to unilateral action against the regime for fear of severe retaliation, but are willing to take action collectively in the belief that the regime will lack sufficient resources to punish the entire collective. Thus, an individual $v \in V$ will decide to join the revolution if he or she believes that it has grown sufficiently large to reduce the risk of retaliation from the regime to an acceptable level. It is reasonable, therefore, to assume that individuals will decide to join the revolution only after a large enough fraction of their neighbours in the social network have done so. This behaviour is captured by the following transition rule: if individual $v \in V$ has k neighbours $\{w_j\}_{j=1}^k$, if $s_v(t) = 0$, and if

$$\sum_{j=1}^{k} s_{w_j}(t) \geq \theta_v k, \tag{3.1}$$

i.e., if v is inactive at time t and has at least a fraction θ_v of its neighbours that are active, then node v transitions from state 0 to state 1 at time $t' = t + \xi_1$, where $\xi_1 > 0$ is the first arrival time of a Poisson process with rate c_1. We say that nodes that satisfy Eq. (3.1) can "see" the revolution, i.e. the revolution is visible to them. Alternatively, we say that these nodes are "considering joining the revolution". Parameter c_1 determines the timescale of

the growth process. In other words, c_1 determines the rate at which individuals join the revolution, and hence, we call c_1 the protesters' enthusiasm parameter.

While most of the linear threshold models that have been studied in the context of social spreading or contagion processes [34, 36, 35, 98, 75, 187] consider evolution in discrete time by choosing ξ_1 constant, we have specified ξ_1 so that Eq. (3.1) evolves in continuous time. Specifying ξ_1 in this way is often considered in biology and epidemiology [6, 12], since it (a) facilitates the comparison of the ABM with continuous-time population-level models (see Section 3.4), and (b) eliminates the problem of choosing a suitable time-step for iterating the discrete-time process (see Gillespie's Algorithm in Appendix C.2). Furthermore, we note that by choosing ξ_1 to be the first arrival time of a Poisson process, we are assuming that the decision making process is a Markov, or memoryless, process. Specifically, conditioned on the state of the system at time t, we assume that the likelihood of an individual joining the revolution in the time interval $[t, t + \Delta t]$ is independent of the state of the system at any time $\tilde{t} < t$. In words, the future is independent of the past, given the present.

3.3.2 Decay Process: Active ($s_v = 1$) to Inactive ($s_v = 0$)

We assume that once an individual $v \in V$ has become active the regime will attempt to arrest or disperse him or her, thus returning v to an inactive state. As in the previous section, we assume that this is a memoryless process. We further assume that the regime can only arrest or disperse protesters so long as the total fraction of active protesters remains less than the regime's finite *police capacity* $\beta \in (0, 1)$. The transition rule can then be characterized by the following: if $s_v(t) = 1$ and

$$\frac{1}{N} \sum_{w \in V} s_w(t) < \beta, \tag{3.2}$$

i.e., if v is active at time t and the fraction of active individuals is less than the regime's police capacity β, then the node v transitions from state 1 to state 0 at time $t' = t + \xi_2$, where $\xi_2 > 0$ is the first arrival time of a Poisson process with rate c_2, which determines the timescale of the decay process. In other words, c_2 determines the rate at which the regime can disperse protesters, and hence, we call c_2 the policing efficiency parameter.

This fully specifies the evolution of the ABM, which we simulate using Gillespie's algorithm, see Appendix C.2.

3.4 Population-Level ODE Approximation for the ABM

Due to the difficulty of analyzing and computational cost of simulating the ABM, we consider the following derivation of a one-dimensional ODE approximation that can be applied to lend insight into how the ABM process behaves on different networks. We begin by defining $r_a(t) = r_a(t|t_0)$ to be the fraction of nodes (in the ABM) that are expected to be in state 1, i.e. the expected fraction of active nodes at time t conditioned on information at time t_0:

$$r_a(t) = \frac{1}{N}\sum_{i=1}^{N} \mathbb{E}\left[s_{v_i}(t)|t_0\right].$$

The fraction of nodes that are expected to be in state 0 (inactive) at time t, conditioned on information at time t_0, is then $1 - r_a(t)$. We now write the change in r_a from time t_0 to time t as

$$\Delta r_a(t) = r_a(t) - r_a(t_0) = g(t|t_0) - d(t|t_0), \tag{3.3}$$

where $\Delta t = t - t_0$, and the expected growth and decay of the fraction of active nodes are modelled by the non-negative growth and decay functions $g(t) = g(t|t_0)$, and $d(t) = d(t|t_0)$, respectively. In order to obtain the one-compartment model we will need to approximate the quantities $g(t)$ and $d(t)$ in terms of $r_a(t_0)$, $\{\theta_{v_i}\}_{i=1}^{N}$, and β.

For notational convenience we begin by considering the case where $\forall v \in V : \theta_v = \theta$. The pool of individuals that can go from active to inactive at time t_0 is $r_a(t_0)$. Active nodes become inactive at the first arrival time of a Poisson process with rate c_2, provided that the fraction of active nodes does not exceed the regime's police capacity, i.e. $r_a(t_0) < \beta$. Thus,

$$d(t) = \left[c_2\ r_a(t_0)\ \Delta t + o(\Delta t)\right]\ p(r_a(t_0);\beta) = c_2\ r_a(t_0)\ p(r_a(t_0);\beta)\ \Delta t + o(\Delta t), \tag{3.4}$$

where we take $p(r_a(t_0);\beta) = \mathbb{I}_{\{r_a(t_0)<\beta\}}$, and where $\mathbb{I}_{\{\cdot\}}$ is the indicator function, i.e.

$$\mathbb{I}_{\{X\}} = \begin{cases} 1 & \text{if } X \text{ is true} \\ 0 & \text{otherwise} \end{cases}.$$

Next, let

$$\nu(r_a(t_0);\theta) = \frac{1}{N}\left[\sum_{v\in V} \mathbb{I}_{\left\{\sum_{j=1}^{k_v} s_{w_{v,j}}(t_0) \geq \theta k_v\right\}}\right], \tag{3.5}$$

where $\{w_{v,j}\}_{j=1}^{k_v}$ denotes the neighbours of individual $v \in V$. In words, we let $\nu(r_a(t_0);\theta)$ be the fraction of the total population at time t_0 that can see the revolution. We call $\nu(r_a(t_0);\theta)$ the visibility function of the ABM. In the chosen notation for ν it is emphasized that ν depends on $r_a(t_0)$, i.e. on the fraction of nodes that are active at time t_0. At time $t = t_0$, the pool of individuals that can go from inactive to active, i.e. that are considering joining the revolution, is therefore approximately $(1 - r_a(t_0))\ \nu(r_a(t_0);\theta)$. Since inactive nodes that can see the revolution become active at the first arrival time of a Poisson process with rate c_1, we have

$$g(t) = c_1\ (1 - r_a(t_0))\ \nu(r_a(t_0);\theta)\ \Delta t + o(\Delta t). \tag{3.6}$$

Combining Equations (3.3)-(3.6) and dividing by Δt gives

$$\frac{\Delta r_a}{\Delta t} = c_1\ (1 - r_a(t_0))\ \nu(r_a(t_0);\theta) - c_2\ r_a(t_0)\ p(r_a(t_0);\beta) + o(1).$$

Taking the limit as $\Delta t \to 0$ and suppressing the t_0 argument yields

$$\frac{dr_a}{dt} = c_1\ (1 - r_a)\ \nu(r_a;\theta) - c_2\ r_a\ p(r_a;\beta). \tag{3.7}$$

This establishes a general population-level ODE model that approximates the linear threshold ABM, with visibility function $\nu(r_a;\theta)$. To close the model, the visibility function has to be specified, guided by Equation (3.5). We observe that the model of Section 2.2.1, given in Eq. (2.1), turns out to be of the form (3.7) as derived for the ABM, when $\nu(r;\theta)$ is chosen to be the step-visibility function (SVF) $v_s(r;\alpha)$, i.e. when $\nu(r;\theta)$ in Eq. (3.7) is taken to be

$$\nu(r;\theta) = v_s(r;\alpha) \equiv \mathbb{I}_{\{r>1-\alpha\}}. \tag{3.8}$$

This is the simplest closure for Eq. (3.7). In the following section we introduce two novel candidates for the visibility function that explicitly take into account the structure of the underlying network: the binomial and empirical visibility functions (BVF and EVF, respectively). These two visibility functions can be seen to be equivalent in the limit of large network and sample size, as shown in Section 3.4.2.

We conclude this section by noting that the above derivation can easily be modified for the case where the θ_v are not uniform. In this case the growth and decay functions g and d, respectively, can be computed by averaging over the distribution of the θ_v's. As such, the procedures outlined below for the SVF, BVF, and EVF, can also easily be modified for the case of non-uniform θ_v.

3.4.1 Visibility Functions for Population-Level ODE Models

Binomial Visibility Function

Consider the following derivation. Suppose that $v \in V$ is a node in the network $G = G(V, E)$ with degree k, and that the fraction of nodes active in the revolution is r. Also, assume that the states of v's neighbours are active with probability r independently so that the probability of v having j neighbours active in the revolution is

$$\binom{k}{j} r^j (1-r)^{k-j}.$$

The probability that the fraction of v's neighbours exceeds the linear threshold θ is, therefore,

$$\sum_{j=\lceil \theta k \rceil}^{k} \binom{k}{j} r^j (1-r)^{k-j} = \text{BinCDF}(k - \lceil \theta k \rceil; k, 1-r), \tag{3.9}$$

where $\text{BinCDF}(x; n, p)$ is the cumulative distribution function for the binomial distribution with n trials and probability of success p evaluated at x. This quantity can be used to approximate the function $\nu(r; \theta)$ (the expected fraction of the population that can see the revolution). In the network-based justification of the SVF model given in Section 2.2.3 we argued that Eq. (3.9) with $k = \lfloor \phi \rfloor$, i.e. Eq. (3.9) with k taken to be the average degree, is a sigmoidal function that can be approximated by step-visibility function $v_s(r; \alpha)$ for some appropriate visibility parameter α. One possible improvement on adopting

$$\nu(r; \theta) = v_s(r; \alpha(\theta))$$

in Eq. (3.7) (resulting in the SVF model) is to actually adopt

$$\nu(r; \theta) = \text{BinCDF}(\lfloor \phi \rfloor - \lceil \theta \phi \rceil; \lfloor \phi \rfloor, 1-r)$$

instead. However, we can improve on this even further in terms of using more relevant information about the real graph by choosing $\nu(r; \theta)$ to be equal to

$$v_b(r; \theta, \rho) = \sum_k \rho_k \text{BinCDF}(k - \lceil \theta k \rceil; k, 1-r) = \sum_k \sum_{j=0}^{k - \lceil \theta k \rceil} \rho_k \binom{k}{j} (1-r)^j r^{k-j}, \tag{3.10}$$

where ρ_k is the degree distribution. We call the function $v_b(r; \theta, \rho)$ in Eq. (3.10) the *binomial visibility function* (BVF), see Fig. 3.5. Note the steep sigmoidal shape of $v_b(r; \theta, \rho_k)$ for

the two empirical networks of Fig. 3.5, further justifying the choice of a step visibility function as an appropriate approximation in the SVF model of Section 2.2.1 [109]. We note that computing the binomial visibility function on some grid of $\mu+1$ discrete r-values $r \in \{r_i\}_{i=1}^{\mu+1}$ requires calculating a double sum for every $i = 1, \dots \mu + 1$. In practice, this can be very expensive, especially for degree distributions ρ_k with fat tails. For this reason, the following section introduces a complementary but less expensive approximation to the visibility function that is equivalent to the binomial visibility function of Eq. (3.10) in the limit of large network and sample size.

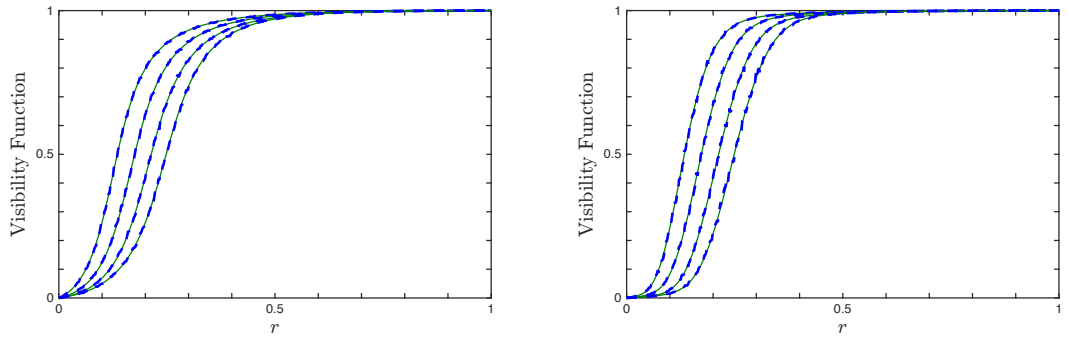

(a) Facebook social network (G_F) (b) Physical contact network (G_P)

Figure 3.5: Binomial visibility function (solid green) and empirical visibility function (dashed blue) for Facebook social (G_F) and physical contact (G_P) networks generated for linear thresholds $\theta = \{0.13, 0.17, 0.21, 0.25\}$ (curves from left to right). The binomial and empirical visibility functions are computed on the grid $r \in \{0, \frac{1}{512}, \frac{2}{512}, \dots, 1\}$.

Empirical Visibility Function

We now propose to reconstruct the binomial visibility function $v_b(r; \theta, \rho_k)$ empirically by sampling the underlying network directly. Specifically, for fixed θ we approximate v_b by a discrete function $v_e \in \mathbb{R}^{\mu+1}$ where $\forall j = 1, \dots, \mu + 1 : v_{e,j} \approx v_b(\frac{j-1}{\mu})$. Since $v_b(0) = 0$ and $v_b(1) = 1$, we set $v_{e,1} = 0$ and $v_{e,\mu+1} = 1$. We calculate $v_{e,j}$ for $j = 2, \dots, \mu$ as follows.

1. Seed the network with fraction $r_j = \frac{j-1}{\mu}$ active nodes in expectation. (Generate N random numbers drawn from the uniform distribution on $[0, 1]$: $rand_i$ for $i = 1, \dots, N$. If $rand_i < \frac{j-1}{\mu}$ then node i is active, otherwise node i is inactive.)

2. Determine the fraction of nodes that can see the revolution.

3. Repeat 1-2 a total of rep times and set $v_{e,j}$ to be the average over the realizations. (We choose $rep = 100$ in our numerical experiments.)

We denote the linear interpolation between the $v_{e,j}$ by the function $v_e(r;\theta)$, i.e.

$$v_e(r;\theta) = \begin{cases} v_{e,j} & \text{if } r = \frac{j-1}{\mu} \text{ for some } j = 1,\dots,\mu+1 \\ v_{e,j}\frac{j/\mu - r}{1/\mu} + v_{e,j+1}\frac{r-(j-1)/\mu}{1/\mu} & \text{if } r \in [\frac{j-1}{\mu}, \frac{j}{\mu}] \text{ for some } j = 1,\dots,\mu \end{cases}, \tag{3.11}$$

which we call the *empirical visibility function* (EVF), see Fig. 3.5.

We now make several brief remarks about the procedure for calculating the empirical visibility function, outlined above. Firstly, we remark that this procedure is easily extended to the case where θ_v are specified for each individual node. However, to keep our analysis as simple as possible we continue to restrict ourselves to the case where $\theta_v \equiv \theta$. In the next section we show that the empirical and binomial visibility functions are equivalent in the limit of large network and sample size. Figure 3.5 shows that they can also be expected to give similar results for finite network and sample size. Finally, we emphasize that compared to the binomial visibility function, the procedure for calculating the empirical visibility function is in practice significantly less costly to implement, especially for networks with fat-tailed degree distributions.

3.4.2 Equivalence of Binomial and Empirical Visibility Functions

We briefly argue that the empirical and binomial visibility functions are equivalent in the limit of large network and sample size, i.e. as $N, rep \to \infty$, see also Fig. 3.5. Let $v \in V$ be a node in the network $G = G(V,E)$ with degree k. Now, fix θ and suppose that we are calculating $v_{e,j} \approx v_b(\frac{j-1}{\mu}) = v_b(r_j)$ via the algorithm presented in Sec. 3.4.1. The l^{th} iteration of Step 1 of this algorithm selects $\mathcal{N}_l$ nodes uniformly at random to be active and results in $\mathcal{V}_l$ nodes that can see the revolution. Note that by the Law of Large Numbers [52]

$$\frac{\mathcal{N}_l}{N} \xrightarrow{p} r_j = \frac{j-1}{\mu} \text{ as } N \to \infty, \text{ and}$$

$$v_{e,j} = \frac{1}{rep}\sum_{l=1}^{rep} \frac{\mathcal{V}_l}{N} \xrightarrow{p} \mathbb{E}\left[\frac{\mathcal{V}_l}{N}\right] = \mathbb{E}\left[\frac{\mathcal{V}_1}{N}\right] \text{ as } rep \to \infty,$$

where $\xrightarrow{p}$ denotes convergence in probability. Since the $\mathcal{N}_l$ activated nodes are chosen uniformly at random, the statuses of the neighbours of v are independent. Therefore, the

probability that v can see the revolution is

$$\sum_{l=\lceil \theta k \rceil}^{k} \binom{k}{l} \left(\frac{\mathcal{N}_l}{N}\right)^l \left(1 - \frac{\mathcal{N}_l}{N}\right)^{k-j} = \text{BinCDF}\left(k - \lceil \theta k \rceil; k, 1 - \frac{\mathcal{N}_l}{N}\right).$$

It follows that the expected fraction of nodes that have degree k and can see the revolution is

$$\mathbb{E}\left[\rho_k \text{BinCDF}\left(k - \lceil \theta k \rceil; k, 1 - \frac{\mathcal{N}_l}{N}\right)\right],$$

and hence, the expected fraction of nodes that can see the revolution is

$$\mathbb{E}\left[\frac{\mathcal{V}_l}{N}\right] = \mathbb{E}\left[\sum_k \rho_k \text{BinCDF}\left(k - \lceil \theta k \rceil; k, 1 - \frac{\mathcal{N}_l}{N}\right)\right].$$

We conclude the proof by observing that by the Continuous Mapping Theorem [18] $\frac{\mathcal{N}_l}{N} \xrightarrow{p} r_j$ as $N \to \infty$ implies

$$\mathbb{E}\left[\frac{\mathcal{V}_l}{N}\right] \xrightarrow{p} \sum_k \rho_k \text{BinCDF}\left(k - \lceil \theta k \rceil; k, 1 - r_j\right) = v_b(r; \theta, \rho_k) \text{ as } N \to \infty,$$

and hence,

$$v_{e,j} \to v_b(r_j; \theta, \rho) \text{ as } N \to \infty, rep \to \infty.$$

3.4.3 Comparison of ABM, SVF, and BVF/EVF Models

Before we compare the ability of the SVF/BVF/EVF models to approximate the ABM we make several remarks. First, for a given θ in the ABM the SVF model requires that we specify α. We will propose an optimization procedure for choosing such an optimal α, which we call $\hat{\alpha}$. However, in order to preserve the continuity of our narrative we choose to defer this discussion to the end of this section. Second, we note that since the BVF and EVF are equivalent in the limit of large network and sample size, and are very close in practice (see Fig. 3.5), the observations we make below about the behaviour of the BVF model also apply to the behaviour of the EVF model. Finally, our numerical results will show that the behaviour of the BVF/EVF model has parameter regimes that induce behaviours analogous to those observed in the SVF model. Since the BVF/EVF have sigmoidal shape, this can be understood from the analysis presented in Section 2.3 for visibility and policing functions of general sigmoidal shape. This analysis shows that sigmoidal visibility functions (such as BVF and EVF) lead to parameter regimes that are very similar to the regimes of the SVF described in Fig. 2.3.

Numerical Comparison of ABM, SVF, and BVF/EVF Models

We now present a numerical comparison of ABM simulations with SVF and BVF/EVF models[6]. Figures 3.6-3.10 show some representative results when the SVF and BVF models are used to approximate the ABM. As illustrated in Fig. 3.6, for many combinations of parameters the SVF model is able to approximate the behaviour of the ABM in parameter Region IIIe (left panels of Fig. 3.6) and parameter Region III1 (right panels of Fig. 3.6)[7]. However, we observe that under certain circumstances stochastic effects and details of the network structure become important and the SVF model fails to approximate, even in a qualitative sense, the behaviour of the ABM. Specifically, the SVF model fails to capture the behaviour of the ABM when the initial condition is small ($r_0 << 1$), when the initial condition is near the boundary between the basins of attraction of $r = 0$ and $r = 1$ in Fig. 2.3b, and when parameters c_1, c_2, α, and β are close to the boundary between parameter Regions III0, IIIe, or III1 in Fig. 2.3a, see Figs. 3.7, 3.8, and 3.9, respectively. These are well-known limitations of population-level ODE models that attempt to approximate detailed network dynamics [6, 12].

[6] All computations are performed in Matlab. ABM simulations are performed using Gillespie's Algorithm, see Appendix C.2. Solution to the BVF/EVF models are computed using the Matlab function *ode45*. We describe in detail the solution to the SVF model at the end of this section, see also Appendix C.3

[7] For additional details on the analysis of the SVF model and a description of the parameters regimes III0, IIIe, and III1 see Section 2.3.

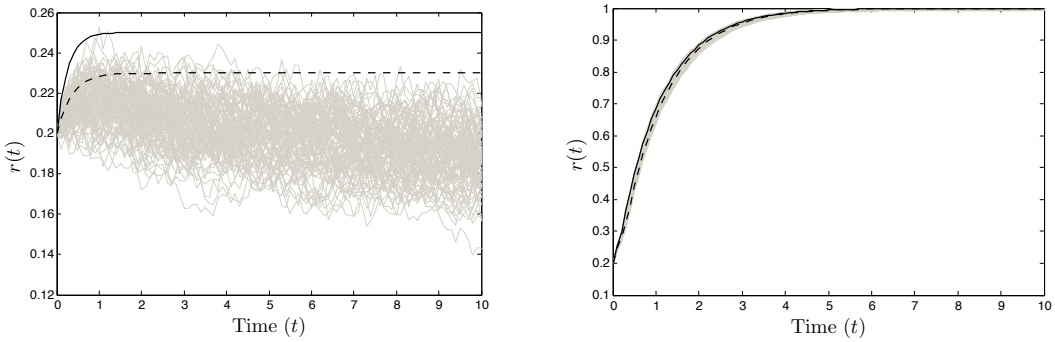

(a) Facebook social network (G_F, $1 - \hat{\alpha} = 0.131$)

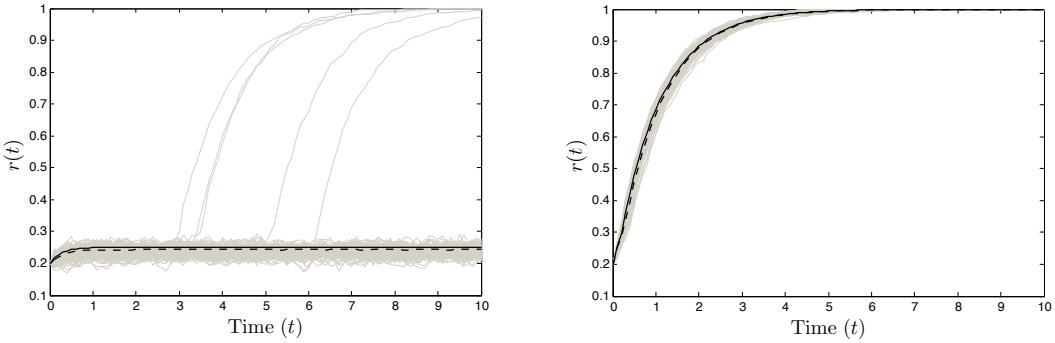

(b) Physical contact network (G_P, $1 - \hat{\alpha} = 0.134$)

Figure 3.6: Parameter values for which the SVF and BVF models approximate the ABM satisfactorily. (Left) Parameter Region IIIe ($c_2 = 3$, $1 - \hat{\alpha} < c^* < \beta$). (Right) Parameter Region III1 ($c_2 = 1$, $1 - \hat{\alpha} < \beta \leq c^*$). Time traces of $rep = 100$ ABM simulations (grey), solution to the SVF model (solid black), and BVF model (dashed black) with parameters $\theta = 0.11$, $\beta = 0.3$, $c_1 = 1$, and $r_0 = 0.2$.

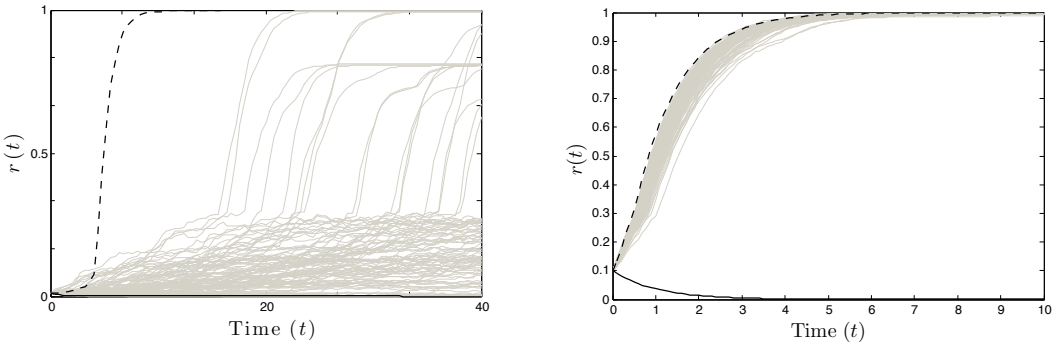

(a) Facebook social network (G_F, $1 - \hat{\alpha} = 0.131$)

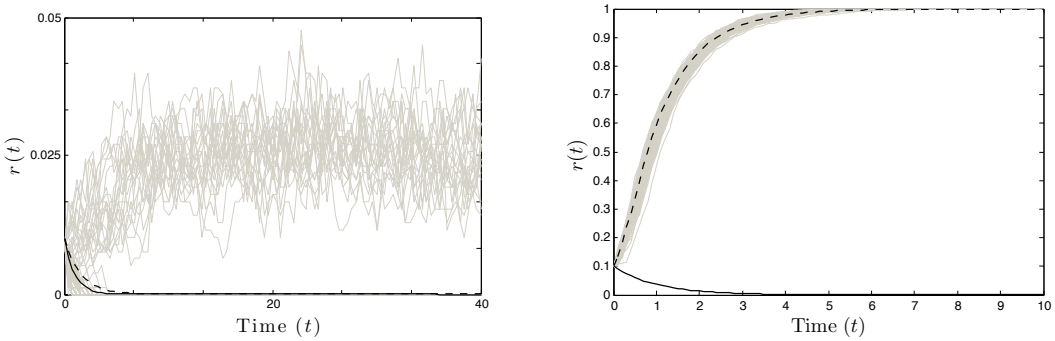

(b) Physical contact network (G_P, $1 - \hat{\alpha} = 0.134$)

Figure 3.7: The SVF model can fail to approximate the ABM when the initial condition is small ($r_0 = 0.01$, left), or is near the boundary between the basins of attraction of $r = 0$ and $r = 1$ ($r_0 = 0.1 \approx 1 - \hat{\alpha}$, right). Time traces of $rep = 100$ ABM simulations (grey), solution to the SVF model (solid black), and solution to the BVF model (black dashed) with parameters $\theta = 0.11$, $\beta = 0.3$ and $c_1 = c_2 = 1$ (Region III1, $1 - \hat{\alpha} < \beta \leq c^*$).

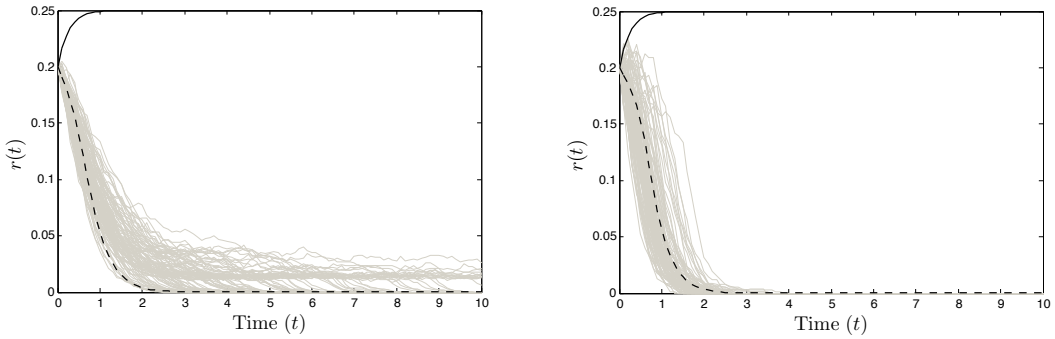

Figure 3.8: The SVF model can fail to approximate the ABM when the initial condition is near the boundary between the basins of attraction of $r = 0$ and $r = 1$, i.e. when $r_0 \approx 1 - \alpha$. (Left) Facebook social network (G_F, $1 - \hat{\alpha} = 0.179$) and (right) physical contact network (G_P, $1 - \hat{\alpha} = 0.180$). Time traces of $rep = 100$ ABM simulations (grey), solution to the SVF model (solid black), and solution to the BVF model (dashed black) with parameters $\theta = 0.17$, $\beta = 0.3$, $c_1 = 1$, and $c_2 = 3$, and initial condition $r_0 = 0.2$ (Region IIIe, $1 - \hat{\alpha} < c^* < \beta$). For these parameters, the SVF model has solution in Region IIIe, while the ABM and BVF models have solution in Region III0.

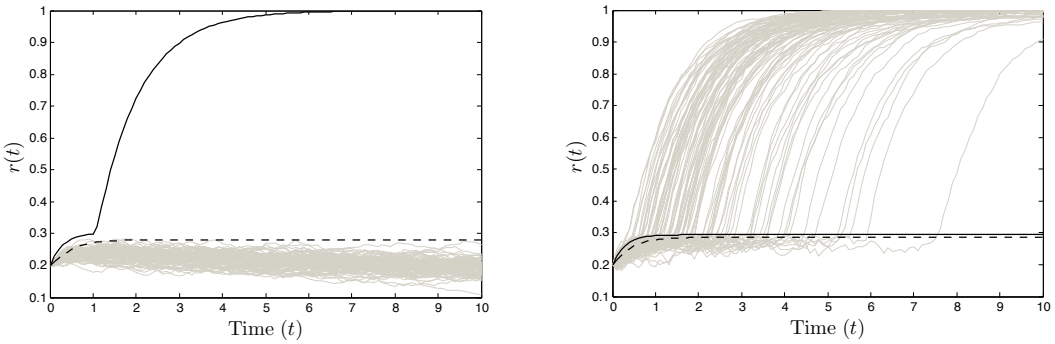

Figure 3.9: The SVF model can fail to approximate the ABM when parameters are close to the boundary between regions. (Left) Facebook social network (G_F, $1 - \hat{\alpha} = 0.163$) with parameters in Region III1 ($c_2 = 2.3$, $1 - \hat{\alpha} < \beta \leq c^*$). For these parameters, the SVF model has solution in Region III1, while the ABM and BVF models have solution in Region IIIe. (Right) Physical contact network (G_P, $1 - \hat{\alpha} = 0.165$) with parameters in Region IIIe ($c_2 = 2.4$, $1 - \hat{\alpha} < c^* < \beta$). For these parameters, the SVF and BVF models have solution in Region IIIe, while the ABM has solution in Region III1. Time traces of $rep = 100$ ABM simulations (grey), solution to the SVF model (solid black), and solution to the BVF model (dashed black) with parameters $\theta = 0.15$, $\beta = 0.3$, and $c_1 = 1$ and with initial condition $r_0 = 0.2$.

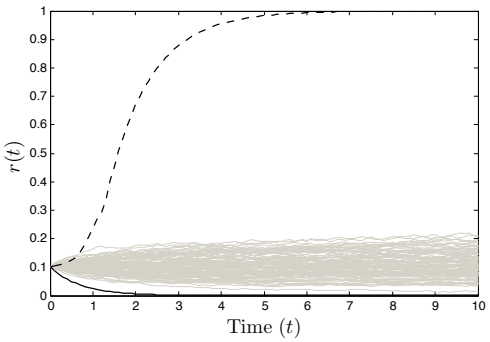

(a) Both SVF and BVF models fail to approximate the dynamics of the ABM on the Facebook social network (G_F, $1-\hat{\alpha}=0.163$) with parameters in Region III0 ($c_2=1.4$, $1-\hat{\alpha}<\beta\leq c^*\approx 0.417$).

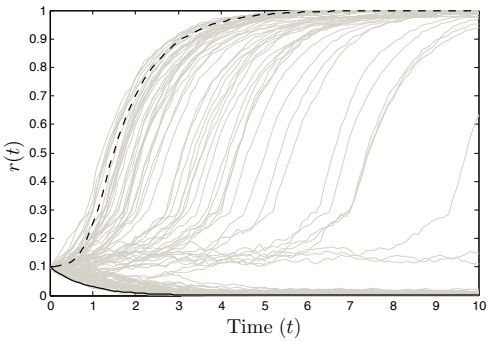

(b) The SVF and BVF models provide complementary approximations of the dynamics of the ABM on the physical contact network (G_P, $1-\hat{\alpha}=0.165$) with parameters in Region III0 ($c_2=1.2$, $1-\hat{\alpha}<c^*\approx 0.455<\beta$).

Figure 3.10: The SVF and the BVF models can (a) both fail to approximate the dynamics of the ABM, or (b) act as complementary approximations to the dynamics of the ABM. Time traces of $rep=100$ ABM simulations (grey), solution to the SVF model (black), and solution to the BVF model (dashed black), with parameters $\theta=0.15$, $\beta=0.3$, and $c_1=1$ and with initial condition $r_0=0.1$.

We observe four types of outcomes when comparing the ability of the BVF model to approximate the ABM to the ability of the SVF model to approximate the ABM:

1. The BVF model may succeed better in approximating the dynamics of the ABM than the SVF model, even if the improvement is only qualitative in nature, see Fig. 3.7 (right panels), 3.8, and 3.9 (left panel),

2. The BVF model may complement the dynamics predicted by the SVF model, see Fig. 3.10b,

3. The BVF model may produce qualitatively similar predictions to the SVF model, see Figs. 3.6, 3.7b (left panel) and 3.9 (right panel), or

4. Both the BVF and the SVF models fail to approximate the dynamics of the ABM either qualitatively or quantitatively, see Figs. 3.7b (left panels) and 3.10a.

For each of these outcomes, the approximation of the ABM by the BVF model is no worse than the approximation by the SVF model, and, in many cases the approximation

of the ABM by the BVF model is much better than the approximation by the SVF model. Therefore, we conclude that the BVF model, and hence the EVF model, represents an improvement over the SVF model with respect to their ability to approximate the behaviour of the ABM. This is no surprise, since the BVF/EVF model takes important information of the real network into account.

Optimizing α for the SVF model

We now outline the procedure we used in Figs. 3.6-3.10 for choosing an appropriate α in the SVF model (i.e. using $\nu(r_a;\theta) = v_s(r_a;\alpha)$ in Eq. (3.7)) for a given θ of the ABM. Specifically, we simulate the ABM on the physical contact and Facebook networks for a fixed set of model parameters and find $\hat{\alpha}$, the optimal value of α for Eq. (3.8) that minimizes the difference between the ABM realizations and the output of the SVF model (Eqs. (3.7)-(3.8)). We begin by fixing the initial condition $r_a(0) = r_0$[8] and the parameters $\theta_v \equiv \theta$, β, c_1, and c_2. To compare realizations of the ABM with the solution to the SVF model we use Gillespie's Algorithm (summarized in Appendix C.2) to simulate $rep = 100$ realizations of the ABM, recording r_a at each $t \in \{0, 0.01, 0.02, \ldots, 1\}$, and then find the α that minimizes the difference (in L_2 norm) between the ABM realizations and the solution to the SVF model. This optimization procedure is performed in Matlab using the command *fminsearch*, see Appendix C.3 for Matlab code. The results of this procedure for ABM simulations performed on the physical contact and Facebook social networks with parameters $r_0 = 0.25$, $\theta = 0.15$, $\beta = 0.3$, $c_1 = 1$, and $c_2 = 9$ are given in Fig. 3.11. For these parameters we find the optimal α is given by $\hat{\alpha} = 0.835$ and $\hat{\alpha} = 0.834$ for the Facebook and physical contact networks, respectively. Equivalently, the fitted value for the visibility threshold $1 - \alpha$ is given by $1 - \hat{\alpha} = 0.165$ and $1 - \hat{\alpha} = 0.166$, respectively. Holding the initial condition $r_0 = 0.25$ and parameters $\beta = 0.3$, $c_1 = 1$ and $c_2 = 9$ constant, we are able to repeat the above fitting procedure to find fitted values of α for $\theta \in \{0.105, 0.11, 0.115, \ldots, 0.25\}$. These results are displayed in Fig. 3.12 and illustrate that for a large range of individual thresholds θ the Facebook social network has a slightly lower fitted visibility threshold $1 - \hat{\alpha}$ than does the physical contact network.

Observe that, for the parameters employed in Figs. 3.11-3.12 ($\theta \in \{0.105, 0.11, 0.115, \ldots, 0.25\}$, $\beta = 0.3$, $c_1 = 1$, and $c_2 = 9$), the SVF model lies in the stable police state dynamical regime (Region III0, $c^* = c_1/(c_1 + c_2) \leq 1 - \hat{\alpha} < \beta$). This is critical to the fitting procedure, since the parameter α appears in the solution to the SVF model only when (a) the parameters α, β, c_1, and c_2 lie in Region III0, and

[8]Note that for simulations of the ABM an initial condition $r_a(0) = r_0$ is implemented by choosing $\lceil r_0 N \rceil$ nodes uniformly at random to be active.

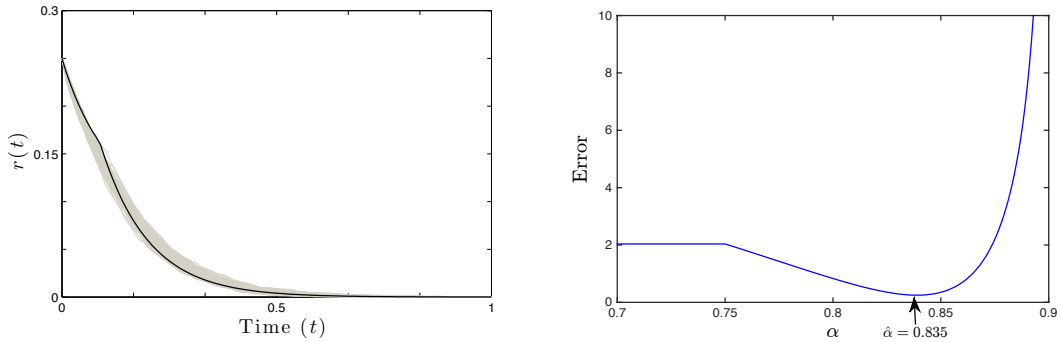

(a) Facebook social network (G_F, $\hat{\alpha} = 0.835$)

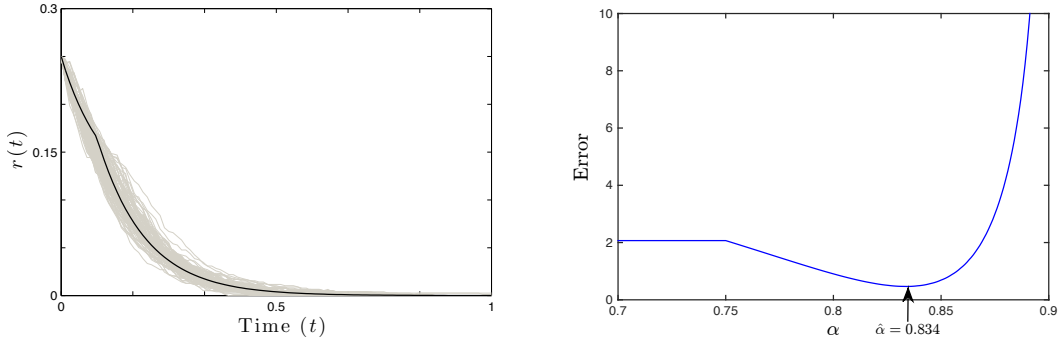

(b) Physical contact network (G_P, $\hat{\alpha} = 0.834$)

Figure 3.11: The optimal visibility parameter $\hat{\alpha}$ is found by fitting the SVF model to $rep = 100$ ABM simulations with parameters $\theta = 0.15$, $\beta = 0.3$, $c_1 = 1$, $c_2 = 9$, and $r_0 = 0.25$ (Region III0, $c^* \leq 1 - \hat{\alpha} < \beta$). (Left) Time traces of ABM simulations (grey) and solution to SVF model with optimally fitted $\hat{\alpha}$ (black). (Right) L_2^2 error between ABM realizations and the SVF solution versus α.

(b) $r_0 \in (1 - \alpha, \beta)$, see Section 2.3.2. Therefore, for a given θ, fitted values of $\hat{\alpha}$ used in attempts to approximate the ABM by the SVF model are computed using parameters $\beta = 0.3$, $c_1 = 1$, and $c_2 = 9$, and initial condition $r_0 = 0.25$, as in Figs. 3.11-3.12. These values of $\hat{\alpha}$ are also used for the SVF model when parameters lie in Regions IIIe and III1 and when the initial condition $r_0 \notin (1 - \alpha, \beta)$.

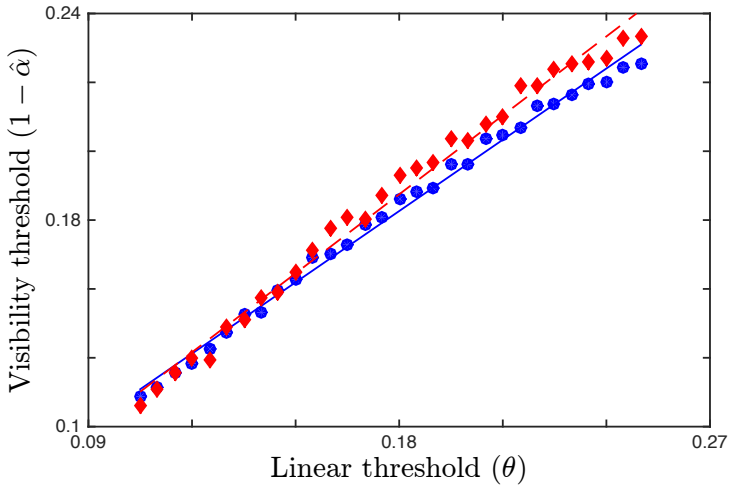

Figure 3.12: Result of finding optimal visibility parameter $\hat{\alpha}$ as a function of linear threshold θ by fitting Eq. (3.8) (the SVF) to $rep = 100$ ABM simulations with parameters $\beta = 0.3$, $c_1 = 1$, $c_2 = 9$, and $r_0 = 0.25$ (Region III0, $c^* \leq 1 - \hat{\alpha} < \beta$). (Blue circles) Facebook social network (G_F). (Red diamonds) Physical contact network (G_P). Line of best fit for networks given by solid ($m = 0.690 \pm 0.022$, $b = 0.058 \pm 0.008$) and dashed ($m = 0.764 \pm 0.032$, $b = 0.050 \pm 0.006$) lines, respectively (note: $\pm$ indicates 95% confidence interval).

3.5 A Higher Order ODE Model: The Degree Approximation (DA) Model

In the previous section we introduced the BVF/EVF models as extensions to the SVF model. In particular, we showed that the BVF/EVF models are better able to approximate the ABM than the SVF model. However, both the BVF/EVF and SVF models are one-dimensional ODEs. This section now compares the performance of the BVF/EVF models as approximations to the ABM to the performance of a higher order (i.e., higher dimensional) compartmental ODE model which we modify from [128]. We call this model the *degree approximation* (DA) model because it compartmentalizes individuals based on their state as well as their degree. We demonstrate numerically that the BVF model, and hence the EVF model, is no worse at approximating the aggregate behaviour of the ABM than the DA. This highlights the usefulness of the BVF/EVF model, especially since the DA is in practice much more computationally expensive to solve and much harder to analyze than the BVF/EVF model.

3.5.1 Degree Approximation Model Specification

Before we detail the approach which we modify from [128], it is useful to introduce some notation. We define V_k to be the set of nodes with degree k and N_k to be the set of nodes that have at least one neighbour of degree k,

$$\begin{aligned} V_k &= \{v \in V : d_v = k\}, \text{ and} \\ N_k &= \{w \in V : \exists v \in V_k \text{ such that } (v, w) \in E\}, \end{aligned}$$

respectively, where we denote the degree distribution of N_k by $\rho_{k,j}$, i.e.

$$\begin{aligned} \rho_{k,j} &= \frac{\text{number of nodes of degree } j \text{ that have at least one neighbour of degree } k}{\text{number of nodes that have at least one neighbour of degree } k} \\ &= \frac{|V_j \cap N_k|}{|N_k|}. \end{aligned}$$

As mentioned above, the degree approximation aggregates individuals according to their state and degree. Thus, conditioning on the state of the system at time t_0, we define $r^{(k)}(t) = r^{(k)}(t|t_0)$ as follows

$$r^{(k)}(t) = \frac{1}{|V_k|} \sum_{v \in V_k} \mathbb{E}\left[s_v(t)|t_0\right]$$

to be the fraction of nodes with degree k that are expected to be in state 1 at time t. Since nodes and edges are neither created nor destroyed, the fraction of nodes with degree k that are expected to be in state 0 at time t is given by $1 - r^{(k)}(t)$. Analogously, we define $r_d(t) = r_d(t|t_0)$ as follows

$$r_d(t) = \frac{1}{N} \sum_k |V_k| r^{(k)}(t) = \sum_k \rho_k r^{(k)}(t)$$

to be the fraction of nodes expected to be in state 1 at time t.

Applying the notation introduced above with the rules characterizing the dynamics of the process we find an equation analogous to Eq. (3.3)

$$\Delta r^{(k)}(t) = r^{(k)}(t) - r^{(k)}(t_0) = g^{(k)}(t|t_0) - d^{(k)}(t|t_0), \tag{3.12}$$

where the expected growth and decay of the fraction of active nodes is modelled by the nonnegative growth and decay functions $g^{(k)}(t) = g^{(k)}(t|t_0)$ and $d^{(k)}(t) = d^{(k)}(t|t_0)$, respectively. In order to close the degree approximation we need to approximate the quantities $g^{(k)}(t)$ and $d^{(k)}(t)$ in terms of $r^{(k)}(t)$, $\{\theta_v\}_{v \in V}$, and β.

For notational convenience we first suppress the time argument, and consider the case where $\theta_v \equiv \theta$. In order to approximate $g^{(k)}(t)$ or $d^{(k)}(t)$ in terms of $r^{(k)}$, θ, and β, we are now required to make an assumption about how the states of nodes sharing a common neighbour are correlated. Specifically, as for the BVF and EVF models, for any fixed $v \in V_k$ we assume that the states of any two neighbours of v are independent. This assumption implies that the probability of v having exactly j active neighbours is $\binom{k}{j}\bar{n}_k^j(1-\bar{n}_k)^{k-j}$, where $\bar{n}_k = \sum_l \rho_{k,l} r^{(l)}$ is the fraction of nodes in N_k with state 1. So, the probability of v having at least $\lceil \theta k \rceil$ active neighbours is

$$\sum_{j=\lceil \theta k \rceil}^{k} \binom{k}{j} \bar{n}_k^j (1-\bar{n}_k)^{k-j} = 1 - \text{BinCDF}(\lceil \theta k \rceil - 1; k, \bar{n}_k) = \text{BinCDF}(k - \lceil \theta k \rceil; k, 1-\bar{n}_k).$$

It follows that

$$g^{(k)}(t) = c_1 \ (1 - r^{(k)}(t_0)) \ \text{BinCDF}(k - \lceil \theta k \rceil; k, 1-\bar{n}_k)\Delta t + o(\Delta t). \tag{3.13}$$

As in Eq. (3.4), we choose

$$d^{(k)}(t) = c_2 \ r^{(k)}(t_0) \ p(r_d(t_0); \beta)\Delta t + o(\Delta t). \tag{3.14}$$

Combining equations (3.12)-(3.14) and dividing by Δt gives

$$\frac{\Delta r^{(k)}}{\Delta t} = c_1 \ (1 - r^{(k)}(t_0)) \ \text{BinCDF}(k - \lceil \theta k \rceil; k, 1-\bar{n}_k) - c_2 \ r^{(k)}(t_0) \ p(r^{(k)}(t_0); \beta) + o(1).$$

Now, taking the limit as $\Delta t \to 0$ yields

$$\frac{dr_k}{dt} = c_1(1 - r^{(k)})\text{BinCDF}(k - \lceil \theta k \rceil; k, 1-\bar{n}_k) - c_2 \ r^{(k)} \ p(r^{(k)}; \beta). \tag{3.15}$$

3.5.2 Comparison of ABM, BVF and Degree Approximation

In Section 3.4.3 we showed that the approximation of the ABM by the BVF/EVF model is no worse, and usually much better, than the approximation of the ABM by the SVF model. In this section we show that the approximation of the ABM by the BVF/EVF model is usually no worse than the approximation of the ABM by the DA given in Eq. (3.15). Specifically, we observe that the BVF and DA models produce similar qualitative predictions for most choices of parameters θ, β, c_1, and c_2, see for example Fig. 3.13. Also, although this is an atypical result, we present one set of parameters where the BVF model

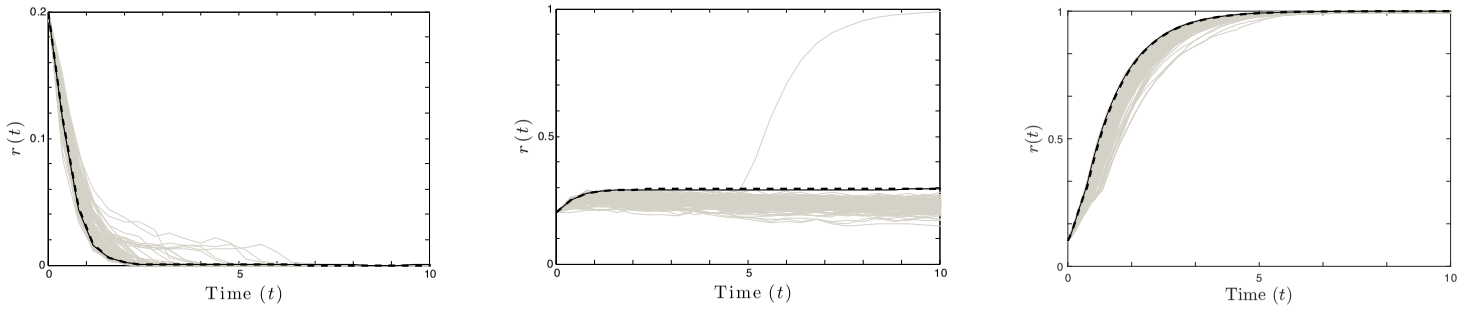

(a) Facebook social network (G_F). Parameters $c_1 = 1$, $\beta = 0.3$, and (Left) $\theta = 0.19$, $c_2 = 3.0$, $r_0 = 0.2$. (Middle) $\theta = 0.15$, $c_2 = 2.2$, $r_0 = 0.2$. (Right) $\theta = 0.11$, $c_2 = 1.0$, $r_0 = 0.1$.

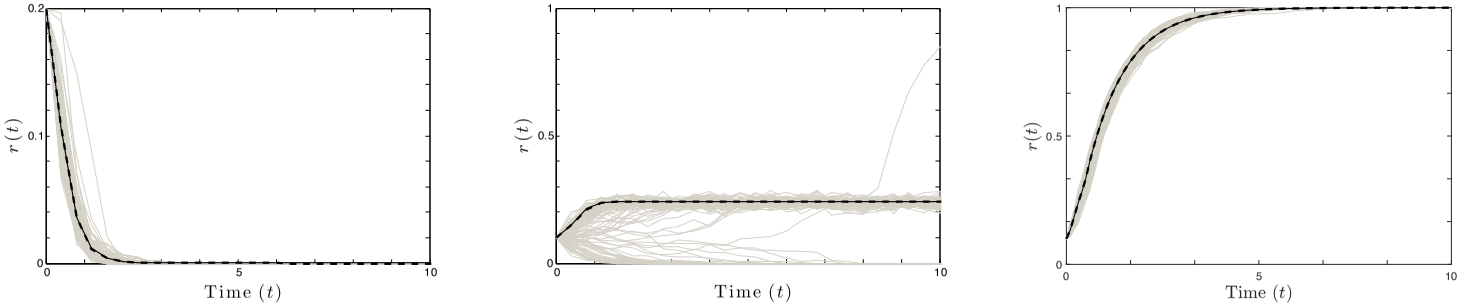

(b) Physical contact network (G_P). Parameters $c_1 = 1$, $\beta = 0.3$, and (Left) $\theta = 0.19$, $c_2 = 3.0$, $r_0 = 0.2$. (Middle) $\theta = 0.11$, $c_2 = 3.0$, $r_0 = 0.1$. (Right) $\theta = 0.11$, $c_2 = 1.0$, $r_0 = 0.1$.

Figure 3.13: The BVF model and the DA produce qualitatively similar predictions. From left to right ABM simulations die out (Region III0), reach an intermediate equilibrium (Region IIIe), or reach the fully realized revolution (Region III1). Time traces of $rep = 100$ ABM simulations (grey), solution to BVF model (dashed black) and DA (solid black).

outperforms the DA on the Facebook social network (G_F), and where the BVF model complements the DA on the physical contact network (G_P), see Fig. 3.14. We observe that, for the specific parameters and networks used in the simulations of Figs. 3.13-3.14, the DA does not significantly outperform the BVF model. It was confirmed in extensive additional simulations that this occurs generically for large parts of the parameter space of the models. This is worth emphasizing once more, since the BVF model is much easier to analyze and much less costly to solve than the DA. This supports our proposition that the BVF (or EVF) model is a powerful tool for approximating and analyzing the ABM and a general technique for formulating tractable models for spreading processes on social networks that take real network characteristics into account.

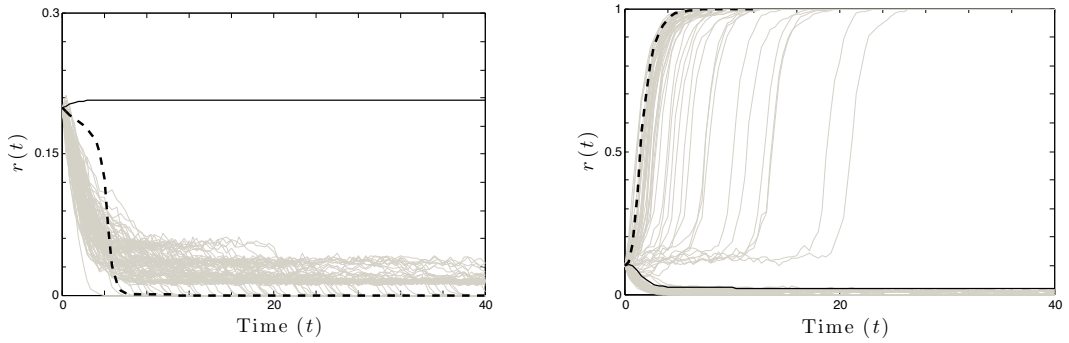

(a) Facebook social network (G_F) with parameters $c_1 = 1$, $c_2 = 3$, $\theta = 0.15$, and $\beta = 0.3$, with initial condition $r_0 = 0.2$.

(b) Physical contact network (G_P) with parameters $c_1 = 1$, $c_2 = 1.2$, $\theta = 0.15$, and $\beta = 0.3$, with initial condition $r_0 = 0.1$.

Figure 3.14: Example of parameters for which (a) The BVF model outperforms the DA on the Facebook social network (G_F), and (b) the BVF model and the DA complement each other on the physical contact network (G_P). Time traces of $rep = 100$ ABM simulations (grey), solution to the BVF model (dashed black) and the DA (solid black).

3.6 Extending the Basic Reproduction Number (R_0) from Epidemiology to the Linear Threshold ABM

In epidemiological models the *basic reproduction number* R_0 is defined to be the number of secondary infections caused by a single infected individual introduced into a population that is entirely susceptible [84]. Note that in this context exposure to one infected individual is

sufficient for a new infection. If $R_0 < 1$ then an infected individual is, on average, unable to replace himself and the outbreak terminates on its own in short order. On the other hand, if $R_0 > 1$ then an infected individual is, on average, more than able to replace himself and the outbreak may spread into a full-blown epidemic. We define the basic reproduction number for our linear threshold model of political revolution analogously: the basic reproduction number is the average number of individuals that become active in the revolution due directly to the introduction of a single active individual into a population that is otherwise completely inactive. In this section, we show that the basic reproduction number in our linear threshold ABM is related to the initial slope of the BVF/EVF by the equation

$$R_0 = c^* v_b'(0; \theta, \rho), \tag{3.16}$$

where $c^* = c_1/(c_1 + c_2)$, and discuss efficient ways to compute R_0. We remark that our attempt to derive a threshold value, R_0, that predicts whether or not a small initial protest will grow into a full scale revolution, is related to work done by [69, 187]. These works estimate final outbreak size for an activation-only linear threshold model (i.e. a linear threshold model where nodes, once active, cannot become inactive) on tree-like networks, incorporating information about the degree distribution and initial condition. We remark that, although these works present interesting and useful methodology, they cannot be directly applied to our linear threshold ABM model, since our model includes a mechanism for activated nodes to become inactive, and since our model is applied to social networks with high clustering coefficients.

We begin by deriving an expression for R_0. Suppose that we activate individual $v \in V$ in the population, and suppose that individual v has degree k. Further, suppose that this individual has $\mathcal{V}$ neighbours who can see the revolution (i.e., the linear threshold criterion is satisfied) once individual v is activated. Let $w \in V$ be a neighbour of v's that can see the revolution, and let τ_1 and τ_2 be the first arrival time of Poisson processes with rates c_1 and c_2, respectively. The probability that individual w becomes active before individual v becomes inactive is

$$\begin{aligned}
\mathbb{P}(\tau_1 < \tau_2) &= \int_0^\infty \mathbb{P}(\tau_1 < t | \tau_2 = t)\, \mathbb{P}(\tau_2 = t)\, dt \\
&= \int_0^\infty [1 - \exp(-c_1 t)] c_2 \exp(-c_2 t) dt \\
&= \frac{c_1}{c_1 + c_2} = c^*.
\end{aligned}$$

Since we only consider activations that result directly from v's activation (and not activations caused by a combination of the activation of v and the subsequent activation of

v's neighbours), for our purposes the behaviour of each of v's neighbours is independent. Thus, the expected number of v's neighbours which become active is $\mathcal{V}c^*$.

Recall that the social network $G = G(V, E)$ has N nodes, M edges, degree distribution ρ_k and secondary degree distribution $\rho_{k,j}$. For the single activated individual v with degree k the expected number of neighbours who can see the revolution is

$$\mathbb{E}[\mathcal{V}] = k \sum_j \rho_{k,j} \mathbb{I}_{\{j\theta \leq 1\}}.$$

Averaging over all possible degrees k now yields the basic reproduction number

$$R_0 = c^* \mathbb{E}\left[\mathcal{V}\right] = \frac{c_1}{c_1 + c_2} \sum_{k,j} k\rho_k \rho_{k,j} \mathbb{I}_{\{j\theta \leq 1\}} = c^* \sum_{k=1}^{\infty} \sum_{j=1}^{\lfloor \theta^{-1} \rfloor} k\rho_k \rho_{k,j}.$$

It is now possible to prove the relationship given in Eq. (3.16). Writing

$$R_0 = \frac{c_1}{c_1 + c_2} \sum_{j=1}^{\lfloor \theta^{-1} \rfloor} \sum_{k=1}^{\infty} k\rho_k \rho_{k,j},$$

it suffices to show that

$$j\rho_j = \sum_{k=1}^{\infty} k\rho_k \rho_{k,j},$$

since

$$\begin{aligned} v_b(r; \theta, \rho) &= \sum_{k=1}^{\infty} \rho_k \mathrm{BinCDF}(k - \lceil \theta k \rceil; k, 1 - r) \\ &= \sum_{k=1}^{\infty} \sum_{j=0}^{k - \lceil \theta k \rceil} \rho_k \binom{k}{j} (1-r)^j r^{k-j} \\ &= \sum_{k=1}^{\infty} \rho_k \left[\binom{k}{0} r^k + \ldots + \binom{k}{k - \lceil \theta k \rceil} (1-r)^{k - \lceil \theta k \rceil} r^{\lceil \theta k \rceil} \right] \\ \implies v_b'(0; \theta, \rho) &= \sum_{k=1}^{\infty} k\rho_k \, \mathbb{I}_{\{\theta k \leq 1\}} = \sum_{k=1}^{\lfloor \theta^{-1} \rfloor} k\rho_k. \end{aligned} \tag{3.17}$$

We note that together with Eq. (3.16), Eq. (3.17) gives an exact and cheap way to compute the basic reproduction number R_0.

From the undirected social network $G = G(V, E)$ from the directed social network $G' = G(V, E')$, where $E' = \{(v, w) : \{v, w\} \in E\}$. In this case, the number of edges emanating from all nodes of degree k is $k\rho_k N$. It follows that the number of edges emanating from nodes of degree k and incident on nodes of degree j is $k\rho_k N\rho_{k,j}$. Thus, the number of edges incident on nodes of degree j is

$$N \sum_{k=1}^{\infty} k\rho_k \rho_{k,j}.$$

Equivalently, the number of edges incident on nodes of degree j is

$$j\rho_j N.$$

Putting these two expressions together, we find

$$j\rho_j = \sum_{k=1}^{\infty} k\rho_k \rho_{k,j},$$

which completes the proof.

Intuitively, this makes sense since the number of nodes that can see the revolution when one node is activated can also be approximated by

$$N\, v_b\left(\frac{1}{N}; \theta, \rho_k\right) = \frac{v_b\left(\frac{1}{N}; \theta, \rho_k\right) - \overbrace{v_b\left(0; \theta, \rho_k\right)}^{=0}}{\frac{1}{N}} \approx v_b'\left(0; \theta, \rho_k\right).$$

3.7 Application: Linear Threshold Propagation on Online versus Offline social networks

In this section we investigate experimentally the differences in spreading behaviour that occur under the linear threshold model for the spread of political revolutions when applied to some empirical online and offline social networks, searching for some initial quantitative evidence that political revolutions may be facilitated by the network structure of online social networks of social media.

3.7.1 Network Structure: Online versus Offline social networks

It is often assumed that the connectivity of modern online social networks was an important factor in the spread of political revolutions in the past decade, e.g., in the Arab Spring revolutions of 2011 [50, 100, 102, 177], while traditional offline social networks (using in-person physical contact, or phone or mail interaction for safe communication) featured a different connectivity structure that was often severely restricted by censorship of the regime. Unfortunately, representative samples of the offline, traditional communication networks that were in existence prior to the adoption of new media technologies are unavailable at sufficient scale for countries affected by the Arab Spring, or indeed for any country. As opposed to the online networks of social media which are, by their nature, digitally stored and available, the offline social networks of pre-Internet societies have not been recorded at scale, simply because it was impractical in terms of cost and effort. This is a serious roadblock when investigating the effects of the structure of new media networks on the dynamics of political revolution, and if one wants to compare with propagation on pre-Internet social networks, it is necessary to identify proxy networks that are likely to be reasonable approximations to pre-Internet social networks, in terms of network structure.

In this section, we use simulations of our ABM model on two empirical networks as a starting point to investigate differences in propagation properties that may arise between online and offline social networks within the linear threshold propagation model. We choose the small physical contact network G_P between individuals from [158] as a representative for offline social networks, and we choose the Facebook network G_F as a representative for online social networks.

3.7.2 Linear Threshold Agent-Based Model on Facebook Social G_F and Physical Contact G_P Networks

We now briefly examine the differences in how the political revolution spreads on the Facebook social and physical contact networks via direct simulation of the ABM. In particular, choosing parameters $\theta = 0.1$, $\beta = 0.3$, $c_1 = 1$, and $c_2 = 0.2$ in the unstable police state region (Region III1, $1-\hat{\alpha} < \beta \leq c^*$), we simulate $rep = 100$ realizations of the ABM for each initial condition $r_0 \in \{0, 0.0015, 0.003, \ldots, 0.0495\}$ from time $t = 0$ until time $t = 20$. The average final size (at $t = 20$) for each initial condition $\langle r_{final}|r_0\rangle = \langle r_a(20)\rangle|_{r_0}$ is recorded and displayed in Fig. 3.15, which demonstrates that for these parameters the political revolution described by our ABM propagates to a much greater extent on the Facebook social network, i.e. the online proxy network, than on the physical contact network, i.e. the offline proxy network.

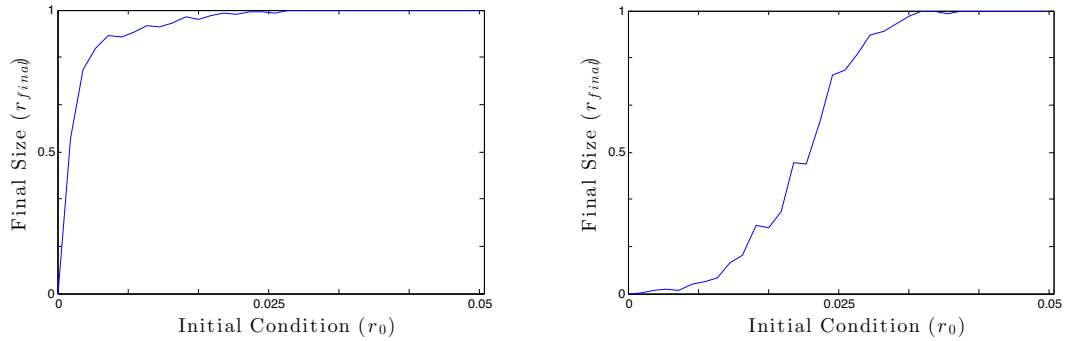

(a) Facebook social network (G_F) (b) Physical contact network (G_P)

Figure 3.15: Average final size r_{final} of $rep = 100$ realizations of ABM simulations on Facebook social (G_F) and physical contact (G_P) networks with $r_0 \in [0, 0.05]$, $\theta = 0.1$, $\beta = 0.3$, $c_1 = 1$, and $c_2 = 0.2$ (Region III1, $1 - \hat{\alpha} < \beta \leq c^*$). Realizations are allowed to run until time $t_{final} = 20$. Basic reproduction numbers R_0 for Facebook social and physical contact networks are $R_0 = 1.12 > 1$ and $R_0 = 0.35 < 1$, respectively.

As an important illustration of the usefulness of the basic reproduction number R_0 we defined in Section 3.6, we find that this is consistent with a difference in R_0 for the two networks: the basic reproduction number is $R_0 = 1.12$ (above the value of 1 which is expected to be required for propagation) for the Facebook network, and $R_0 = 0.35 < 1$ for the physical contact network.

These results indicate that the offline social network is less conducive to spreading the revolution in the ABM than the online social network: it has a smaller basic reproduction number, and in simulations a larger initial population of revolutionaries is required to spread the revolution. This is consistent with the initial slopes of the BVF/EVF illustrated in Fig. 3.5: the initial slope of the BVF/EVF for the Facebook network is always greater than that of the physical contact network. This provides some initial quantitative evidence that the spread of revolutions under a linear threshold process may occur more easily on modern online social networks than on traditional offline networks. While this is an interesting first observation, the next section discusses limitations and further investigations that are required to address this intriguing but complex question more comprehensively.

3.8 Discussion and Conclusion

In this chapter we developed a linear threshold agent based model (ABM) to model the spread of a political revolution in a dictatorial regime. We showed that this model is consistent with the previous simple step visibility function (SVF) ODE model developed in Chapter 2 [109]. The SVF model is useful when little information is known about the underlying network structure. For example, without having to specify an underlying communication network, the SVF model provides a potential mechanism by which networks with increased visibility α are more susceptible to political revolution. Using the relationship between these two models as a template we developed a hierarchy of models of varying complexity that approximate the behaviour of the ABM, see Table 3.1 and Fig. 3.1. Of all the models we have identified, we find that the BVF and EVF models (models of moderate complexity) offer the optimal combination of low computational complexity (cost), ease of analysis, and ability to approximate the behaviour of the ABM. Specifically, we find that for most parameters and initial conditions the BVF/EVF model is better able to approximate average ABM behaviour than the SVF. Also, for most networks the BVF/EVF model is much less costly to solve and much easier to analyze than the degree approximation from [128]. Importantly, the analysis of the simple ODE models in terms of stability of solutions for various parameter regimes directly gives insight in the qualitative dynamics of the linear threshold ABM for the spread of political revolutions.

We extended the concept of the basic reproduction number R_0 from epidemiology [84] to the linear threshold ABM, we showed how it is related to the slope of the empirical or binomial visibility functions at $r = 0$, and we provided efficient ways to compute or estimate it. Analogously to epidemiological models, when $R_0 > 1$ we expect the political revolution to spread, and when $R_0 < 1$ we expect it to die out. Thus, computing this quantity for a network can give an indication of how the ABM will behave on that network without the need to perform simulations, as we have demonstrated for empirical networks.

The Facebook and physical contact networks we consider as case studies, i.e. the online and offline proxy networks, provide initial support to the hypothesis that the adoption of online social media may facilitate the spread of political revolutions by effectively changing the connectivity structure of the population in a way that makes linear threshold spreading more effective. Specifically, we find that for certain parameters the online proxy network is more susceptible to the linear threshold spreading process than the offline proxy network. Moreover, we find that the different behaviour of these two networks is consistent with the basic reproduction number calculated for these two networks.

In addition to studying the ABM and its approximations on more and larger online and offline networks, much work remains to be done in the actual modelling of the political

revolution process. For example, throughout this manuscript we have assumed that individuals are either active or inactive, i.e. there are not intermediate states of involvement in the revolution, that the linear threshold of individuals is constant for the population, that the underlying communication networks are static, that the graph is undirected, and that the nodes that are initially activated are chosen uniformly at random. Each of these assumptions represents a major simplification of reality that needs to be addressed to study further aspects of the spread of political revolutions on online social networks. For example, models for the formation of extreme positions have been proposed where individuals can take on a continuum of states [150]. Or, consider that in our model individuals estimate the current participation in the revolution by sampling their neighbours. Since individuals with greater sample size, i.e. larger degree, can form more accurate estimates of the current participation in the revolution, they should be willing to join the revolution at a lower linear threshold θ than individuals who are more uncertain in their estimate of revolution size. We would therefore expect the linear threshold to vary from individual to individual as a function of their network degree. Or, consider that one of the principal strengths of new media is the ability of individuals to search for both content and like-minded individuals. Thus, we would expect that the underlying communication network should be changing on the same time scale as the revolution. These changes in the nature of the communication network may be effectively modelled by considering revolutions/cascades on multiplex [90], interdependent [54], or geographic [22] networks. Finally, the growing availability and analysis of real world data generated from protest movements [72] and information diffusion processes [139] that identify characteristics of seed and "super-spreader" nodes may shed light on what initial conditions may be the most appropriate. For example, one might consider that the nodes that are predisposed to be active at the initial stages of a revolution, e.g. nodes that represent activists, may also have larger degree (if, for example, they are charismatic and allowed to accrue followers) or smaller degree (if, for example, they are the specific target of regime censorship). There is significant value in the kind of parsimonious ABM model on static networks that we have considered in this chapter, because much can be learned from this type of model and it is easier to analyze and interpret than more complex models. Nevertheless, extensions along the lines sketched above are important avenues for further study.

Chapter 4

The Influence of Societal Individualism on a Century of Tobacco use: Modelling the prevalence of smoking

This chapter is based on research published in *BioMed Central (BMC) Public Health* [108].

4.1 Background and Motivation

In the fifty years since the first report of the Surgeon General's Advisory Committee on Smoking and Health [181] the smoking epidemic has been responsible for more than 20 million deaths in the United States alone [180, 186], and continues to be responsible for over 6 million deaths worldwide each year [95, 119]. The strong social component of the dynamics of smoking prevalence has been modelled mathematically [183, 29, 161, 111, 156], and examined statistically through analysis of social network data [40] and survey data [88, 179, 94]. However, whereas previous works tend to focus on the micro-level, in this chapter we investigate how social aspects of smoking affect its prevalence at the societal level.

Significant inter-country differences exist in smoking prevalence [133]. For example, Fig. 4.1 shows smoking prevalence estimates over most of the past century for Sweden and the USA, obtained from surveys and cigarette consumption data (see Section 4.4.1).

In both countries, smoking prevalence increased rapidly starting from the early decades of the 20$^{\text{th}}$ century and reached a peak in the 1960s–1980s era when the adverse health effects of smoking became widely known [181], after which smoking prevalence declined rapidly. However, there are conspicuous differences between the curves: the rate of smoking adoption and cessation before and after the peak is much greater in the US than in Sweden, and the peak in prevalence in the US occurs much earlier than in Sweden.

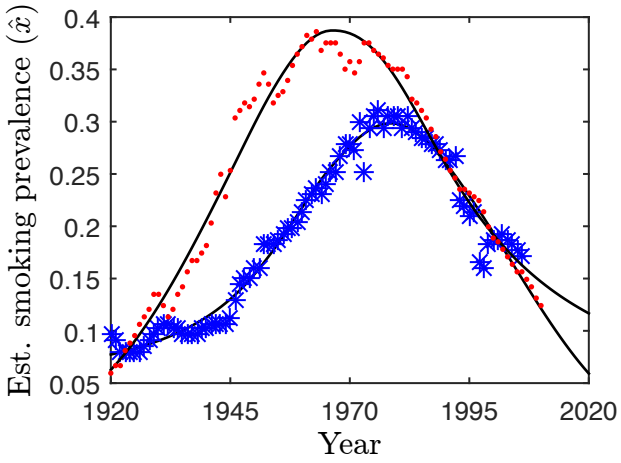

Figure 4.1: Estimated smoking prevalence $\hat{x}$ versus time for the United States (red dots) and Sweden (blue asterisks). The solid lines give the curves of best fit for Eq. (4.1).

Considerable time and resources have been devoted to identifying the factors that contribute to smoking prevalence. Major factors include differences in beliefs about the harm of smoking [45], socio-economic status [13, 91], cost [67], regulation/tobacco control policies [171, 127, 96], and gender [134]. However, we note that these advances in the understanding of the factors contributing to smoking prevalence are based primarily on micro-level data, methods that inform general hypotheses, and non-mathematical descriptive models. Indeed, comprehensive and quantitative cross-national analyses of how all these factors affect smoking prevalence are rare [133]. Existing studies that compare national trends in smoking prevalence, as well as the factors that contribute to these trends, tend to take a descriptive [115, 118] and/or statistical [133] approach, and do not address the mechanism underlying the key decision of whether or not to smoke in a quantitative manner [132].

In this chapter we present a new model for the social spreading of smoking. We aim to create and test a *tractable mathematical model*, that is, a model for qualitative dynamics from which insight (including causation) can be drawn. This differs from the statistically-driven approach often used in areas such as econometrics and medicine, where correlations

may be uncovered and analyzed without formulating first-principle-based dynamic mathematical models. The statistical approach is difficult to apply here because the amount of available data on historical smoking dynamics is small. Our model-based approach has much in common with simple explanatory mathematical models that have been successful in, e.g., epidemiology and population dynamics.

Our model incorporates the concepts of *individual utility from smoking*, i.e. the utility an individual derives directly from the act of smoking (including awareness of health effects), and *social utility from smoking*, i.e. the utility an individual derives indirectly from smoking through social interactions with other smokers (peer influence and social inertia). Together these two quantities determine the *total utility from smoking*. Our model assumes that an individual's decision to smoke is based on the desire to maximize total utility. By invoking this decision-making mechanism in a simple mathematical model, our approach differs from the approaches of the previous mathematical [29, 161, 111, 156] and descriptive/statistical [133, 115, 118, 132] models. Whereas previous mathematical models generally require the calibration of many parameters (leading to difficulties in analysis, interpretation, and overfitting), we propose a simple approach based on principles of social psychology and sociology whose predictions can be directly compared to smoking prevalence data. Whereas previous descriptive and statistical models lack an underlying decision-making mechanism, we propose a model with a decision-making mechanism that is capable of incorporating factors previously identified as contributing to smoking prevalence. Specifically, we note that monetary cost, beliefs about the harm/health effects of smoking, and regulation/tobacco control policies are all implicitly accounted for in the concept of individual utility from smoking. Our simple model applies to the population level, focusing on major effects that may influence the temporal dynamics of smoking across societies. It proposes a mechanism for smoking adoption and cessation that hinges on the balance between individual and social utility (which both encompass other more fine-grained factors). Matching the model to real-world data reveals that the balance between social and individual utility indeed is an important factor in the temporal dynamics of smoking, differentiating between countries in a way that is consistent with known measures of societal individualism. This lends support to the compelling hypothesis that the balance between individual and social utility, which we will show to be related to societal individualism, is indeed an important society-level driver for the temporal dynamics of smoking prevalence. This is consistent with previous findings that the level of individualism/collectivism of a society may have fundamental implications for its biology [38, 57], as well as its behaviour [88, 37, 59, 148, 166].

The model we propose is explained in Section 4.2 and analyzed in Section 4.3. Section 4.4 describes the data and parameter fitting procedure used in calibrating the model.

In the context of societal individualism/collectivism, the parameter in our model that controls the relative importance of individual versus social utility (see *relative conformity parameter*, defined in Section 4.2) is interpreted as follows: the greater the relative contribution of individual utility to total utility (at the expense of social utility), the more individualistic the society is interpreted to be. Conversely, the greater the relative contribution of social utility to total utility (at the expense of individual utility), the more collectivistic the society is interpreted to be. As described in detail in Section 4.5, this allows us to test the model's predictions against independently collected smoking prevalence and individualism/collectivism data sets in three separate phases, see Fig. 4.2. First, we compile smoking prevalence data spanning the past century for seven countries belonging to the Organization for Economic Co-operation and Development (OECD) and find good agreement between this data and the fitted model (see Phase (i) in Fig. 4.2 and Section 4.5.1). Second, the country-specific parameter in our model that controls the relative importance of individual versus social utility, i.e. the parameter that we interpret as the degree of societal individualism/collectivism (see Section 4.2), and that we fit to smoking prevalence data, is found to be significantly correlated to an established measure of societal individualism for each country (Hofstede's IDV [86]), in agreement with the predictions of the model (see Phase (ii) in Fig. 4.2 and Section 4.5.2). Thirdly, the central role played by societal individualism/collectivism in our model motivates us to investigate directly the role that individualism (as measured by Hofstede's IDV) plays in observed historical tobacco use data. Our model predicts that more individualistic societies will show faster adoption and cessation of smoking. We investigate this in historical tobacco use data, and find that IDV is significantly correlated to the average rate of increase in smoking prevalence (s_x) in seven OECD countries for which historical smoking prevalence data is available, and that it is significantly correlated to the peak year of tobacco consumption (t_{max}) for 25 countries in which tobacco consumption data are available, in agreement with model predictions (see Phase (iii) in Fig. 4.2 and Section 4.5.3). These findings are interpreted according to our modelling framework, and provide evidence for the compelling hypothesis that individualism/collectivism has an important influence on the dynamics of smoking prevalence at the aggregate, population level. We next apply our model to investigate the sensitivity of the model to the relative conformity parameter, i.e. to the parameter that controls the relative importance of individual versus social utility. Specifically, Section 4.6 considers how the smoking epidemic in the United States might have evolved if the United States were slightly less individualistic, i.e. if the Unites States relative conformity parameter was equal to that of the United Kingdom.

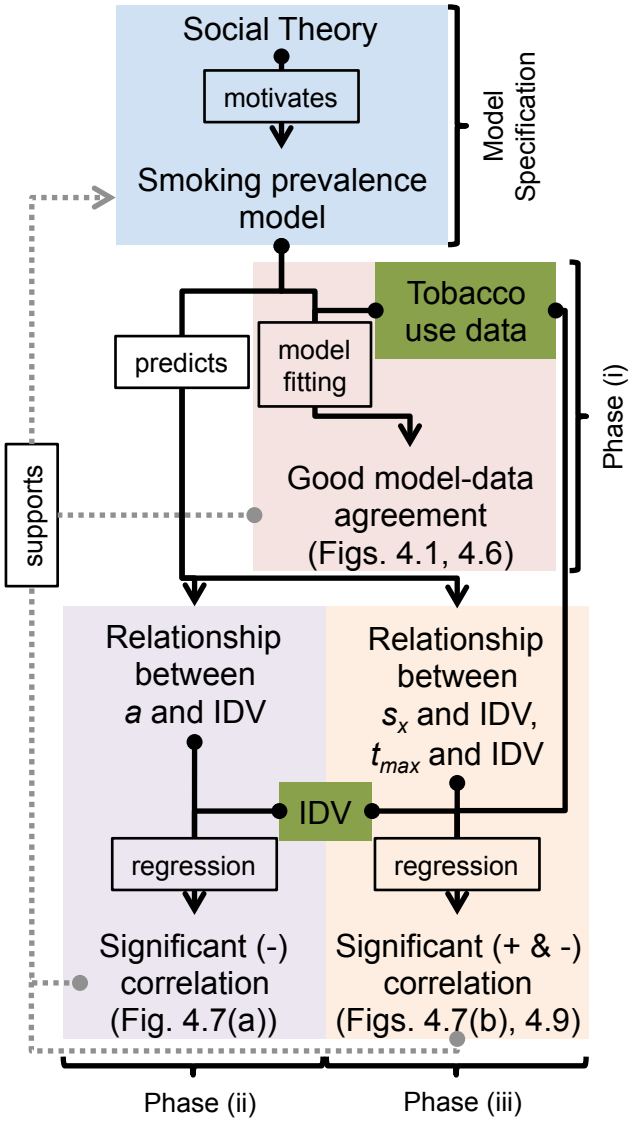

Figure 4.2: Schematic representation of the development and testing of the smoking prevalence model given in Eqs. (4.1)-(4.2). Logic flows from filled circles to arrow heads. (Green rectangles) Data sets used in this study. (Blue rectangle) Model Specification, see Section 4.2. Model predictions are tested in three phases. (Red rectangle) Phase (i), (purple rectangle) Phase (ii), and (yellow rectangle) Phase(iii), see Sections 4.5.1-4.5.3, respectively.

We discuss the impact of confounding variables, model limitations, and possible areas for future research in Section 4.7. We conclude this chapter in Section 4.8, which summarizes our findings, and also briefly discusses the implications of our model with respect to tobacco control policies and to the broader phenomenon of cultural influence on social spreading processes.

4.2 Model Specification

We begin formulating our model by observing that individuals derive utility from smoking via two mechanisms. First, they derive utility directly from the act of smoking (*individual utility*). Second, they derive utility from social interaction with other smokers (*social utility*). We note that social utility commonly manifests itself in the form of peer influence or peer pressure [26, 36]. We then proceed using a modelling framework that explicitly accounts for the effect of competition between individual and social utilities, and that was first applied to explore the temporal dynamics of language death and religious affiliation as binary choice problems [1, 2]. Specifically, we propose the model

$$\frac{dx}{dt} = b\left[(1-x)\ x^a u_x - x\ (1-x)^a(1-u_x)\right], \tag{4.1}$$

where $x = x(t) \in [0,1]$ is the fraction of smokers in the population (i.e., the prevalence) at time t, $u_x \in [0,1]$ is the individual utility from smoking, and the constant $b > 0$ determines the timescale of the equation. The interpretation of the positive term in Eq. (4.1) is, therefore, that non-smokers $1-x$ take up smoking at a rate proportional to the *total utility* derived from smoking, $x^a u_x$, which is the weighted product of the individual utility from smoking u_x and the social utility from interactions with other smokers x, with weighting determined by the constant parameter $a > 0$. Observe that the ratio of marginal total utility from social utility and the marginal total utility from individual utility is

$$\frac{\frac{d}{dx}[x^a u_x]}{\frac{d}{du_x}[x^a u_x]} = a\frac{u_x}{x}.$$

Since societies with large a weigh changes in social utility more heavily than changes in individual utility when calculating total utility, we call a the *relative conformity parameter*. We therefore interpret societies with large a to be more collectivistic (or less individualistic) than societies with small a. The interpretation of the negative term in Eq. (4.1), which models smoking cessation, follows analogously: smokers x cease smoking at a rate proportional to the total utility derived from non-smoking, $(1-x)^a(1-u_x)$, which is the

weighted product of the individual utility from non-smoking $u_y = 1 - u_x$ and the social utility from interactions with other non-smokers $1-x$, where we have normalized individual utilities from smoking u_x and from non-smoking u_y such that $u_x, u_y \geq 0$ and $u_x + u_y = 1$. We note that this modelling framework is conceptually consistent with the findings presented in [88]: that personal attitudes about smoking have a stronger influence on smoking behaviour in individualistic countries than in collectivistic countries.

Next, we observe that a combination of factors, including advances in our understanding of the health effects of smoking and public policy initiatives designed to curb smoking, have likely reduced individual utility from smoking (u_x) over the past century. Thus, in a significant departure from previous work that treats individual utility as a constant [1, 2], we account for this decline in individual utility by using the cumulative number of scholarly articles on the health effects of smoking ($n(t)$) as a proxy for the reduction in individual utility over the past century. Specifically, following the principle of temporal discounting [66], we assume that each additional article published is discounted by the factor $\delta \in (0, 1)$ so that for year t

$$u_x(t) = u_\infty + \delta^{n(t)} \, (u_0 - u_\infty), \tag{4.2}$$

where u_0 and u_∞ are the limiting individual utilities from smoking when there is no knowledge and perfect knowledge of the adverse effects of smoking, respectively. We remark that since individual utility should be decrease with increasing knowledge of the effects of smoking, we expect to find $u_\infty < u_0$. Since we normalized $u_x, u_y \geq 0$ and $u_x + u_y = 1$, this further implies that $0 \leq u_\infty < u_0 \leq 1$. Here, u_0, u_∞ and δ are parameters to be fitted to observational data.

4.3 Model Analysis

We now present a brief analysis of the behaviour of the model specified by Eqs. (4.1)-(4.2). Throughout this section we assume that $a, b > 0$, $\delta \in (0, 1)$, $x_0 = x(t_0), u_0, u_\infty \in [0, 1]$, and $n(t) \geq 0$ is a monotonically increasing function. Given these assumptions, we observe that Eqs. (4.1)-(4.2) have at least two equilibria: $x = 0$ and $x = 1$.

4.3.1 Case 1: Relative conformity parameter $a = 1$

We start by considering the case where $a = 1$. In this case Eqs. (4.1)-(4.2) become

$$\frac{dx}{dt} = b \; x(1 - x) \; \left[2u_\infty - 1 + 2\delta^{n(t)} \, (u_0 - u_\infty)\right], \tag{4.3}$$

a separable ordinary differential equation with solution

$$x(t) = \frac{\exp\left(b(2u_\infty - 1)(t - t_0) + 2b(u_0 - u_\infty)\int_{t_0}^{t} \delta^{n(\tau)} d\tau\right)}{\exp\left(b(2u_\infty - 1)(t - t_0) + 2b(u_0 - u_\infty)\int_{t_0}^{t} \delta^{n(\tau)} d\tau\right) + \frac{1-x_0}{x_0}}.$$

We observe that

$$\lim_{t\to\infty} x(t) = \begin{cases} 0 & \text{if } u_\infty < \frac{1}{2} \\ x_\infty(x_0) & \text{if } u_\infty = \frac{1}{2} \\ 1 & \text{if } u_\infty > \frac{1}{2} \end{cases},$$

where

$$x_\infty(x_0) = \lim_{t\to\infty} \frac{\exp\left(2b(u_0 - u_\infty)\int_{t_0}^{t} \delta^{n(\tau)} d\tau\right)}{\exp\left(2b(u_0 - u_\infty)\int_{t_0}^{t} \delta^{n(\tau)} d\tau\right) - \frac{x_0-1}{x_0}}.$$

4.3.2 Case 2: Relative conformity parameter $a \neq 1$

We next consider the case where $a \neq 1$. We define the function

$$x^*(t) = \frac{\left(\frac{u_x(t)}{1-u_x(t)}\right)^{\frac{1}{1-a}}}{1 + \left(\frac{u_x(t)}{1-u_x(t)}\right)^{\frac{1}{1-a}}},$$

where $u_x(t)$ is specified in Eq. (4.2). Next, we observe from Eq. (4.1) that

$$0 = \frac{dx}{dt} = b\left[(1-x)x^a u_x - x(1-x)^a(1-u_x)\right] \iff x = 0, x = 1, \text{ or } x = x^*(t).$$

Using this information we sketch the phase diagram for Eqs. (4.1)-(4.2) in Fig. 4.3. Figure 4.3 also illustrates how Eqs. (4.1)-(4.2) might be capable of reproducing solutions qualitatively similar to smoking prevalence profiles illustrated in Fig. 4.1. Specifically, we observe that if a solution $x(t)$ to Eqs. (4.1)-(4.2) has a maximum value, then that maximum occurs at the intersection of the curves $x(t)$ and $x^*(t)$.

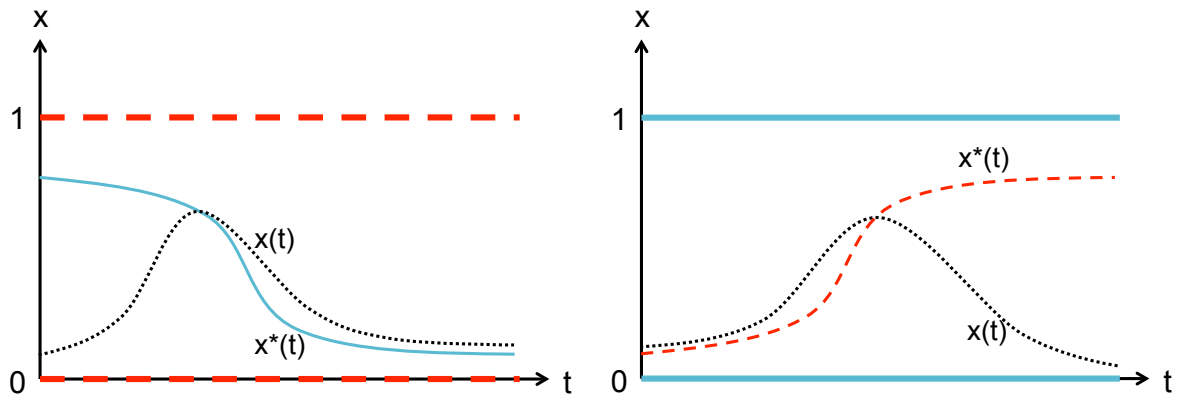

Figure 4.3: Phase diagram for the smoking prevalence mode given by Eqs. (4.1)-(4.2). If a solution $x(t)$ to Eqs. (4.1)-(4.2) (dotted black lines) has a maximum, then its maximum occurs at the intersection of the curves $x(t)$ and $x^*(t)$. (Left) When $a < 1$ we have that $x = 0$ and $x = 1$ are unstable equilibria (dashed red lines) and the vector field points towards $x^*(t)$ (solid blue line). (Right) When $a > 1$ we have that $x = 0$ and $x = 1$ are locally asymptotically stable equilibria (solid blue lines) and the vector field points away from $x^*(t)$ (dashed red line).

4.4 Methods

4.4.1 Data

We note that Eq. (4.1) subject to Eq. (4.2) requires the fitting of four parameters per country ($x_0 = x(t_0)$, a, u_0, and u_∞) and two parameters b and δ that we take equal for all countries in the data set (see Section 4.4.2). We determine these parameters by fitting them to estimated historical smoking prevalence data and proxy data on the health effects of smoking. We summarize the methods used to obtain these data below.

Smoking Prevalence and Cigarette Consumption Data

We consider smoking prevalence $x(t) \in [0, 1]$ for 24 OECD countries which we download from the OECD iLibrary online statistical database [130] in Excel format. We also consider manufactured cigarette consumption (in grams) per person per day $c(t)$ for the same 24 OECD countries plus Romania (which is a non-OECD country) [63, 64]. When available, cigarette consumption data is downloaded directly from the International Smoking

Statistics (Web Edition) website [63] in Excel format. Cigarette consumption data for countries not included in the International Smoking Statistics (Web Edition) are retrieved from the International Smoking Statistics (2nd Ed.) [64] by manually transferring these entries into Excel. These data are summarized in Table 4.1. We make these data available in *Additional file 2* of [108], see Appendix D.1.

Unfortunately smoking prevalence data is limited to, on average, only 21.5 observations over a period of 31.4 years spanning 1960–2012 [130]. As such, it misses much of the crucial period in the earlier parts of the 20th century during which smoking steadily gained popularity in many countries. However, historical national cigarette consumption data is available for the same 24 OECD countries plus Romania for an average of 78.4 observations over a period of 82.2 years spanning 1900-2012 [63, 64]. Since our model is specified in terms of smoking prevalence, we estimate smoking prevalence from cigarette consumption in order to exploit the much richer cigarette consumption data for model fitting purposes. First, we assume a linear relationship between smoking prevalence $x(t)$ and smoking consumption $c(t)$

$$x(t) = Cc(t) + B. \tag{4.4}$$

Next, we calculate estimates $\widehat{C}$ and $\widehat{B}$ by regressing smoking prevalence $x(t)$ on tobacco consumption $c(t)$ for all years for which both measurements are available. The results of this regression are summarized in Table 4.2, which illustrates that the assumption that x and c are linearly related does not hold equally well for all countries. In order to restrict ourselves to the cases where the assumption of linearity between x and c is valid, we restrict ourselves to the seven OECD countries with $R^2 \geq 0.7$, $p < 0.001$, and $n_{obs} \geq 15$: Australia, Canada, France, New Zealand, Sweden, the United Kingdom, and the United States. We display the raw data for these seven OECD nations in Fig. 4.4. The smoking prevalence for these seven OECD countries is then estimated from tobacco consumption using the relationship

$$\hat{x}(t) = \widehat{C}c(t) + \widehat{B}. \tag{4.5}$$

Note that survey-based prevalence data are susceptible to noise stemming from variations in the survey methodology. In particular, prior to performing the linear regression of x on c for France, we removed the outlier $x(1960) = 0.32$ since it is inconsistent with the rest of the data for France, see Fig. 4.4(c). Specifically, the Grubbs test on $x/\hat{x}$ indicates that the 1960 data point is a significant outlier ($p < 0.05$). This can also be seen intuitively: from $t = 1960$ until the next measurement at $t = 1965$ smoking prevalence drops from $x(1960) = 0.32$ to $x(1965) = 0.25$ (a decrease of 21.9%), while cigarette consumption steadily increases from $c(1960) = 3.6$ to $c(1965) = 4.1$ (an increase of 13.9%). Given the population in France in 1960 (45.5 million) and in 1965 (48.6 million) [129], this would

Table 4.1: Summary of data on smoking prevalence x and cigarette consumption c.

			Smoking prevalence (x)		Cigarette consumption per person per day (c)		
No.	Country	Abbrev.	Obs. Period	No. of Obs.	Obs. Period	No. of Obs.	Source
1	Australia	AUS	1964–2010	16	1920–2010	91	[63]
2	Austria	AUT	1972–2006	5	1923–2004	82	[63]
3	Belgium	BEL	1997–2008	4	1921–2011	91	[63]
4	Canada	CAN	1964–2011	29	1920–2010	91	[63]
5	Denmark	DNK	1970–2010	41	1920–2010	91	[63]
6	Finland	FIN	1978–2011	34	1920–2009	90	[63]
7	France	FRA	1960–2010	22	1900–2010	93	[63]
8	Greece	GRE	1998–2009	6	1920–1995	76	[64]
9	Hungary	HUN	1994–2009	4	1920–2012	87	[63]
10	Iceland	ICE	1987–2012	26	1932–1995	64	[64]
11	Ireland	IRE	1973–2007	14	1920–1995	76	[64]
12	Israel	ISR	1996–2010	8	1967–1995	29	[64]
13	Italy	ITA	1980–2012	23	1905–2010	73	[63]
14	Japan	JPN	1965–2011	47	1920–2007	88	[63]
15	Netherlands	NLD	1966–2011	39	1923–1995	67	[64]
16	New Zealand	NZL	1976–2012	28	1920–2009	90	[63]
17	Norway	NOR	1973–2012	40	1927–2011	85	[63]
18	Poland	POL	1996–2009	4	1925–1995	43	[64]
19	Portugal	PRT	1987–2006	4	1940–1995	56	[64]
20	Romania	ROM	–	0	1920–1995	52	[64]
21	Spain	SPA	1985–2011	11	1920–2010	87	[63]
22	Sweden	SWE	1980–2011	32	1920–2006	87	[63]
23	Switzerland	CHE	1992–2007	4	1934–2009	76	[63]
24	United Kingdom	GBR	1960–2010	38	1905–2009	105	[63]
25	United States	USA	1965–2011	36	1920–2010	91	[63]

Table 4.2: Result from Eq. (4.4) regression of smoking prevalence x on cigarette consumption c.

Country	$\widehat{C} \times 10^2$	$\widehat{B} \times 10^2$	R^2	p	n_{obs}
Australia	4.5 ± 1.3	-0.3 ± 8.8	0.80	3.2×10^{-6}	16
Austria	0.0 ± 4.9	24.2 ± 32.4	0.00	0.99	4
Belgium	2.6 ± 20.3	13.0 ± 81.5	0.13	0.64	4
Canada	3.5 ± 0.5	6.3 ± 3.8	0.87	3.0×10^{-13}	28
Denmark	0.0 ± 9.2	40.5 ± 44.4	0.00	0.99	41
Finland	2.0 ± 0.7	15.8 ± 2.8	0.55	1.0×10^{-6}	32
France	1.8 ± 0.5	19.1 ± 2.5	0.72	6.3×10^{-7}	22
Greece	–	–	–	–	0
Hungary	1.9 ± 1.6	17.4 ± 11.2	0.93	3.5×10^{-2}	4
Iceland	4.9 ± 1.2	0.9 ± 7.0	0.93	2.6×10^{-5}	9
Ireland	5.4 ± 1.1	-4.0 ± 7.4	0.93	1.7×10^{-6}	11
Israel	–	–	–	–	–
Italy	4.8 ± 2.5	-0.3 ± 13.2	0.47	6.1×10^{-4}	21
Japan	1.3 ± 3.2	25.7 ± 27.2	0.02	0.43	43
Netherlands	4.8 ± 3.2	20.5 ± 15.0	0.32	4.7×10^{-3}	23
New Zealand	2.0 ± 0.3	18.8 ± 1.4	0.86	2.6×10^{-12}	27
Norway	-7.2 ± 4.3	50.1 ± 10.6	0.24	1.6×10^{-3}	39
Poland	–	–	–	–	0
Portugal	–	–	–	–	1
Romania	–	–	–	–	0
Spain	6.0 ± 6.2	-7.4 ± 41.7	0.38	5.7×10^{-2}	10
Sweden	5.4 ± 0.6	4.3 ± 2.3	0.92	1.7×10^{-15}	27
Switzerland	2.8 ± 5.6	7.2 ± 38.6	0.69	0.17	4
United Kingdom	5.6 ± 0.7	1.6 ± 4.5	0.88	5.3×10^{-18}	37
United States	3.6 ± 0.3	-0.1 ± 2.3	0.95	1.1×10^{-22}	35

$\pm$ indicates 95% confidence intervals. We report R^2 values for the linear regression of x on c, the p-value of the correlation between x and c, and the number of years for which both x and c measurements are available, n_{obs}.

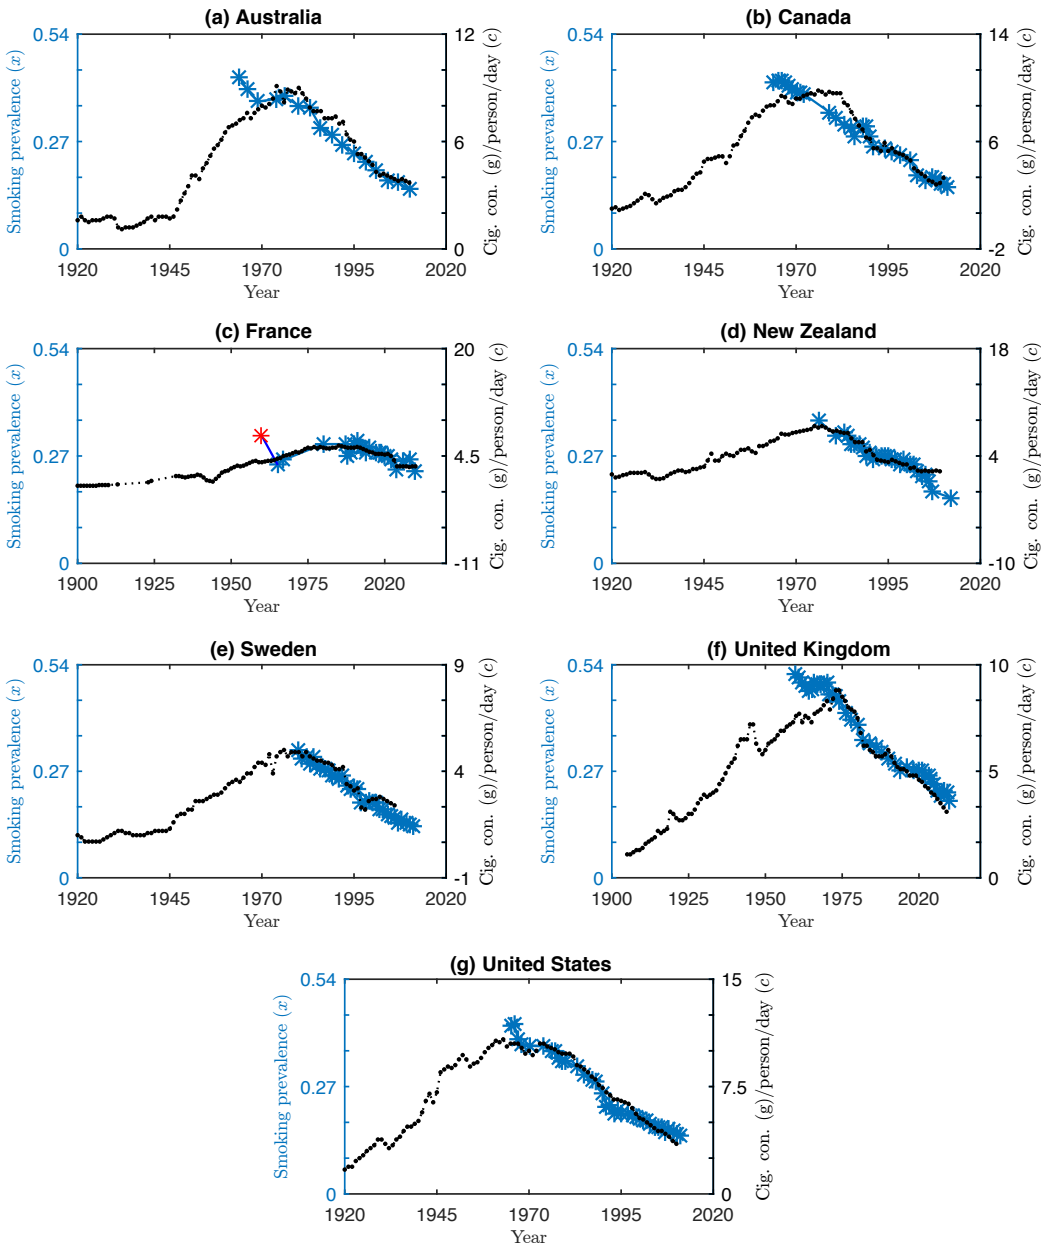

Figure 4.4: Raw smoking prevalence x (blue asterisks, left axis) and raw cigarette consumption c (black dots, right axis) versus time. Raw cigarette consumption data is given in grams per person per day. A single outlier for smoking prevalence (x) for the country of France (panel c) is denoted with a red asterisk.

correspond to an increase in the average mass of cigarettes smoked (in grams) per smoker per day from 11.3 to 16.4 (an increase of 45.1%) over a short 5 year period. This is in sharp contrast with the relatively stable relationship between x and c for France's remaining data points and justifies the exclusion of the outlier $x(1960) = 0.32$. With the outlier removed, France satisfies our data quality requirements for inclusion in the set of seven OECD countries ($R^2 \geq 0.7$, $p < 0.001$, and $n_{obs} \geq 15$).

Our assumption of linearity between smoking prevalence x and cigarette consumption c is not perfect, but it appears to be satisfied at most times in countries where both data sets are available. Quadratic or other higher order terms could be included, but additional unknown parameters would have to be introduced and the limitations of our data set (sparsity, noise) mean that there would be little or no improvement in the model's fit.

Proxy Data $n(t)$: articles published on the health effects of smoking

We calculate the cumulative number of articles published on the health effects of smoking $n(t)$ by performing a search of the online research database Scopus for papers with

(i) tobacco, smok*, or cigar* in the title, and

(ii) death, illness, mortality, risk*, tumour*, tumor*, or cancer in the title, and

(iii) medicine, dentistry, nursing, veterinary, health professions, or multidisciplinary in the subject area, and

(iv) plant*, mosaic, botany, smog, fog, and soot not in the title.

Items (i)-(iii) are search terms included in order to select for papers researching the health effects of smoking, whereas items (iv) are search terms excluded in order to prevent selection of papers researching the tobacco mosaic virus (plant*, mosaic, botany) and the health effects of atmospheric smoke (smog, fog, soot). This provides us with $n(t)$ for integer t, where time t is measured in years. We make the article data available in *Additional file 3* of [108], see Appendix D.2. To calculate $n(t)$ for non-integer and missing values of t we use linear interpolation, see Fig. 4.5(a). Furthermore, Fig. 4.5(b) displays $u_x(t)$ from Eq. (4.2) using $n(t)$ calculated above for various discount factors δ and with $u_0 = 0.51$ and $u_\infty = 0.49$. (For comparison, see Table 4.3 for model-fitted values of δ, u_0 and u_∞.)

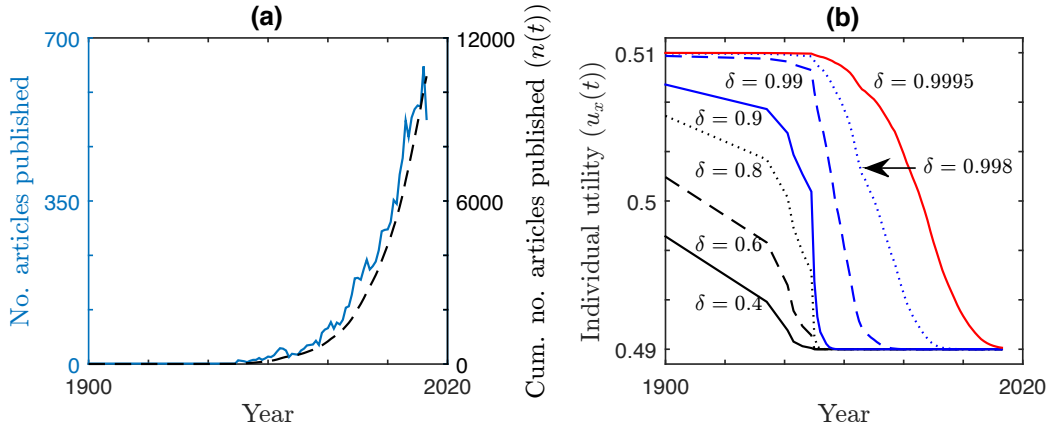

Figure 4.5: Articles retrieved by Scopus with search terms (i)-(iv) and individual utility profiles from Eq. (4.2) for varying values of δ. (a) (Left axis - blue, solid) Number of articles published per year and (right axis - black, dashed) cumulative number of articles published $n(t)$. (b) Discounted utility $u_x(t)$ from Eq. (4.2) with $u_0 = 0.51$ and $u_\infty = 0.49$, using cumulative number of articles published $n(t)$. (Solid black) $\delta = 0.4$, (dashed black) $\delta = 0.6$, (dotted black) $\delta = 0.8$, (solid blue) $\delta = 0.9$, (dashed blue) $\delta = 0.99$, (dotted blue) $\delta = 0.998$, and (solid red) $\delta = 0.9995$.

4.4.2 Model Fitting

We now fit Eq. (4.1) to the estimated prevalence, $\hat{x}(t)$. To reduce the dimensionality of the optimization problem, we assume that certain *universal parameters* are constant across countries. Specifically, we assume that b and δ are universal parameters, and that $x_i(t_{i,0}) = x_{i,0}$, a_i, $u_{i,0}$, and $u_{i,\infty}$ are *local parameters* for country i, where $t_{i,0}$ is the first year for which cigarette consumption data (c), and hence estimated smoking prevalence data ($\hat{x}$), are available. We denote the smoking prevalence estimated above for country i at time t by $\hat{x}_i(t)$. The time series of estimated smoking prevalences for country i is then denoted by the vector $\widehat{X}_i$. Analogously, we denote the time series of smoking prevalences predicted by Eq. (4.1) for country i by $\widetilde{X}_i$. We solve Eq. (4.1) using the Matlab differential equation solver *ode45*.

Using the Matlab function *lsqcurvefit* we now proceed as follows:

1. Holding universal parameters constant, for each country i we find the $x_{i,0}$, a_i, $u_{i,0}$, and $u_{i,\infty}$ that minimize $E_{i,2} = \|\widetilde{X}_i - \widehat{X}_i\|_2^2$.

2. Holding local parameters constant for each country i, we find the b and δ that minimize $E_2 = \sum_i \|\widetilde{X}_i - \widehat{X}_i\|_2^2$.

3. Repeat steps (1) and (2) until either

   (a) the change in the objective function $E_2 = \sum_i \|\widetilde{X}_i - \widehat{X}_i\|^2$ is below tolerance tol, or

   (b) the number of iterations exceeds a limit max_{itn}.

We perform the optimization with the initial guess $u_{i,0} \equiv 0.51$, $u_{i,\infty} \equiv 0.49$, $x_{i,0} = \hat{x}_i(t_{i,0})$, $a_i = 1$, $b = 1$, and $\delta = 0.9985$. We also provide the optimization algorithm *lsqcurvefit* with constraints

$$0 \leq a_i, b \leq 2 \text{ and}$$
$$0 \leq x_{i,0}, u_{i,0}, u_{i,\infty}, \delta \leq 1,$$

and with parameters $tol = 10^{-6}$ and $max_{itn} = 150$. The fitting procedure terminates after 114 iterations, the results of which are recorded in Table 4.3 and Fig. 4.6. For completeness, Table 4.3 also records the average of the absolute value of the difference between $\tilde{X}_i$ and $\hat{X}_i$

$$E_{i,1} = \frac{\|\tilde{X}_i - \hat{X}_i\|_1}{\text{length of } \hat{X}_i}.$$

Complete model simulation code is provided in *Additional file 4* of [108], see Appendix D.3.

4.5 Results: Testing the Model

We now test the model in three phases, as depicted in Fig. 4.2.

4.5.1 Phase (i): Direct test

Figure 4.1 shows the fit of our model to data sets from the United States and Sweden. Additional fits and parameter values are displayed for our set of seven OECD countries in Fig. 4.6 and Table 4.3. The average of the absolute value of the difference between smoking prevalence estimates $\hat{x}$ and the output of Eq. (4.1) ranges from a low of 0.005 for France to a high of 0.018 for the United Kingdom (see $E_{i,1}$ in Table 4.3). The good agreement that we found with all data sets provides support for the model.

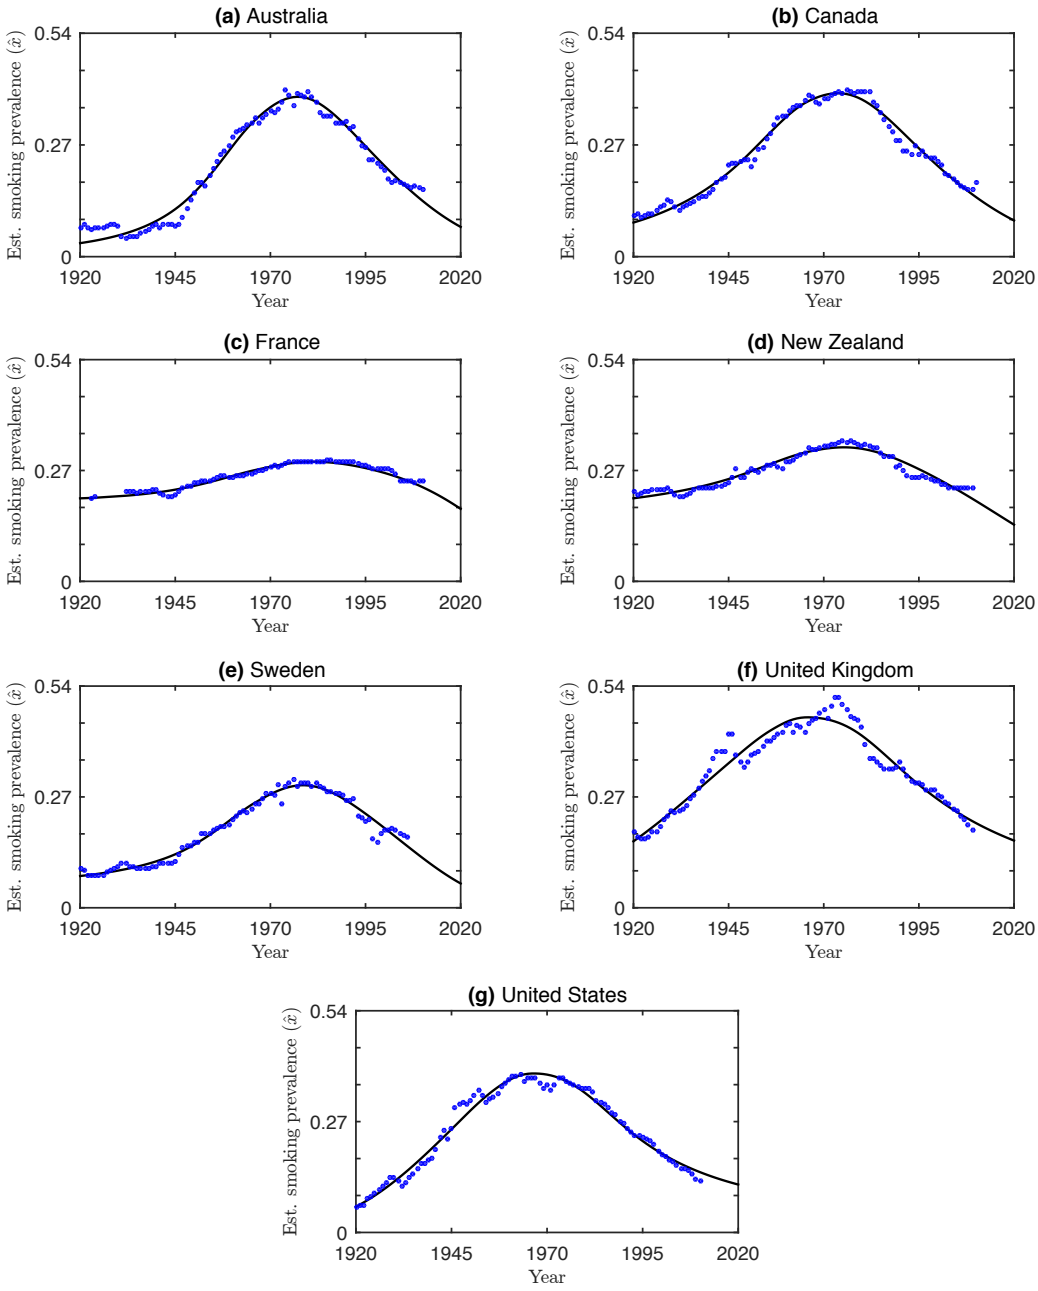

Figure 4.6: The result of fitting Eq. (4.1) to the estimated smoking prevalence $\hat{x}$. (Solid black line) Solution to Eq. (4.1) with parameters from Table 4.3. (Blue dots) Estimated smoking prevalence values $\hat{x}$.

Table 4.3: The result of fitting smoking prevalence model (Eq. (4.1)) to the estimated smoking prevalence $\hat{x}$.

Universal parameters and Total Error (E_2)		
b	δ	E_2
1.049	0.9981	0.163

Country	Local parameters and local error ($E_{i,2}$ and $E_{i,1}$)					
(i)	a_i	$x_{i,0}$	$u_{i,0}$	$u_{i,\infty}$	$E_{i,2}$	$E_{i,1}$
Australia	1.035	0.033	0.551	0.484	0.032	0.015
Canada	1.020	0.083	0.530	0.483	0.020	0.011
France	1.121	0.198	0.543	0.524	0.004	0.005
New Zealand	1.062	0.202	0.525	0.504	0.012	0.010
Sweden	1.076	0.077	0.555	0.503	0.015	0.009
United Kingdom	0.976	0.079	0.513	0.478	0.060	0.018
United States	0.963	0.063	0.513	0.470	0.024	0.013

4.5.2 Phase (ii): Test of model implications for relative conformity parameter a

If the model and its interpretation are correct and the balance between individual and social utility is a relevant factor for the temporal dynamics of smoking prevalence, then we expect that the fitted "relative conformity parameter" a will be different for different countries and will capture something meaningful about the individualism/collectivism of a society. To test this we compare with Hofstede's IDV, an established metric for societal individualism [86] that has been evaluated in most countries. Panel (a) of Fig. 4.7 shows the comparison. The relative conformity parameter a shows significant differences for different countries, and, as expected, a is negatively correlated with Hofstede's IDV (negative because a increases with collectivism while IDV decreases with it). This concordance with independently assessed individualism values supports our model.

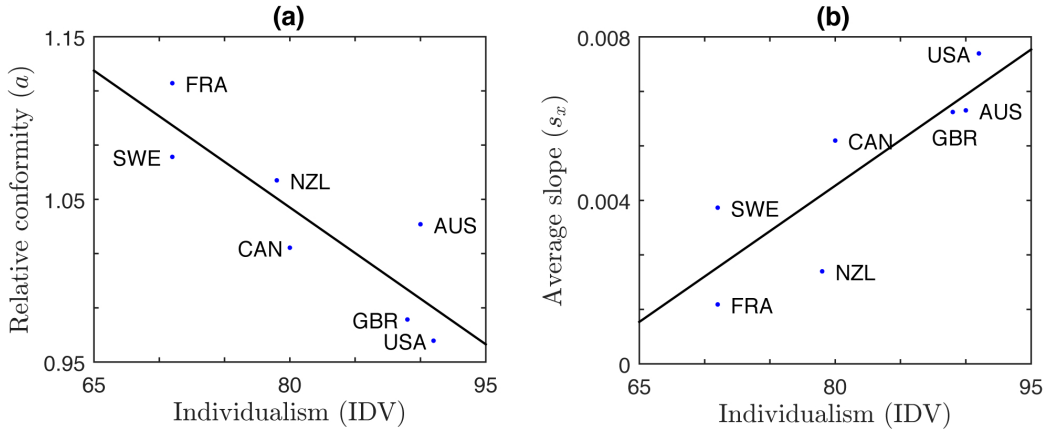

Figure 4.7: Relative conformity a and average slope s_x versus Hofstede's IDV for seven OECD countries. For both panels the line of best fit is given by a solid line. (a) Relative conformity a is negatively correlated with IDV ($\rho = -0.87$, $p = 0.011$, best fit slope: -5.6×10^{-3}). (b) The positive correlation between average slope in smoking prevalence s_x and IDV ($\rho = 0.85$, $p = 0.015$, best fit slope: 2.2×10^{-4}) is consistent with slower change in collectivistic societies: the average slope is greater in individualistic societies and smaller in collectivistic societies.

4.5.3 Phase (iii): Test of model implications for slope and peak year

Besides the correlation of a with collectivism, we note that another prediction is implicit in model (4.1). As the relative conformity parameter increases, the model requires that changes in smoking prevalence occur more slowly (this is true for solutions to Eq. (4.1) for the range of a and u values corresponding to the observational data). Put another way, societies with higher levels of individualism should experience faster changes in smoking prevalence. Intuitively, when smoking prevalence is low the lack of existing smokers inhibits smoking initiation more strongly in a collectivistic society than in an individualistic society. Thus, we expect the average rate of increase in a collectivistic society to be smaller than in an individualistic society. In contrast, when smoking prevalence is high, and once the deleterious health effects of smoking become widely known and negatively impact individual utility from smoking, the presence of existing smokers inhibits smoking cessation more strongly in a collectivistic society than in an individualistic society. In both cases collectivism acts as a break on change in the status quo (higher cultural inertia [190, 83]). Panel (b) of Fig. 4.7 and panel (a) of Fig. 4.8 demonstrate that this is indeed the case: the

average slope s_x of the smoking prevalence curves leading up to the peak increases with Hofstede's IDV and decreases with a, respectively. Here we define the average slope s_x to be

$$s_x = \frac{\hat{x}(t_{max}) - \hat{x}(t_0)}{t_{max} - t_0}, \tag{4.6}$$

where $t_0 = 1920$ is the first year for which smoking prevalence estimates are available in the subset of seven OECD countries, and where t_{max} is the earliest year for which the maximum tobacco consumption was recorded, see Table 4.4. Correlations are significant (see Table 4.5).

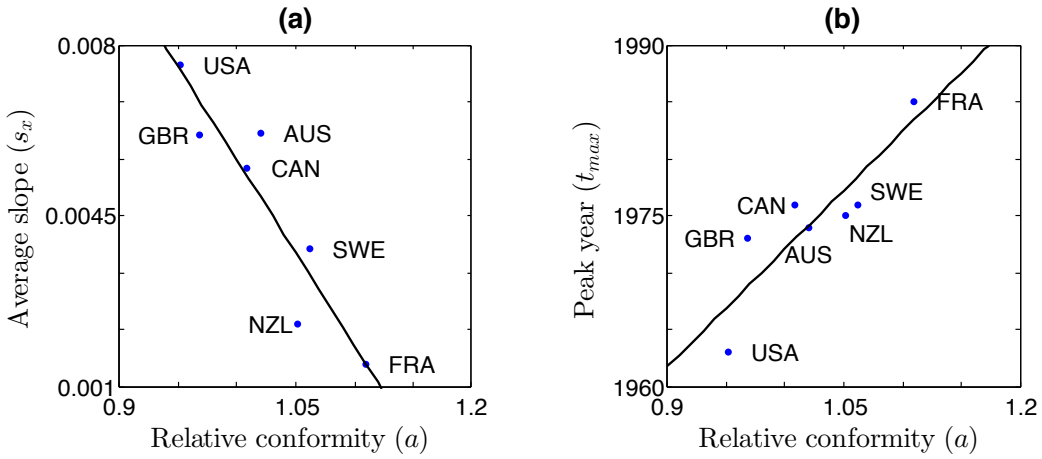

Figure 4.8: Average slope s_x and peak year of smoking prevalence t_{max} versus relative conformity parameter a. (a) Average slope s_x versus relative conformity parameter a ($\rho = -0.92$, $p = 0.003$). (b) Peak year t_{max} versus relative conformity parameter a ($\rho = 0.88$, $p = 0.009$). The line of best fit is given by a solid line.

Table 4.4: Hofstede's Individualism Index IDV and peak year t_{max}.

Country	Abbreviation	IDV	Peak year (t_{max})
Australia	AUS	90	1974
Austria	AUT	55	1979
Belgium	BEL	75	1973
Canada	CAN	80	1976
Denmark	DNK	74	1976
Finland	FIN	63	1963
France	FRA	71	1985
Greece	GRE	35	1986
Hungary	HUN	80	1980
Iceland	ICE	60	1984
Ireland	IRE	70	1974
Israel	ISR	54	1974
Italy	ITA	76	1984
Japan	JPN	46	1977
Netherlands	NLD	80	1977
New Zealand	NZL	79	1975
Norway	NOR	69	2004
Poland	POL	60	1991
Portugal	PRT	27	1994
Romania	ROM	30	1995
Spain	ESP	51	1985
Sweden	SWE	71	1976
Switzerland	CHE	68	1972
United Kingdom	GBR	89	1973
United States	USA	91	1963

Table 4.5: Correlation between IDV, relative conformity a, average slope s_x, and peak year t_{max}.

	7-country subset			25-country set
	a	s_x	t_{max}	t_{max}
IDV	-0.87 (0.011)	0.85 (0.015)	-0.76 (0.047)	-0.53 (0.006)
a	–	-0.92 (0.003)	0.88 (0.009)	–

Correlations between IDV, a, s_x, and t_{max} are recorded for the seven-country subset. Correlation between IDV and t_{max} is recorded for the full set of 25 countries. p-values are in parentheses. All correlations are significant at the 95% confidence level.

We note that fluctuations in the data due to either volatility in tobacco consumption or measurement error may affect reported t_{max}. Smoothing of the data could be applied prior to calculation of peak year, however, the choice of smoothing algorithm is itself arbitrary and unnecessarily complicates our findings without significantly altering the result. For example, consider the seven OECD countries for which we have estimated historical smoking prevalence data $\hat{X}_i$. We observe that the model fitting procedure described in the Methods section results in the time series $\tilde{X}_i$, which we can consider as one possible smoothing of the data $\hat{X}_i$. In this case, the measurement for peak year does not change substantially after smoothing for most countries (see Fig. 4.6(a)-(e)), while the measurement for peak year in the USA would slightly increase from $t_{max} = 1963$ to $t_{max} = 1967$ and the measurement for peak year in the UK would slightly decrease from $t_{max} = 1973$ to $t_{max} = 1966$ (see Fig. 4.6(f)-(g)). These changes would result in no discernible net change in the relationship between peak year and individualism, but would result in added complexity, and hence, in a greater chance of introducing additional error[1].

This reasoning further suggests that the peak year for smoking prevalence t_{max} should be later in collectivistic societies and earlier in individualistic countries. As shown in Fig. 4.9, the raw observational data are consistent with this prediction: t_{max} is significantly negatively correlated with IDV (shown) and significantly positively correlated with a (see Fig. 4.8(b)). Note that our assumption of a linear relationship between national cigarette consumption and smoking prevalence is not needed to establish t_{max}, so Fig. 4.9 is independent of any model assumptions.

[1]Indeed, these changes would increase the statistical significance of our results, but again, we don't believe that they justify the additional complexity and the introduction of additional arbitrary smoothing parameters.

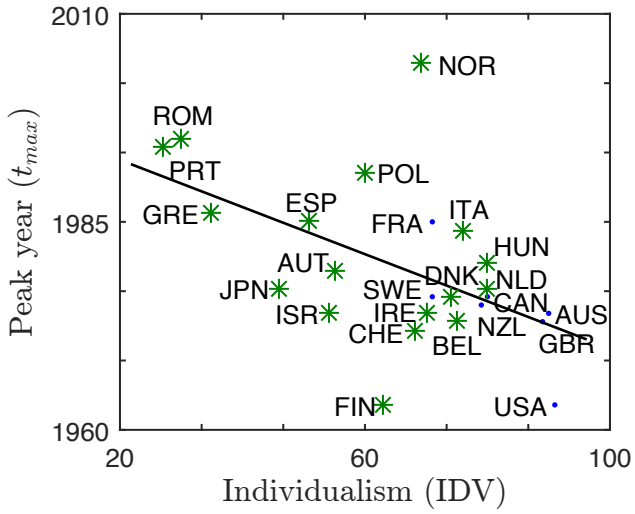

Figure 4.9: Peak year t_{max} in cigarette consumption versus Hofstede's individualism index IDV for 24 OECD countries and Romania. The negative correlation between peak year t_{max} and IDV in a set of 25 countries ($\rho = -0.524$, $p = 0.008$) is consistent with slower change in collectivistic societies: the peak year in tobacco consumption tends to occur later in collectivistic societies and earlier in individualistic societies. The seven OECD countries considered for the mathematical model are indicated by dots ($\rho = -0.76$, $p = 0.047$), and the remaining 18 countries are indicated by asterisks. The line of best fit, calculated using all 25 countries, is given by a solid line. See Table 4.4 for country abbreviations.

4.6 Application: A counterfactual scenario

In this section we illustrate the effect size of individualism/collectivism on the dynamics of the smoking epidemic by considering a simple counterfactual scenario. Specifically, holding all other fitted parameters constant, we consider how the smoking epidemic in the United States might have evolved if the United States (IDV=91 and $a = 0.963$) were about 2% less individualistic (IDV=89 and, using the slope from Fig. 4.7(a), $a = 0.974$). Fig. 4.10 plots an estimate for the number of cigarettes smoked per year (in trillions) versus time. Integrating the difference between the number of cigarettes smoked per year versus time for the United States with fitted ($a = 0.963$, solid line) and counterfactual ($a = 0.974$, dashed line) relative conformity implies that, according to our model, if the United States had 2% lower individualism during the 90 year period from 1920–2010 then there would have been approximately 7×10^{12} fewer cigarettes smoked. This is equivalent to a 16% decrease in the number of cigarettes smoked.

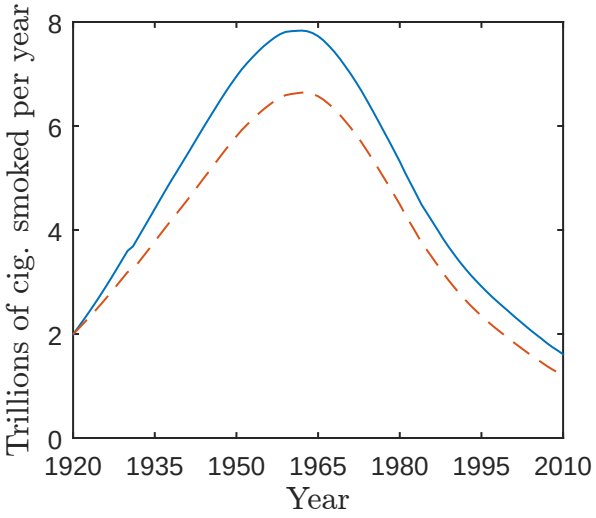

Figure 4.10: Solution to Eq. (4.1) for the United States with $a = 0.963$ (solid) and $a = 0.974$ (dashed). Parameters x_0, b, u_0, u_∞, and δ are as reported for the United States in Table 4.3.

The number of cigarettes smoked per year is estimated as follows. First, we observe that for year t the number of cigarettes smoked per smoker per day is[2] $\mathcal{C}_d(t) = c(t)/x(t)$. Therefore, using Eq. (4.5) we find that the number of cigarettes smoked per smoker per year $\mathcal{C}_a(t) = 365 \times \mathcal{C}_d(t)$ can be bounded. For example, in the case of the United States, where $\hat{B} < 0$, we find that $\mathcal{C}_a(t)$ is bounded by

$$\begin{aligned} 1.02 \times 10^4 = 365 \times \hat{C}^{-1} \times \frac{365 \text{ days}}{\text{year}} \\ \leq \mathcal{C}_a(t) = \frac{1 - \hat{B}/\hat{x}(t)}{\hat{C}} \times \frac{365 \text{ days}}{\text{year}} \\ \leq \frac{1 - \hat{B}/\min \hat{x}(t)}{\hat{C}} \times \frac{365 \text{ days}}{\text{year}} = 1.04 \times 10^4. \end{aligned}$$

Since the lower and upper bounds are relatively tight, we estimate the number of cigarettes smoked per smoker per year to be the average of the lower and upper bounds

$$\bar{\mathcal{C}}_a \approx \frac{2 - \hat{B}/\min \hat{x}(t)}{2\hat{C}} \times \frac{365 \text{ days}}{\text{year}} \approx 1.0 \times 10^4.$$

We cross-check this estimate with the direct estimate of $\bar{\mathcal{C}}_a$ taken by averaging $c(t)/x(t)$ for all times where both measurements are available in the raw data (data shown in Fig. 4.6).

[2] Assuming 1.002 cigarettes per gram, as in [63, 64].

These two estimates agree to two significant figures. Finally, we estimate the number of cigarettes smoked per year to be

$$\hat{x}(t) \times N_{pop}(t) \times \bar{\mathcal{C}}_a,$$

where $N_{pop}(t)$ is the total population at time t. The total population for the United States is taken from US census estimates [146, 147] and is given in *Additional file 5* of [108], see Appendix D.4.

We emphasize that in the counterfactual scenario described above we have only changed a for the United States while keeping all other fitted parameters constant, merely to illustrate that the effect of small changes in a in the model can be large. Therefore, the broad variation in the fitted a across countries, as illustrated in Fig. 4.7(a), can indeed be expected to lead to a large effect size on cigarette consumption. Note, however, that the results from this counterfactual scenario do not imply that less individualism automatically means lower cigarette consumption, since countries with lower IDV (higher a) than the United States also tend to differ for other fitted parameters and quantities in the model, resulting in substantially different solutions to Eq. (4.1).

4.7 Discussion

Before discussing the limitations of our model, it is worth discussing the potential effect of confounding variables on our model. Specifically, we argue that the effect of confounding variables on our results are limited, since most potential candidates for confounding variables are actually accounted for implicitly in our modelling framework. Consider, for example, two possible candidates for confounding variables: the wealth (per capita GDP) and the strength of tobacco control policies in each country. Specifically, consider the trend in Fig. 4.9 where wealthier countries have, on average, earlier peak year in tobacco consumption (t_{max}). This relationship is easily explained by our model, since (a) our model predicts a negative correlation between peak year and Hofstede's IDV and (b) IDV and wealth (per capita GDP) are highly positively correlated [86]. An alternative explanation for wealthier nations having earlier peak year, however, would be that individuals who are more wealthy are better able to afford cigarettes and, in aggregate, are better able to implement strong anti-tobacco policies: in theory, the former would lead to a more rapid increase in smoking prevalence and the latter would lead to a more rapid decline in smoking prevalence. Although this alternative explanation might seem to be in competition with our model, we argue that it is in fact accounted for implicitly in our modelling

framework: both wealth and the strength of tobacco control policies are contributing factors to individual utility from smoking. Furthermore, we note that although the precise timing of anti-tobacco policies is not included in the model, it is reasonable to assume that these initiatives are implemented more frequently and more intensely as the health effects of smoking are better understood - a phenomenon which is modelled using Eq. (4.2) and proxy data on smoking related publications. In summary, since most potential confounding variables are actually accounted for in, and not in competition with, our modelling framework, the exposure of our results to the effects of confounding variables is limited.

Despite the good match between model predictions and data, a number of limitations remain. First, due to limitations in the quantity and quality of the available smoking prevalence and tobacco consumption data, we are only able to fit the parameters in our model to seven countries, all of which have advanced/developed economies. There is no *a priori* reason to believe that, given adequate sources of data, our model would not generalize to less developed countries with lower income. Indeed, Fig. 4.9 supports the position that the behaviour in less developed countries is consistent with our mathematical model. Nevertheless, our inability to directly apply our model to a larger set of more diverse countries due to a lack of good data remains a limitation of our work and an area open to future research. Second, we have made an implicit "mean-field" approximation in taking social utility to be a function of the overall smoking prevalence x, rather than the local smoking prevalence among contacts in an individual's social network. Similarly, we have taken individual utility to be uniform across the population (though not in time), whereas a more detailed model might allow for variation with, e.g., gender, age and income. We note, however, that the success of our model in reproducing the trends in smoking prevalence in a manner consistent with its interpretation in the context of individualism/collectivism, despite these limitations, is generally supportive of the modelling framework we have developed. In particular, our results and the data indicate that, when averaging over gender, age and income in a country, a strong net influence remains from societal individualism on the aggregate temporal dynamics of smoking prevalence. Furthermore, if the mechanism in our model did not reflect the reality of the decision-making process for smoking then, even if it still somehow managed to fit the smoking prevalence data, we would not expect to simultaneously find correlation of (a) the relative conformity parameter a with Hofstede's individualism measure IDV (Fig. 4.7(a)), (b) average slope s_x with IDV (Fig. 4.7(b)), and (c) peak year t_{max} with IDV (Fig. 4.9). Moreover, we would not expect that the sign of these correlations would be consistent with the predictions of the model.

Our findings suggest that the correlation of individualism with faster societal change (as a consequence of a sudden change in personal utility) results from a causative influence as predicted by our model. As already mentioned, other factors such as income levels

also correlate with individualism. We certainly cannot exclude that there may be other causative factors. For example, our model in its current form is incapable of explaining differences in smoking prevalence between genders and why these inter-gender differences vary between countries [133, 134]. Nevertheless, we remark that many previously proposed causative factors for differences in observed inter-country smoking dynamics can be accounted for within our modelling framework. In particular, beliefs about the harmful effects of smoking [45], the price of cigarettes [67], socioeconomic status and inequality [13, 91], and government regulation [171, 127, 96] have all been cited as potential factors affecting the differences observed in inter-country smoking dynamics. Each of these factors can be interpreted within our modelling framework. For example, beliefs about the harmful effects of smoking, as well as the price of cigarettes, both likely contribute directly to individual utility derived from smoking (u_x) and from non-smoking (u_y). Moreover, socioeconomic status may affect individual utility from smoking indirectly by affecting an individual's tolerance for risk and/or how they discount future rewards and costs (i.e. how they discount their future health status) [76]. Addressing the model's inability to account for gender differences in smoking prevalence and explicitly quantifying the relationship between other causative factors and model parameters in more elaborate models are potential areas for future work.

4.8 Conclusion

Despite the above mentioned limitations, the quantitative mathematical model proposed in this chapter, which we derived from basic principles well-documented in the sociology and social psychology literature, appears to match real-world smoking prevalence data from a variety of countries well (to our knowledge, the largest historical data set of this type ever compiled), and all predictions of the model appear to be supported by the data. Indeed, we emphasize the strong support of the model by the data, since the model was calibrated (in Phase (i)) and its predictions were tested (in Phases (ii)-(iii)) using two separate data sets (tobacco use data and Hofstede's IDV, respectively). In particular, the model predicts that the level of individualism or collectivism of a society may significantly affect the temporal dynamics of smoking prevalence: the strong influence of the individual utility of smoking (and its decrease due to increased awareness of adverse health effects) is predicted to lead to faster adoption and cessation of smoking in individualistic societies than in more collectivistic societies. The significance of this effect can be illustrated by considering the counterfactual scenario of how the smoking prevalence might have evolved in the United States had the United States been less individualistic. Specifically, we estimate that a

reduction in the IDV of the United States of 2% would have resulted in a 16% decrease in the total number of cigarettes smoked between 1920 and 2010, see Section 4.6.

It has previously been argued that social support mechanisms in collectivistic societies make it more likely that a person will stop smoking [166, 174] based on findings that social support (supportive counselors) can help people to adhere to decisions to quit smoking [94]. In contrast to this behaviour at the individual level, we find that aggregate smoking prevalence decreases more slowly in collectivistic societies. Since the aggregate smoking prevalence is a function of both smoking adoption and cessation, our model suggests that this may be so because social inertia/peer pressure will either inhibit the decision to stop smoking, or encourage the decision to start smoking, more strongly in collectivistic societies than in individualistic societies.

These results suggest that it may be effective to consider culture-dependent strategies for combating the ongoing smoking epidemic. For example, interventions to discourage smoking can be tailored differently in societies or social groups whose cultures differ in how they value individualism versus collectivism [142]. Specifically, consider how the goal of many tobacco control policies is to reduce the individual utility from smoking, often by increasing the cost of cigarettes through sin taxes or by requiring warnings on cigarette packages illustrating the danger of smoking to health. Our results suggest that these tactics will be more successful in individualistic societies and less successful in collectivistic societies. In contrast, tactics that may be more successful in collectivistic societies might focus on social dangers resulting from smoking, for example by emphasizing the association between smoking and low social status [13, 91], or emphasizing the large number of individuals who have already quit. More broadly, these results demonstrate that differences in culture can measurably affect the dynamics of a social spreading process, and that a mathematical model can help to illuminate this phenomenon. We welcome future work comparing a variety of social contagion phenomena across societies. Our model suggests that the increased cultural inertia in collectivistic societies may potentially lead to slower change across a wide spectrum of spreading processes (those where important changes occur in personal utility), a hypothesis that could be supported or rejected by further study.

Chapter 5

Modelling the Dynamics of the Body Mass Index (BMI) at the Individual and Population Levels: Evidence for the social spread of obesity

5.1 Introduction

Over the past 35 years there has been a near doubling in the worldwide prevalence of obesity [60, 131]. This drastic increase is of major concern: obesity is a risk factor for many chronic illnesses [56, 112] and therefore is responsible for significant economic costs both directly (e.g., through increased healthcare expenses) [30, 58] and indirectly (e.g., through loss of productivity) [74, 175].

The body mass index (BMI), defined as the mass (in kilograms) divided by the height squared (in meters^2), is a standard measure of relative body weight used to classify individuals as underweight (BMI $\leq$ 18.5), normal weight (18.5 $<$ BMI $\leq$ 25), overweight (25 $<$ BMI $\leq$ 30), or obese (BMI $>$ 30)[1] [49]. Because the risks from being overweight or obese increase as BMI increases [8, 17, 143], a significant amount of research has been done on the distribution of BMIs in the population [62, 123, 140, 141, 155, 160, 164]. These studies have found that the distribution of BMIs in various populations are right-skewed

[1]Obesity is further subdivided into moderately obese (obesity class I: 30 $<$ BMI $\leq$ 35), severely obese (obesity class II: 35 $<$ BMI $\leq$ 40), and very severely obese (obesity class III: BMI $>$ 40).

(see Fig. 5.1 for an example), and that the observed increase in obesity prevalence over time can be explained by an increase in the right-skewness of the BMI distribution and by an overall rightward shift of the BMI distribution, with many authors arguing for increased right-skewness as the primary cause of the increase in the prevalence of obesity [62, 140, 141, 160, 164]. These findings have significant implications for public health intervention policies: if the primary cause of increased obesity prevalence is increased right-skewness then intervention strategies targeting the entire population may not be as effective as strategies that target the most overweight individuals (i.e., the right tail of the BMI distribution), and vice versa [3, 141, 154, 193].

Although there has been speculation as to why the BMI distribution is right skewed [168], a mechanistic explanation for this phenomenon that links the behaviour of individuals to the right-skewness of the distribution has yet to be proposed. Specifically, while it has been observed that the distribution of BMIs in a population is approximately log-normal [141], explanations for the shape of this distribution are mostly conceptual. For example, consider the "runaway train" argument [168] which posits that there are multiple forces inducing weight gain that positively reinforce each other (e.g. weight gain may result in/from difficulty participating in physical activity, increased caloric intake due to low self-esteem or body dissatisfaction, and increased exposure to obesogenic environments, etc...) and only very weak forces opposing weight gain (e.g. social discrimination, physical discomfort, reduction in appetite, and insulin resistance, etc...). Thus, individuals are susceptible to a positive feedback loop that is responsible for large amounts of weight gain. The "runaway train" argument is an intuitively appealing explanation for the right-skewness of the BMI distribution, under the additional assumption that not all individuals are equally susceptible to the "runaway train" phenomenon[2].

Despite the increasing interest in studying how the distribution of BMIs changes over time, we remark that there is no established standard approach for comparing two different BMI distributions [62]. Most common approaches are graphical [61, 62], for example, plotting the BMI distribution functions, plotting the cumulative BMI distribution functions, or generating a Tukey mean-difference plot[3]. While these techniques are useful, they unfortunately suffer from significant limitations. In particular, these methods are only capable of qualitatively describing changes in BMI distributions, they may not detect subtle

[2]It is conceivable that if all individuals were equally susceptible to the "runaway train" phenomenon then the result would be a rightward shift of the overall distribution with little, if any, increase in right-skewness.

[3]In the Tukey mean-difference plot the difference in values of a particular percentile is plotted against the mean in values of that same percentile. For example if the 50^{th} percentile of BMI distribution A is 25 and the 50^{th} percentile of the BMI distribution B is 28 then this results in a data point with coordinates (26.5, -3), since $0.5(25 + 28) = 26.5$ and since $25 - 28 = -3$.

changes in BMI distributions, and may be difficult to apply when comparing many BMI distributions simultaneously.

In addition to investigations on the shape of the BMI distribution, a significant amount of research has also been done to identify major contributing factors to the obesity epidemic, for example see the "Foresight Obesity Map" [24]. These can roughly be divided into factors intrinsic to the individual (e.g., energy consumption, activity level, age, gender, ethnicity, education, socioeconomic status, environment, etc...) [93, 125, 160, 164, 169, 170, 189] and extrinsic factors (e.g., social/peer influences [39, 145, 176, 182], or microbial influences [178], etc...). Although the effects of intrinsic factors have been well established, there is some uncertainty over the importance of external factors, especially social and peer influences [41].

In this chapter we present a mechanistic mathematical model for BMI dynamics with implications at both the population (macro) and individual (micro) level that is consistent with three independent sources of BMI data. Our model makes two main contributions at the macro-level. First, our model predicts a functional form for the BMI distribution that fits empirical data better than two standard distribution functions commonly used to describe right-skewed data (i.e. the log-normal and skew-normal distributions). Second, our model proposes a potential mechanistic explanation for the right-skewness of the BMI distribution. In addition, we remark that the parameters of our model relate the behaviour of individuals to the mean, standard deviation, and skewness of the BMI distribution in a straightforward way, and therefore, allows for the comparison between BMI distributions from different years or different regions in a consistent manner. For example, in contrast to the graphic techniques described above, our model allows us to disentangle the effect of (a) overall rightward shift of the BMI distribution and (b) increased right skewness of the BMI distribution on the observed increase in the prevalence of obesity over the past three decades. Our model also makes a significant contribution at the micro-level. Specifically, our model predicts a functional form for both the average and standard deviation of year-over-year change of individuals' BMI, which we verify by fitting to empirical data. At both the individual and population levels our model allows us to investigate the role of social and peer influences in BMI dynamics, and provides additional evidence in support of the hypothesis that social effects make a contribution to BMI dynamics.

We note that our work differs from previous statistical studies [39, 145, 176, 182] that investigate the role of social and peer influences in that we propose a mechanism through which social and peer influences can affect dynamics of the BMI, and differs from previous compartmental [48, 185] and network [11, 73] mathematical models in that our model closely replicates BMI data from three independent data sets at both the macro- and micro-level. Interestingly, while the "runaway train" phenomenon may indeed be an im-

portant driver of individuals' weight gain, our model illustrates that this phenomenon is not necessary to explain the right skewness of the BMI distribution.

Finally, as part of the above studies, we compile and present analysis of an entirely new BMI data set more abundant than any previously reported. BMI measurements calculated from anonymized medical records for patients of the Northwestern Medicine system of hospitals and clinics are available from 1997 through 2014. Although this data does not form a representative sample of the population, it is nevertheless extremely valuable since it represents the largest known collection of BMI panel data. Specifically, the Northwestern Medicine medical record data set contains measurements for 329,543 distinct individuals whose year-over-year change in BMI can be calculated (1,017,518 data points in total).

The remainder of this chapter is organized as follows. We specify our mathematical model for the dynamics of individuals' BMI in Section 5.2. Section 5.3 then introduces three BMI data sets that will be used to calibrate and test our model: the National Health and Nutrition Examination Survey (NHANES), the Northwestern Medicine medical records data set (NU), and the Behavioral Risk Factor Surveillance System (BRFSS). The population and individual level implications of our model are outlined in Section 5.4, and we test these implications in Section 5.5 by comparing to the NHANES, Northwestern Medicine, and BRFSS data sets. Section 5.6 then concludes with a brief discussion.

5.2 Model Specification

The specification of our model for the dynamics of individuals' BMI proceeds in two stages. In Section 5.2.1 we specify the deterministic component of our model, which models the deliberate actions taken by individuals in order to modify their BMI. In Section 5.2.2 we specify the stochastic component of our model, which models random fluctuations in individuals' lifestyle (e.g. calorie consumption and activity level) that affect the dynamics of their BMI.

5.2.1 Deterministic Component of BMI dynamics

Intrinsic Dynamics

A model for a contagion process can be divided into two key components: intrinsic dynamics (self) and extrinsic dynamics (interaction with others) [1, 2, 108]. The intrinsic

dynamics of weight change in healthy adults are thought to follow a "return to equilibrium" pattern where, absent social effects, individuals tend to fluctuate about a natural equilibrium, or "setpoint" [25, 65, 80, 81, 97, 163]. This setpoint weight may depend upon many factors including genetics, average exercise and eating habits, etc... Though the setpoint may vary gradually over the course of an individual's life, we approximate it as a constant on the shorter time scale over which our model applies.

We start with a simple deterministic model for the intrinsic dynamics by assuming exponential decay to equilibrium as

$$\frac{dx}{dt} = k_1(x^\star - x) \ , \tag{5.1}$$

where $x = x(t)$ represents the BMI of an individual at time t, $x^\star$ (Setpoint) represents the individual's BMI setpoint, and the constant $k_1 > 0$ determines the rate of exponential relaxation to equilibrium weight (note that we assume constant height in adults over time, so changes in BMI—defined as the ratio of weight to height squared—are proportional to weight changes). We will later add stochastic noise terms to account for natural fluctuations in lifestyle, see Section 5.2.2.

Another way to deduce this same model for intrinsic dynamics is to assume that individuals tend to maximize some *individual utility function* $u_I(x) = u_I(x; x^\star)$. This function could in general have a complicated shape, but by assumption must be peaked at $x = x^\star$, i.e., by assumption $u_I(x)$ must have a local maximum at $x = x^\star$. Assuming that this utility function is smooth, the Taylor series expansion around $x = x^\star$ is

$$u_I(x) = u_I(x^\star) + \underbrace{u_I'(x^\star)}_{=0}(x - x^\star) + \frac{1}{2}\underbrace{u_I''(x^\star)}_{<0}(x - x^\star)^2 + \dots \ , \tag{5.2}$$

where $u_I'(x^\star) = 0$ and $u_I''(x^\star) < 0$ because we have assumed $u_I(x)$ is peaked at $x = x^\star$. Note that in what follows we set the additive constant $u_I(x^\star)$ to zero for convenience, without loss of generality. We assume that an individual changes their BMI in such a way as to increase his or her utility. A further mild assumption is that the rate of change of BMI will be proportional to the rate of increase of utility, or

$$\frac{dx}{dt} = k_2\frac{du_I}{dx} \ , \tag{5.3}$$

leading to the same intrinsic dynamics as model (5.1) when the Taylor series is truncated at quadratic order (the lowest order that gives nontrivial dynamics)[4].

[4]In this case, Equations (5.1) and (5.3) are identical when $k_1 = -k_2 u_I''(x^\star)$.

Extrinsic Factors

The extrinsic dynamics of weight change are more difficult to model. Some theories suggests that individuals can become accustomed to the average BMI of peers under exposure to different peer environments [167] and, to reduce disparity, may adjust their weights. We assume that there exists some *social utility function* $u_S(x) = u_S(x; x_{\text{peer}})$ which captures this proposed peer-influence phenomenon: the social utility should peak when an individual reaches the same BMI as his or her peer(s), x_{peer}. Similarly to the intrinsic dynamics, we expect this utility to be well approximated by a quadratic function (at least locally) and therefore propose that for a single neighbour the social utility function takes the form

$$v_S(x; x_{\text{peer}}) = \frac{1}{2}\, v_S''(x_{\text{peer}}; x_{\text{peer}})\, (x_{\text{peer}} - x)^2 \ ,$$

where x_{peer} is the BMI of some peer who influences the individual under consideration, and where we assume that $v_S''(x_{\text{peer}}; x_{\text{peer}}) < 0$ is a constant. When multiple peers simultaneously influence an individual, the net social utility for individual i becomes

$$u_S(x_i) = u_S(x_i; \vec{x}) = \sum_{j=1}^{N} A_{ij} v_S(x_i; x_j) = \frac{1}{2} \underbrace{v_S''(x_j; x_j)}_{= v_S''(x_{\text{peer}}; x_{\text{peer}})} \sum_{j=1}^{N} A_{ij} (x_j - x_i)^2 \ ,$$

where x_i is the BMI of individual i, N is the number of individuals in the population, $\vec{x} = (x_1, x_2, \ldots, x_N)^T$, and A_{ij} represents the strength of social influence of individual j on individual i.

Combining both the intrinsic and extrinsic aspects of the proposed social contagion process, we propose the following net utility for individual i:

$$\begin{aligned} u(x_i) &= c_I\ u_I(x_i) + c_S\ u_S(x_i) \\ &= \frac{1}{2}\, c_I\ u_I''(x^\star; x^\star)\ (x^\star - x_i)^2 + \frac{1}{2} c_S\ v_S''(x_{\text{peer}}; x_{\text{peer}})\ \sum_{j=1}^{N} A_{ij} (x_j - x_i)^2 \ , \end{aligned} \tag{5.4}$$

where the constants c_I and c_S set the relative importance of individual versus social factors. Thus, using the same logic as in Eq. (5.3) above, we find the model for the deterministic dynamical system is given by $dx_i/dt \propto du/dx_i$, i.e.,

$$\frac{dx_i}{dt} = \underbrace{k_I(x^\star - x_i) + k_S \sum_{j=1}^{N} \left[A_{ij}(x_j - x_i) - \frac{1}{2} \frac{\partial A_{ij}}{\partial x_i} (x_j - x_i)^2 \right]}_{= a_N(x_i; \vec{x})} \ , \tag{5.5}$$

where we denote $k_I > 0$ and $k_S > 0$. In order to specify A_{ij} we assume that individuals interact most strongly with others who are most similar to themselves, i.e. individuals with similar BMI interact more strongly than individuals with different BMI [33, 44, 121]. Consistent with this assumption, we choose

$$A_{ij} = A(x_i, x_j) = \frac{1}{N}\phi_{x_i,\sigma}(x_j), \tag{5.6}$$

where N is the population size, $\sigma > 0$ is a fixed parameter, and $\phi_{\mu,\sigma}(x)$ is the probability density function of a normal random variable with mean μ and standard deviation σ evaluated at x. This has the effect of imposing stronger interaction among more similar individuals.

5.2.2 Stochastic Component of BMI dynamics

To account for the natural variation of individuals' BMI over relatively short time scales, e.g. due to fluctuations in diet and activity level, we append a noise term to Eq. (5.5), yielding the following stochastic differential equation model.

$$\frac{dx_i}{dt} = a_N(x_i; \vec{x}) + \xi(x_i, t) \ , \tag{5.7}$$

where $a_N(x_i; \vec{x})$ is the right-hand-side of Eq. (5.5), $\xi(x_i, t)$ is a stochastic random variable, and where the stochastic differential equation is specified according to the Itô interpretation. In order to simplify our model, we assume the separable form

$$\xi(x_i, t) = b(x_i)\eta(t) \ ,$$

where $\eta(t)$ represents Gaussian white noise, i.e., $\langle \eta(t) \rangle = 0$ and $\langle \eta(t)\eta(\tilde{t}) \rangle = \delta(t - \tilde{t})$, and where $b(x_i)$ is a deterministic function of x_i. Substituting this assumption into Eq. (5.7) yields the Langevin equation

$$\frac{dx_i}{dt} = a_N(x_i; \vec{x}) + b(x_i)\eta(t) \ . \tag{5.8}$$

We assume that each individual i behaves according to their *own* Langevin equation of the form given in Eq. (5.8), resulting in a system of coupled stochastic differential equations for finite populations, i.e. for $N < \infty$. However, although each individual behaves according to their *own* Langevin equation we note that in the limit of large population size, i.e.

as $N \to \infty$, each individual Langevin equation (5.8) can be approximated by a *single* Langevin equation, i.e. by

$$\frac{dx_i}{dt} = a(x_i) + b(x_i)\eta(t) \ , \tag{5.9}$$

where

$$\begin{aligned} a(x_i) &= k_I(x^\star - x_i) + \lim_{N\to\infty} a_N(x_i;\vec{x}) \\ &= k_I(x^\star - x_i) + k_S \lim_{N\to\infty} \sum_{j=1}^{N} \left[A_{ij}(x_j - x_i) - \frac{1}{2}\frac{\partial A_{ij}}{\partial x_i}(x_j - x_i)^2 \right] \\ &= k_I(x^\star - x_i) + k_S \int_0^\infty \left[\phi_{x_i,\sigma}(x_j)\ (x_j - x_i) - \frac{1}{2}\frac{\partial \phi_{x_i,\sigma}(x_j)}{\partial x_i}\ (x_j - x_i)^2 \right] p(x_j;t)dx_j \end{aligned} \tag{5.10}$$

and where $p(x,t)$ is the probability density function for BMI x at time t. In the limit of large population size the dynamics of individual i's BMI (x_i) is independent of the BMI of any other single individual in the population. In other words, the contribution of the BMI of individual $j \neq i$ (x_j) to the dynamics of individual i's BMI (x_i) is null, i.e. $a(x_i)$ really is a function of only a single variable (x_i) and does not depend on any other single x_j (for $j \neq i$). Because in the limit of large population size each individual behaves identically according to a single Langevin equation, i.e. to Eq. (5.9), we can approximate the dynamics of the BMI density $p(x;t)$ by the following advection-diffusion equation (also known as a Fokker-Planck equation) [68].

$$\frac{\partial p}{\partial t}(x,t) = -\frac{\partial}{\partial x}[p(x,t)a(x)] + \frac{1}{2}\frac{\partial^2}{\partial x^2}[p(x,t)b^2(x)] \ , \tag{5.11}$$

where we have dropped the subscript i for convenience of notation, $a(x)$ (see Eq. (5.10)) is called the advection term, and $b(x)$ is called the diffusion term. In order to fully specify the model, it remains to determine a functional form for $b(x)$. Before doing so, however, it is useful to review the physical interpretation of $a(x)$ and $b(x)$ in Eqs. (5.9) and (5.11).

The advection term $a(x)$ in the Langevin equation (5.9) represents a deterministic force acting on an individual's BMI. In the context of our model, this force is a result of an individual's desire to adjust their BMI in order to maximize their utility. In the absence of the diffusive term, i.e. when $b(x) \equiv 0$, the Langevin equation (5.9) implies that each individual's BMI will converge to the root of $a(x)$. Analogously, in the absence of the diffusive term, the advection-diffusion equation (5.11) implies that the probability density function for BMI will converge pointwise to Dirac's delta function centred at the root of $a(x)$.

The diffusion term $b(x)$ in the Langevin equation represents random noise acting on an individual's BMI. In the context of our model, this noise is a result of shocks to the individual that interfere with the planned trajectory towards their ideal weight. In the absence of the advective term, i.e. when $a(x) \equiv 0$, the Langevin equation (5.9) implies that each individual's BMI will resemble a random walk, i.e. a Wiener process or Brownian motion. Analogously, in the absence of the advective term the advection-diffusion equation (5.11) implies that the probability density function for BMI will spread out with diverging variance, eventually converging pointwise to the zero function.

When both advection and diffusion terms are present, the equilibrium BMI distribution $p_{eq}(x)$, which is the solution to

$$0 = -\frac{\partial}{\partial x}[p_{eq}(x)a(x)] + \frac{1}{2}\frac{\partial^2}{\partial x^2}[p_{eq}(x)b^2(x)] \ , \tag{5.12}$$

results from the balance between the advection and diffusion terms, i.e. between $a(x)$ and $b(x)$, respectively. Specifically, advection acts to concentrate the distribution around the root of $a(x)$, and diffusion acts to spread out the distribution.

In addition to the discussion above, we can also understand the role of the advection and diffusion terms in the Langevin equation, i.e. Eq. (5.9), as follows. Let $W(t)$ denote a Wiener process defined by $dW(t) = \eta(t)dt$. Writing Eq. (5.9) in terms of differentials, we find

$$dx = a(x)dt + b(x)\eta(t)dt = a(x)dt + b(x)dW \ .$$

Then conditioning on the initial BMI yields

$$\begin{aligned} \text{mean}(dx) &= \mathbb{E}\left[dx\right] = \mathbb{E}\left[a(x)dt\right] + \mathbb{E}\left[b(x)dW\right] = a(x)dt, & (5.13) \\ \text{variance}(dx) &= \mathbb{E}\left[dx^2\right] - \mathbb{E}\left[dx\right]^2 \\ &= \mathbb{E}\left[a^2(x)dt^2\right] + 2\mathbb{E}\left[a(x)b(x)dWdt\right] + \mathbb{E}\left[b(x)^2dW^2\right] - (a(x)dt)^2 \\ &= b(x)^2dt \ , & (5.14) \end{aligned}$$

since $\mathbb{E}\left[dW\right] = 0$ and $\mathbb{E}\left[dW^2\right] = dt$. It follows that, on average, change in individuals' BMI follows the advective term $a(x)$, and the variance in the change in individuals' BMI follows $b^2(x)$. Equivalently, we expect the standard deviation of individual i's BMI to scale with $b(x_i)\sqrt{dt}$.

Consistent with data [81], we make the simplifying assumption that fluctuations in an individual's BMI are of roughly constant magnitude when measured as a percentage of their current BMI. We therefore take

$$b(x) = \sqrt{k_b}\ x \ , \tag{5.15}$$

where $k_b > 0$ is a constant parameter, so that the standard deviation is proportional to the current BMI.

The mathematical model presented above has implications at both the population (macro) and individual (micro) levels which we detail fully in Section 5.4. In Section 5.4.1 we outline the population level implications of our model, which follow from solving Eq. (5.12) for the equilibrium BMI distribution $p_{eq}(x)$. In Section 5.4.2 we outline the individual level implications of our model, which follow from Eqs. (5.13) and (5.14). However, since all of the implications discussed in Section 5.4 are tested against empirical BMI data, we briefly summarize the BMI data used in this study in the following section.

5.3 Data

In order to calibrate the model developed in the previous section we require two different types of BMI data. At the population level we consider the empirical BMI distribution, modelled by the solution to Eq. (5.12). We can compute empirical BMI distributions from three independently collected data sets: the National Health and Nutrition Examination Survey (NHANES), medical records for patients of the Northwestern Medicine system of hospitals and clinics (NU), and the Behavioral Risk Factor Surveillance System (BRFSS). At the individual level we consider the average and standard deviation in year-over-year change in individuals' BMI, see Eqs. (5.13) and (5.14), respectively. We can compute the year-over-year change in individuals' BMI from two independently collected data sets: the NHANES and NU data sets. We briefly describe the NHANES, NU, and BRFSS data sets briefly below.

5.3.1 National Health and Nutrition Examination Survey

The National Health and Nutrition Examination Survey (NHANES) refers to a series of studies designed to collect a representative sample of health and nutrition data for both adults and children (approximately 5,000 individuals total per year) in the United States [32]. NHANES data are available for survey years 1999-2000, 2001-2002, ..., 2011-2012. Directly measured BMI data are available from measurements taken during a physical exam. These data are used to compute empirical BMI distributions for each survey year. In addition, during an interview individuals were asked to self-report their current weight and height, as well as their weight from the preceding year. These measurements allow us to calculate self-reported change in BMI over the year preceding the interview. Note: we

only use NHANES data for individuals 18 years or older at the time of the survey. For additional information, see Appendix E.1.1.

5.3.2 Northwestern Medicine Medical Records

The Northwestern Medicine (NU) data set consists of medical records for patients of the Northwestern Medicine system of hospitals and clinics. Anonymized medical records were available for 750,081 distinct patients from 1997 through 2014, with the majority of records coming from later years. We calculate a total of 2,056,400 BMI data points from weight and height data for individuals in this data set that are at least 18 years of age or older. We use these data to compute the empirical BMI distribution for each year. In addition, we are able to calculate the year-over-year change in BMI for all individuals with patient records in consecutive years. We emphasize that this data set provides the most abundant source of individual level data (comprising a total of 1,017,518 data points for year-over-year change in individuals' BMI from 329,543 distinct individuals). Thus, although these data do not form a representative sample (since they are comprised of medical records), they are nevertheless extremely valuable since they represents the largest data set of its type. For additional details, see Appendix E.1.2.

5.3.3 Behavioral Risk Factor Surveillance System

The Behavioral Risk Factor Surveillance System (BRFSS) refers to a series of telephone surveys designed to collect a representative sample of health data for adults (aged 18 years or older) in the United States [31]. BRFSS data are available for survey years 1984, 1985, ..., 2013. We note that prior to 2011 BRFSS surveys were conducted over land lines only, whereas from 2011 onward BRFSS methodology has been modified to include cell phones as well. We also note that many states did not participate in early BRFSS surveys. Therefore, for the purposes of this study we only consider surveys from 1987 (the first year where a majority of states participated in the BRFSS) onwards. The number of individual records for each BRFSS survey increases from approximately 50,000 in 1987, to approximately 135,000 in 1997, to more than 400,000 from 2007 onward. For each BRFSS survey we extract the BMI of each individual surveyed and use this data to compute the empirical BMI for that year. We note that since these data are gathered using telephone interviews, the weight and height measurements (used in calculating BMI) are all self-reported, in contrast to the NHANES and NU data sets. Also in contrast to NHANES and NU data sets, the BRFSS data does not provide sufficient data for us to compute the year-over-year change in individuals' BMI. For additional details, see Appendix E.1.3.

5.4 Methods for Testing Model Implications

We now detail the implications of our model at the population and individual levels. Specifically, at the population level our model predicts that the BMI distribution can be approximated by the solution to Eq. (5.12) (see Section 5.4.1), whereas at the individual level our model predicts that the mean and variance of the year-over-year change in individuals' BMI should follow Eqs. (5.13) and (5.14) (see Section 5.4.2).

5.4.1 Methods for Population-level Model Implications

Given $a(\cdot)$ and $b(\cdot)$ as in Eqs. (5.10) and (5.15) and under the additional assumption that $k_S = 0$ in Eq. (5.10), i.e. that social interactions between individuals have no effect on the dynamics of individuals' BMI, solving Eq. (5.12) yields the following distribution

$$p_{eq}^{(0)}(x) = p_{eq}^{(0)}(x; k_I/k_b, x^\star) \propto x^{-2(k_I/k_b+1)} \exp\left(-2\frac{k_I}{k_b}\frac{x^\star}{x}\right) . \tag{5.16}$$

We note that since $p_{eq}^{(0)}(x) \sim x^{2-(k_I/k_b+1)}$ as $x \to \infty$, $p_{eq}^{(0)}(x)$ is a scale-free (or power law) distribution. The mean, mode, variance[5], skewness, and mode skewness of this distribution can be expressed in terms of $x^\star$ and $k_0 = k_I/k_b$ and are recorded in Table 5.1, see Appendix E.2.1 for detailed calculations. In contrast, when we assume that $k_S > 0$ in Eq. (5.10), i.e. that social interactions between individuals do have an effect on the dynamics of individual BMIs (through the interaction kernel A_{ij} as specified in Eq. (5.6)), there is no closed form solution to Eq. (5.12). Instead, for $k_S > 0$ the solution to Eq. (5.12), which we denote by $p_{eq}(x) = p_{eq}(x; k_I/k_b, x^\star, k_S/k_b, \sigma)$, must be calculated numerically, see Appendix E.2.2.

Next, suppose that the parameters of the Fokker-Planck equation were constant. In this case, the distribution function $p(x,t)$ will converge asymptotically to the equilibrium distribution, i.e. to $p_{eq}^{(0)}(x; k_I/k_b, x^\star)$ when $k_S = 0$ or $p_{eq}(x; k_I/k_b, x^\star, k_S/k_b, \sigma)$ when $k_S > 0$. Next, suppose that the parameters of the Fokker-Planck equation have their own dynamics. In this case, the functions $p_{eq}^{(0)}(x; k_I/k_b, x^\star)$ and $p_{eq}(x; k_I/k_b, x^\star, k_S/k_b, \sigma)$ are not truly equilibrium distributions (since the parameters k_I/k_b, $x^\star$, k_S/k_b, and σ are not constant). Nevertheless, if we assume that changes in the parameters occur over a much slower time scale than the dynamics of individuals' BMI, then it follows that the BMI distribution

[5] We note that in order for the variance to be non-negative we require that $2k_0 - 1 > 0$. This is condition is satisfied in all the empirical BMI distributions considered in this study.

Table 5.1: Properties of the equilibrium distribution function for the non-social model $p_{eq}^{(0)}(x; k_0, x^\star)$ (Eq. (5.16))

Property	Value
Mean	$x^\star$
Mode	$k_0 x^\star/(k_0+1)$
Variance[5]	$x^{\star 2}/(2k_0-1)$
Skewness	$2\sqrt{2k_0-1}/(k_0-1)$
Mode Skewness	$\sqrt{2k_0-1}(k_0+1)$

should be well approximated by $p_{eq}^{(0)}(x; k_I/k_b, x^\star)$ when $k_S = 0$, i.e. when we assume that social interactions between individuals do not have an effect on the dynamics of individual BMIs, and by $p_{eq}(x; k_I/k_b, x^\star, k_S/k_b, \sigma)$ when $k_S > 0$, i.e. when we assume that social interactions between individuals do have an effect on the dynamics of individual BMIs. We compare the ability of $p_{eq}^{(0)}(x; k_I/k_b, x^\star)$ and $p_{eq}(x; k_I/k_b, x^\star, k_S/k_b, \sigma)$ to fit empirical BMI distributions to each other, and to two other candidate distribution functions commonly used to describe right-skewed data: the log-normal probability distribution function

$$f_{\log}(x; \mu, \sigma) = \frac{1}{x\sqrt{2\pi\sigma^2}} \exp\left[-\frac{(\log x - \mu)^2}{2\sigma^2}\right], \tag{5.17}$$

and the skew-normal probability distribution function

$$f_{\text{skew}}(x; \xi, \omega, \alpha) = \frac{2}{\omega}\phi_{0,1}\left(\frac{x-\xi}{\omega}\right)\Phi_{0,1}\left[\alpha\left(\frac{x-\xi}{\omega}\right)\right], \tag{5.18}$$

respectively, where $\Phi_{\mu,\sigma}(x)$ is the cumulative distribution function for a normal random variable with mean μ and standard deviation σ evaluated at x. For additional details on fitting empirical BMI distributions, see Appendix E.2.3.

5.4.2 Methods for Individual-level Model Implications

Given empirical BMI measurements from individuals at two different time points t_1 and t_2 we can approximate $a(\cdot)$ and $b(\cdot)$ using Eqs. (5.13) and (5.14) with

$$dx_i(t_1) \approx \Delta x_i(t_1) = x_i(t_2) - x_i(t_1).$$

Then

$$\begin{aligned}\hat{a}(x_i(t_1)) &= \text{mean}\left(\Delta x_i(t_1)/\Delta t\right) \\ &\approx a(x_i(t_1)) \\ &= k_I(x^\star - x_i(t_1)) + k_S \underbrace{\sum_{j=1}^{N}\left[A_{ij}\ (x_j(t_1) - x_i(t_1)) - \frac{1}{2}\frac{\partial A_{ij}}{\partial x_i}\ (x_j(t_1) - x_i(t_1))^2\right]}_{=u'_S(x_i(t_1))} ,\end{aligned} \tag{5.19}$$

$$\begin{aligned}\hat{b}(x_i(t_1)) &= \sqrt{\text{variance}(\Delta x_i(t_1))/\Delta t} \\ &\approx b(x_i(t_1)) \\ &= \sqrt{k_b}\ x_i(t_1) ,\end{aligned} \tag{5.20}$$

where $\Delta t = t_2 - t_1$ and the mean and variance are taken over bins comprised of similar $x_i(t_1)$ values. We compare the performance of our model under the assumption that social interactions between individuals do have an effect on the dynamics of individual BMIs, i.e. when $k_S > 0$ in Eq. (5.10), to the performance of our model under the assumption that social interactions between individuals do not have an effect on the dynamics of individual BMIs, i.e. when $k_S = 0$ in Eq. (5.10). In the first case, to estimate the parameters k_I, $x^\star$, k_S, and σ we regress $\hat{a}(x_i(t_1))$ on $x_i(t_1)$, $u'_S(x_i(t_1))$, and a constant vector. In the second case, to estimate the parameters k_I and $x^\star$ ($k_S = 0$ and σ is undetermined) we regress $\hat{a}(x_i(t_1))$ on $x_i(t_1)$ and a constant vector. To estimate the parameter $\sqrt{k_b}$ we regress $\hat{b}(x_i(t_1))$ on $x_i(t_1)$. We provide additional details on estimating the parameters using individual level data in Appendix E.2.4.

We note that the methodology presented in this section can only be applied to NHANES and NU BMI data, because these are the only data sets that have information on how individuals' BMI changes over time.

Testing for Significance of Social Utility Contribution using Synthetic data

By necessity, our social model (i.e. $k_S > 0$ in Eq. (5.10)) will have lower error than our nonsocial model (i.e. $k_S = 0$ in Eq. (5.10)), since it has two additional parameters than the nonsocial model. In other words, if $S(k_I, x^\star, k_S, \sigma)$ denotes the objective function that is minimized in order to estimate the parameters k_I, $x^\star$, k_S, and σ then

$$\hat{S} = S(\hat{k}_I, \hat{x}^\star, \hat{k}_S, \hat{\sigma}) = \min_{(k_I, x^\star, k_S, \sigma)} S(k_I, x^\star, k_S, \sigma) \leq \min_{(k_I, x^\star)} S(k_I, x^\star, 0, -) = \hat{S}\Big|_{k_S=0} .$$

In order to investigate the significance of the social utility contribution to $a(\cdot)$ we estimate the expected error reduction

$$\Delta\hat{S}_{syn} = \hat{S}_{syn}\Big|_{k_S=0} - \hat{S}_{syn}$$

under ideal circumstances, i.e. under the assumption that the data was generated by our social model. If the empirically observed reduction in error $\Delta\hat{S}$ is a similar magnitude to the ideal expected error reduction $\Delta\hat{S}_{syn}$, then this is evidence that the effect of social utility is significant. Conversely, if the reduction in error $\Delta\hat{S}$ is much smaller than the ideal expected error reduction $\Delta\hat{S}_{syn}$, then this is evidence that the effect of social utility is not significant. To compute the expected error reduction $\Delta\hat{S}_{syn}$ we generate a synthetic data set for each initial year t_1 of NHANES and NU data by simulating Eq. (5.9) from $t = t_1$ to $t = t_2$ using estimated of parameters $\hat{k}_I$, $\hat{x}^\star$, $\hat{k}_S$, and $\hat{\sigma}$ and with initial conditions $x_i(t_1)$ taken from the empirical NHANES and NU BMI data, see Appendix E.2.5. The social and the nonsocial models are then fit to this synthetic data set and used to compute $\Delta\hat{S}_{syn,1}$. We repeat this procedure to generate a total of 100 realizations and then set

$$\Delta\hat{S}_{syn} = \frac{1}{100}\sum_{i=1}^{100}\Delta\hat{S}_{syn,i} \ .$$

5.5 Results

5.5.1 Population-level Model Implications

The key prediction of our model at the population level is the distribution of BMIs. Figures 5.1-5.2 displays the results from fitting four distributions to the observed empirical BMI data: (1) our model without social effects $p^{(0)}_{eq}(x; k_I/k_b, x^\star)$ (Eq. (5.16), solid red), (2) our model with social effects $p_{eq}(x; k_I/k_b, x^\star, k_S/k_b, \sigma)$ (Eq. (E.3), dashed red), (3) the log-normal distribution (null model Eq. (5.17), dash-dotted blue), and (4) the skew-normal distribution (null model Eq. (5.18), dotted green). From left to right, Fig. 5.1 displays model distributions and empirical data for the 2011-2012 NHANES survey data, 2011 NU data, and 2011 BRFSS survey data. Each of the three rows illustrates the data from a different perspective. The top row plots the BMI distributions on a linear scale and allows us to confirm visually that the BMI distributions are right-skewed. The middle row plots the BMI distributions on a logarithmic scale and illustrates how both our non-social $p^{(0)}_{eq}(x)$ and social $p_{eq}(x)$ models reproduce the right tail of the BMI distributions better

than either the log-normal or skew-normal distribution. The bottom row of panels displays the deviation from the log-normal distribution. This allows us to illustrate how both our social $p_{eq}(x)$ models reproduces the BMI distributions better than the non-social $p_{eq}^{(0)}(x)$, log-normal, or skew-normal distribution.

In Fig. 5.2 we display from left to right the root mean square error for NHANES survey data, NU data, and BRFSS survey data for all years where data are available. Empirical data and fitted distributions for the remaining years of NHANES, NU, and BRFSS data are displayed in Figs. E.1-E.3. We observe that our nonsocial model (Eq. (5.10) with $k_S = 0$), whose equilibrium distribution $p_{eq}^{(0)}(x)$ is given in Eq. (5.16), outperforms the log-normal distribution, since (a) both $p_{eq}^{(0)}(x)$ and log-normal distributions have two parameters (i.e., two degrees of freedom), and (b) the $p_{eq}^{(0)}(x)$ distribution results in lower error. We also observe that our nonsocial model outperforms the skew-normal distribution, since (a) fitting both distributions results in a similar amount of error and (b) the skew-normal distribution has one more parameter than the $p_{eq}^{(0)}(x)$ distribution.We emphasize that, even without the social component (i.e. with $k_S = 0$), our model is able to capture the BMI distribution better than the log-normal and skew-normal distributions, i.e. two competing distribution functions commonly used to fit right-skewed data. This is the case even though one of the competing distributions, i.e. the skew-normal distribution, has an additional parameter. This provides strong evidence in support of our model, since if our model did not capture the essence of the dynamics of individuals' BMI then it would be unlikely to perform as well as it has against the log-normal and skew-normal distributions.

In addition to our nonsocial model providing a better fit for the BMI distribution than either the skew-normal or log-normal distributions, we note that our nonsocial model provides a potential explanation for the right-skewness observed in the empirical BMI data. Recall that the BMI distribution predicted by our model, which results from solving the advection-diffusion equation (5.11), is determined by the balance of the advection term $a(x)$ (see Eq. (5.10) with $k_S = 0$), which tends to concentrate the BMI distribution function around the root of $a(x)$, and $b(x)$ (see Eq. (5.15)), which tends to diffuse the BMI distribution. We observe, therefore, that in our model the right-skewness follows from the assumption that the standard deviation of individuals' change in BMI is proportional to the individuals' BMI, i.e. $b(x) = \sqrt{k_b}\, x$ (as specified in Eq. (5.15)). In other words, the right-skewness follows from the assumption that an individual's BMI dynamics are subject to random fluctuations in lifestyle (e.g. giving in to the temptation to eat high calorie food, spending extra time engaged in physical activity, etc...) whose magnitude scale roughly with their BMI. Alternatively, we can illustrate just how critical the choice of $b(x) = \sqrt{k_b}\, x$ is to the skewness of the equilibrium distribution by solving Eq. (5.12) with

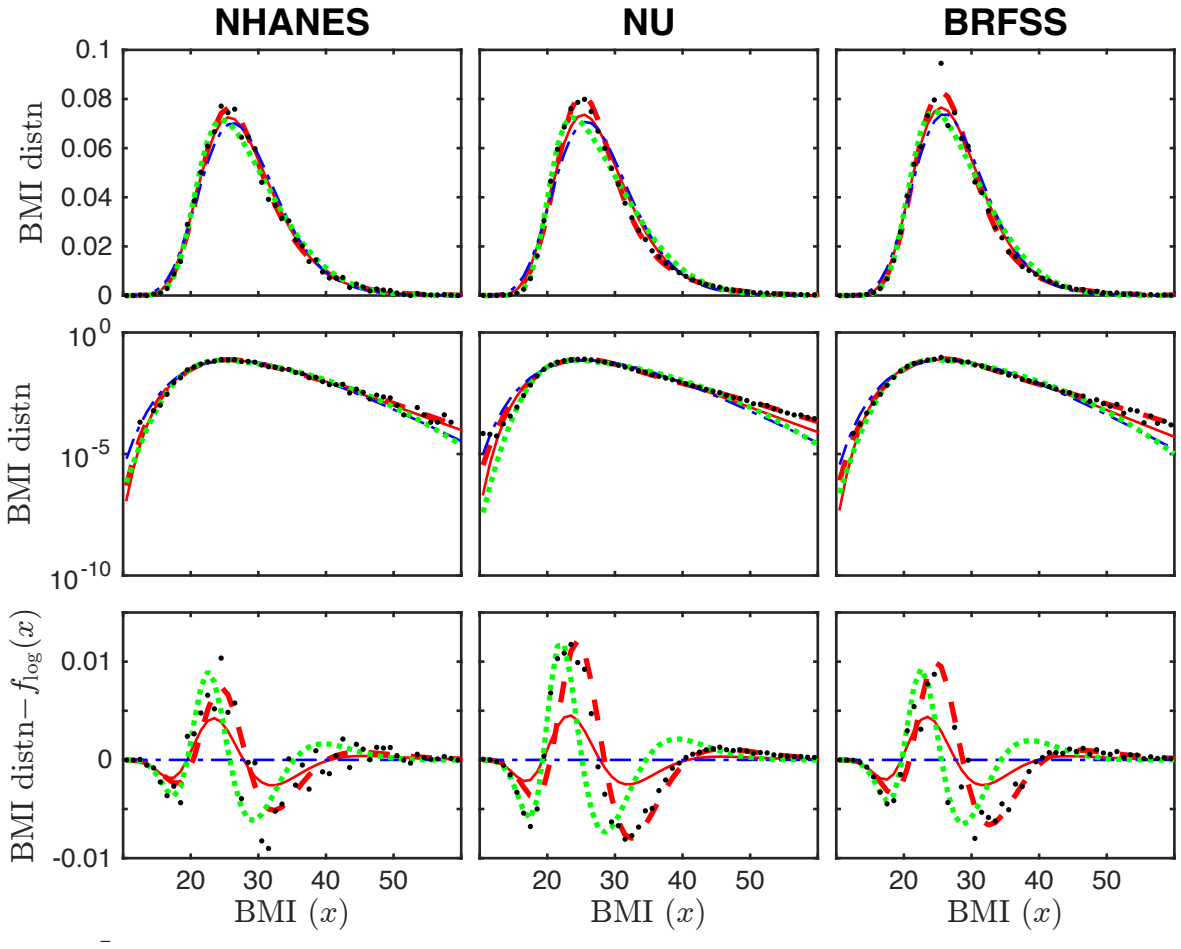

Figure 5.1: Results from fitting BMI distributions. From left to right we display data for NHANES, NU, and BRFSS data sets, respectively. Each panel displays 2011 empirical BMI data (solid black dots) as well as four fitted probability distribution functions: model without social utility (Eq. (5.16)—solid red), model with social utility (Eq. (E.3)—dashed red), log-normal null distribution (Eq. (5.17)—dash-dotted blue), and skew normal null distribution (Eq. (5.18)—dotted green).

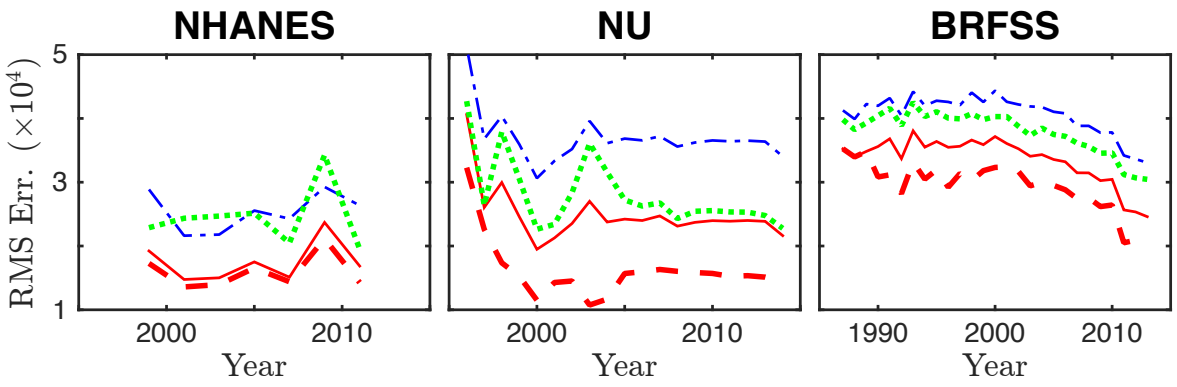

Figure 5.2: Results from fitting four BMI probability distribution functions: model without social utility (Eq. (5.16)—solid red), model with social utility (Eq. (E.3)—dashed red), log-normal null distribution (Eq. (5.17)—dash-dotted blue), and skew normal null distribution (Eq. (5.18)—dotted green). From left to right we display data for NHANES, NU, and BRFSS data sets, respectively. Each panels displays the root mean square error from fitting models to empirical BMI distributions in all years where data is available.

$a(x)$ as in Eq. (5.10) (still with $k_S = 0$) and with $b(x) = \sqrt{k_b} =$ constant. Under these circumstances the equilibrium distribution is a Gaussian centered at the root of $a(x)$, i.e. when $b(x)$ is a constant the equilibrium distribution becomes symmetric and is no longer skewed.

Next, because our social model (Eq. (5.10) with $k_S > 0$), whose equilibrium distribution $p_{eq}(x)$ is given in Eq. (E.3) of Appendix E.2.2, has two more parameters than the log-normal and nonsocial model $p^{(0)}_{eq}(x)$, and one parameter more than the skew-normal distribution, we expect it to result in the lowest error of all the investigated distributions when used to fit empirical BMI distribution data. This is consistent with our findings (see right panels of Fig. 5.1). We note, however, that at this stage we cannot be certain that the inclusion of social utility ($k_S > 0$ in Eq. (5.10)), and the resulting increase in model complexity, is justified by the modest improvement in the goodness of fit. In other words, it is not yet clear whether our social model results in a better fit to empirical BMI distribution data than our nonsocial model because (a) it captures an additional important effect (i.e. the effect of social interactions on the dynamics of the BMI), or (b) it results in overfitting.

We address the issue of overfitting within our social model by computing the Akaike Information Criterion (AIC) [4], and resulting relative likelihood, for each of the proposed distribution functions in each year for which data are available. Representative results for the year 2011 are displayed in Table 5.2, full results are displayed in Tables E.1-E.3 in

Appendix E.3.1. With the exception of years 1994-1996 of the BRFSS data set, we find that the social model $p_{eq}(x)$ is overwhelmingly the most likely of the four models that we tested. Indeed, with the exception of years 1994-1996 of the BRFSS data set the non-social $p_{eq}^{(0)}(x)$, log-normal, and skew-normal distributions are all less than 1.5% as likely as the social model $p_{eq}(x)$. For the years 1994-1996 of the BRFSS data set the non-social model $p_{eq}^{(0)}(x)$ is the most likely of the four distributions tested, but the social model cannot necessarily be excluded by the AIC criteria. Specifically, for the years 1994-1996 of the BRFSS data set the social model $p_{eq}(x)$ is 14% as likely as the non-social model, while the log-normal and skew-normal distributions are less than $2.1 \times 10^{-288}\%$ as likely as the non-social model. Taken as a whole, these data are strongly support the hypothesis that our social distribution results in a better fit to empirical BMI distribution data than the non-social distribution, the log-normal distribution, and the skew-normal distribution because it captures an additional important effect, and not because of overfitting. We remark that we will return to the issue of overfitting in the following section, where we apply our model to individual level data in order to investigate the importance of social interaction between individuals on the dynamics of individuals' BMI.

Table 5.2: Relative likelihood of non-social $p_{eq}^{(0)}(x)$, social $p_{eq}(x)$, log-normal $f_{\log}(x)$, and skew-normal $f_{\text{skew}}(x)$ models for 2011 NHANES, NU, and BRFSS empirical BMI distributions

	Relative Likelihood $\exp[(AIC_{\min} - AIC_i)/2]$			
Data	$p_{eq}^{(0)}(x)$	$p_{eq}(x)$	$f_{\log}(x)$	$f_{\text{skew}}(x)$
NHANES	1.3×10^{-2}	1	1.4×10^{-21}	1.1×10^{-9}
NU	$< 10^{-300}$	1	$< 10^{-300}$	$< 10^{-300}$
BRFSS	$< 10^{-300}$	1	$< 10^{-300}$	$< 10^{-300}$

Finally, we recall that the parameters $k_0 = k_I/k_b$ and $x^\star$ in the distribution $p_{eq}^{(0)}$, given in Eq. (5.16), can be related to properties of the distribution (e.g., the mean, variance, skewness, etc...) in a straightforward way, see Table 5.1. Thus, comparing parameter estimates between different years gives us useful insight into how BMI distributions differ from each other. For additional details, we refer the reader to Appendix E.3.1 where the parameter estimates resulting from fitting $p_{eq}^{(0)}(x)$ to empirical BMI data are plotted for all three data sets.

5.5.2 Individual-level Model Implications

At the individual level our model predicts that both the average and standard deviation of year-over-year change in individuals' BMI are approximately linear functions of BMI. Parameter estimates for $a(x)$ and $b(x)$ (see Eqs. (5.10) and (5.15)) are given in Tables 5.4 and 5.3, respectively, and representative results are displayed in Fig. 5.3, see Figs. E.5-E.6. Table 5.4 also displays results from the synthetic data experiment described in Section 5.4.2 above.

Table 5.3: Parameter $\sqrt{k_b}$ estimated from individual level BMI data (fitting $\sqrt{k_b}$ in Eq. (5.20)). All parameters estimates are significant at the 95% confidence level.

Data	$\widehat{\sqrt{k_b}}$	L_2-Error
NHANES	0.083	0.482
NU	0.071	0.461
Synthetic NHANES	0.077	0.271
Synthetic NU	0.069	0.093

We observe from Fig. 5.3 that our model with social effects (i.e. with $k_S > 0$ in Eq. (5.10)) is able to closely reproduce average dynamics of individual BMIs, i.e. the average change in individuals' BMI is closely modelled by $a(x)$ and the standard deviation of the change in individuals' BMI is closely modelled by $b(x)$, see Eqs. (5.19) and (5.20), respectively. However, since the average year-over-year change in individuals' BMI is roughly linear, we may also infer from Fig. 5.3 that our model without social effects (i.e. with $k_S = 0$ in Eq. (5.10)) is also able to closely reproduce both the average and standard deviation of year-over-year change in individuals' BMI. This is consistent with Table 5.4 which shows that, although the fit resulting from the nonsocial model ($k_S = 0$) results in higher error than the fit from the social model ($k_S > 0$), the error resulting from both models is roughly the same order of magnitude, i.e. they are within 3.5% of each other for both NHANES and NU data sets. Specifically, we observe a 3.1% reduction in error for the NHANES data set and a 2.4% reduction in error for the NU data set. It is worthwhile noting that these two figures are remarkably consistent between themselves, especially when we recall that the NU data set is comprised of medical records and is not a representative sample of the population.

To better understand these numbers, we compare these results to the expected error reduction when going from the nonsocial to the social model under "ideal" conditions, i.e., when synthetic data are directly generated by the social model. The expected error

Table 5.4: Parameters k_I, $x^\star$, k_S, and σ estimated from individual level BMI data. We compare the performance of our full model (fitting parameters k_I, $x^\star$, k_S, and σ in Eq. (5.19)) and our model without social utility (fitting parameters k_I and $x^\star$ in Eq. (5.19) with $k_S = 0$). All parameter estimates are significant at the 95% confidence level.

Data	Model	Parameters				L_2-Error	ΔL_2-Error (%)
		$\hat{k}_I$	$\hat{x^\star}$	$\hat{k}_S$	$\hat{\sigma}$		
NHANES	nonsocial	0.124	28.0	0	–	0.542	–
	social	0.144	27.9	21.5	2.18	0.525	3.1
NU	nonsocial	0.059	28.0	0	–	0.374	–
	social	0.069	28.0	9.44	3.44	0.365	2.4
Synth. NHANES	nonsocial	0.115	28.0	0	–	0.325	–
	social	0.132	28.0	17.5	2.23	0.306	5.9
Synth. NU	nonsocial	0.057	28.0	0	–	0.128	–
	social	0.066	28.0	8.21	3.43	0.106	17.2

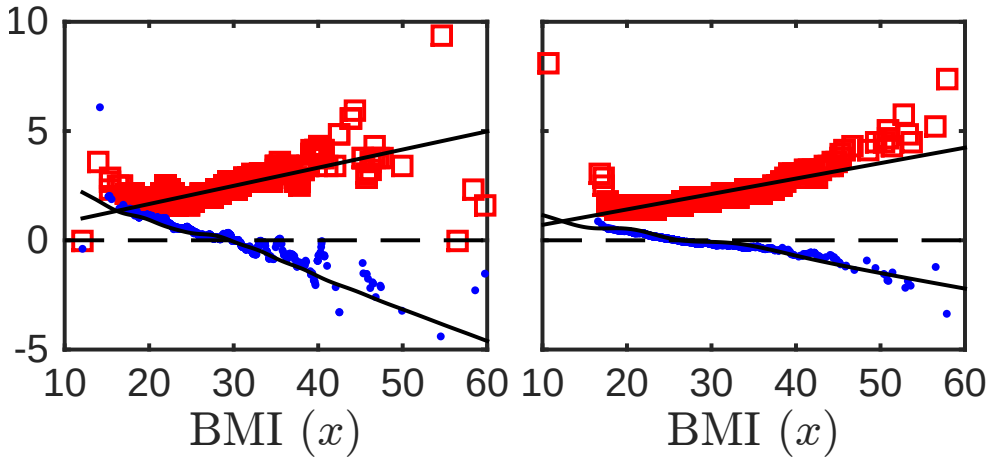

Figure 5.3: Average and standard deviation in year-over-year change in individuals' BMI. The line of best fit for average year-over-year change in individuals' BMI (blue dots) is given by the regression specified in Eq. (5.19) (solid black line). The line of best fit for standard deviation for year-over-year in individuals' BMI (red squares) is given by the regression specified in Eq. (5.20) (solid black line). The line $y = 0$ is given by a dashed black line. 2011-2012 NHANES survey data (left) and for 2011 NU data (right). Note: to improve visualization of data we display only 1,000 data points for each of $\hat{a}\,(x)$ and $\hat{b}\,(x)$, selected uniformly at random.

reduction is estimated to be 5.9% for synthetic NHANES and 17.2% for synthetic NU data, see Table 5.4. In the case of the NHANES data, the comparison is quite favorable since we observe a reduction that is 53% of the ideal expected error reduction (3.1% out of 5.9%). This provides substantial evidence that social effects play an important role in individual BMI dynamics. In the case of NU data, the comparison is more ambiguous since we only observe a reduction that is 14% of the ideal expected error reduction (2.4% out of 17.2%). However, although the result for the NU data may not provide as strong evidence as the NHANES data, the NU results are still largely consistent with the hypothesis that social effects play an important role in BMI dynamics.

Another interesting implication of the data shown in Fig. 5.3 is that the average change in individual BMI (blue dots) is negative for high BMI and positive for low BMI. This indicates that individuals with high BMI are more likely to lose weight than individuals of low BMI (at least over the period of a single year). This is contrary to what one might expect from the "runaway train" argument, i.e., that individuals with high BMI enter a positive feedback loop resulting in ever increasing weight gain.

We note that the results of this section are consistent with those of Section 5.5.1. In fact, to the best of the authors' knowledge, this is the first quantitative mathematical model able to simultaneously reproduce both population and individual level BMI data. We emphasize that the consistency of the results at the population and individual levels mutually support each other, in the sense that if the model did not capture the essential features of the dynamics the BMI then it would be unlikely that our model would be able to simultaneously reproduce both the population and individual level dynamics.

5.6 Discussion

5.6.1 Limitations

We previously noted that our mathematical model is capable of reproducing the right-skewness observed in the empirical BMI distribution data without making use of a "runaway train" type of phenomenon describing individuals' tendency to gain weight (i.e. to increase their BMI). This is somewhat surprising, given the intuitively appealing nature of the "runaway train" argument. The lack of a "runaway train" in our model can be seen by examining the Langevin equation for individuals' behaviour Eq. (5.9) and observing that over an infinitesimal time interval the dynamics of high BMI individuals are not biased toward further increasing their BMI. Specifically, for both our nonsocial and social models

(a) the deterministic contribution of the Langevin equation (i.e. $a(x)$ as in Eq. (5.10)) results in downward pressure on individuals with high BMI (e.g. obese individuals), and (b) the stochastic contribution of the Langevin equation (i.e. $b(x)$ as in Eq. (5.15)) is equally likely to result in shocks that increase or decrease BMI. This first feature of our model is confirmed in the BMI data, for example see Fig. 5.3 which shows that the average change in individual BMI is negative for high BMI and positive for low BMI. This indicates that individuals with high BMI are more likely to lose weight than individuals of low BMI (at least over the period of a single year). Although a "runaway train" phenomenon does not manifest itself in our model, or in the first two moments of the individual level BMI data (e.g. see Fig. 5.3), this does not preclude the "runaway train" phenomenon as an important factor in the dynamics of individuals' BMI. For example, a "runaway train" type of phenomenon might manifest itself in higher moments of the change in BMI dx (see Eqs. (5.19) and (5.20) for the first two moments). This illustrates an important limitation of our mathematical model, which is only able to account for the first two moments of dx (all other moments are assumed to be zero). It is possible that the "runaway train" may manifest itself in higher order moments of dx, for example although high BMI individuals *on average* decrease their BMI over the span of one year this average reduction may be caused by only a few outliers making large reductions in BMI while the remaining high BMI individuals continue to increase in weight.

In addition to the limitations discussed above in the context of the "runaway train" phenomenon, there are several additional aspects of the model that could be improved through future work. For example, in our model we assume that all individuals have the same natural BMI set point $x^\star$. Of course, this is a gross oversimplifications since the natural set point, which is determined by gender, age, race, activity level, diet, etc..., varies significantly between individuals. Therefore, a straightforward generalization of our model might investigate the implications of having individuals' set points drawn from a distribution of possible set points. Alternatively, we assumed a specific form of interpersonal social interaction in Eq. (5.6). A natural question is whether our model is sensitive to this choice, and if so whether other functions may perform as well or better as the one specified in Eq. (5.6).

Finally, we note that each of the data sets that we considered (NHANES, NU, and BRFSS) has their own strengths and weaknesses. In particular, the data available to us are either (a) directly measured (NHANES, NU) or self-reported (BRFSS), or (b) representative samples (NHANES, BRFSS) or non-representative samples (NU). The issue of directly measured or self-reported data is important, since it is possible that individuals asked to self-report are misreporting their weight and height. Indeed, we can illustrate the importance of having directly measured weight and height data by comparing BMI

distributions from NHANES and NU data sets (directly measured data, see Figs. E.1-E.2) to BMI distributions from BRFSS data sets (self-reported data, see Fig. E.3): for many of the empirical BRFSS BMI distributions (self-reported) the data point at BMI = 25.5 (i.e. the data point that represents the fraction of individuals with $25 \leq \text{BMI} < 26$) is a significant outlier whereas this is not the case in NHANES or NU data (directly measured). While there is no way to tell for certain why this is the case, it is interesting to recall that individuals with BMI ≤ 25 are considered normal or underweight, while individuals with BMI > 25 are considered overweight or obese. The issue of non-representative data is also important. In this case, since NU data consist of individuals' medical records, unhealthy individuals may be overrepresented in the NU data set. For example, the first two data points for the NU BMI distributions in Fig. 5.1 (middle, centre panel) illustrates that there is an unusually large fraction of individuals with extremely low BMI in the NU BMI distribution. While the aforementioned limitations of the data are important to note, however, we emphasize that our results are nevertheless consistent across all three data sets. In other words, the conclusions which we draw in this chapter are robust to the limitations in the data.

5.6.2 Conclusions

In this chapter we establish a mathematical model for the dynamics of individuals' BMIs with implications at both the individual and population levels. At the population level we make two main contributions. First, we predict a functional form for the BMI distribution that fits empirical data better than the skew-normal and log-normal distributions, especially in the tail of the distribution. Second, our model provides a mechanistic explanation for the right-skewness observed in empirical BMI distributions: this appears to be a result of random fluctuations with magnitudes that scale roughly with the BMI of the individual. We also note that because there is a straightforward link between the model parameters and the properties of the BMI distribution (see Table 5.1), our model also provides a relatively straightforward way to compare multiple BMI distributions simultaneously. At the individual level we also make two main contributions. Our model reproduces the behavior of both the average and standard deviation in year-over-year change in individuals' BMI (see Fig. 5.3). The ability of this simple mechanistic model to reproduce real-world data, and to connect micro and macro scales, means predictions can be made regarding the impact of proposed interventions. Such an ability may prove valuable in allocating scarce resources for public health improvements.

At both the individual and population levels our model allows us to investigate the role of social and peer influences in BMI dynamics, providing additional evidence in support of

the hypothesis that social effects have an impact on individuals' BMI dynamics. Specifically, at the population level we compute the Akaike Information Criterion (AIC) for our nonsocial and social models, as well as for the log-normal and skew-normal distributions. Using the AIC to compute the relative likelihood we show that the social model is significantly more likely than the alternatives for the vast majority of NHANES, NU, and BRFSS BMI distribution data. At the individual level, consistent with our social model having two more degrees of freedom than our nonsocial model, we find that fitting the empirical micro-level data to our social model results in lower error than fitting to our nonsocial model, i.e. $\Delta\hat{S} = \hat{S}\Big|_{k_S=0} - \hat{S} > 0$. In order to argue that this reduction in error is significant, i.e. that the reduction in error results from capturing an important feature of individuals' BMI dynamics and not from overfitting, we compare the magnitude of $\Delta\hat{S}$ to the expected reduction in error when fitting the nonsocial model versus the social model under ideal circumstances, i.e. $\Delta\hat{S}_{syn} = \hat{S}_{syn}\Big|_{k_S=0} - \hat{S}_{syn} > 0$. We argue that, because $\Delta\hat{S}$ and $\Delta\hat{S}_{syn}$ are of the same order of magnitude, the improvement in model fit $\Delta\hat{S}$ is a result of the social model capturing an additional important feature of individuals' BMI dynamics (and not a result of overfitting), and hence, that social effects play have an impact on individuals' BMI dynamics. Because our model consistently identified social factors as important to individuals' BMI dynamics at both the individual and population levels in several independent data sets, our results represent a significant contribution to ongoing research into whether the obesity epidemic should be considered a social contagion process. These results have significant implications both in the context of the obesity epidemic (e.g. in developing successful intervention strategies) and in the context of other potential social contagion processes (e.g. by establishing a modelling framework in which social contagion processes can be studied).

Chapter 6

Conclusion

In this thesis we investigated the implications of social influence on decision making processes in two different, but related, contexts: collective action problems, and conformity problems. In Chapters 2-3 we studied social influence on a decision making process in a political context. In this context, individuals are given a choice of whether or not to join a political revolution in a dictatorial regime that employs censorship and repression. This represents a collective action problem for individuals: individuals acting unilaterally can easily be punished by the regime, but sufficiently many individuals acting collectively cannot be retaliated against due to the regime's finite resource constraint. In Chapters 4-5 we studied social influence on decision making processes in an epidemiological context, i.e. in the context of prevalence of non-communicable diseases. In this context, individuals are given a choice of whether to join one population subgroup over another (or possibly many others). This represents what we call a conformity problem, since individuals are choosing a population subgroup to which they will conform their behaviour. Specifically, in Chapter 4 we study the conformity problem facing individuals who are deciding whether to smoke or not to smoke, whereas in Chapter 5 we study the conformity problem facing individuals who are deciding how to optimize their body mass index (BMI) for optimal utility. While each chapter of this thesis makes its own contributions to the specific problems that they study, by framing our work in general terms we hope to highlight how the modelling approaches presented in this thesis are widely applicable to many related problems that can be classified under the umbrellas of collective action problems or conformity problems. We now provide a brief summary of the work presented in this thesis and conclude with some closing remarks.

Chapter 2 begins by defining a simple one dimensional ordinary differential equation model whose defining features are the visibility and policing terms which control the growth

and decay of the size of the revolution, respectively. Intuitively, the policing term models the finite resource constraint of the regime, i.e. the regime does not have the capacity to retaliate against or punish large numbers of protesters, whereas the visibility term models the collective action problem faced by individuals choosing to join the revolution, i.e. individuals are willing to join only once the revolution has grown large enough that individuals no longer fear government retribution. This simple model admits a complete mathematical analysis and a clear interpretation of the model in terms of the stability of the regime (stable police state, meta-stable police state, unstable police state, and failed state). We investigate how the model can be applied to the Arab Spring revolutions in Tunisia and Egypt, noting that our model provides possible answers to questions on causes and timing of the Arab Spring revolutions, and the role of the Internet and new media therein. Furthermore, we note that our model is broadly relevant, since its classification of regime stability can be applied (at least conceptually) to current political situations in many different countries, for example we specifically consider the cases of Iran, China and Somalia.

Chapter 3 builds on the model developed in Chapter 2 by specifying a linear threshold agent-based model (ABM) for the spread of a political revolution in a dictatorial regime that is consistent with the step visibility function (SVF) model of Chapter 2. Using the relationship between the ABM and the SVF model as a template we developed two models of moderate complexity, i.e. the binomial visibility function (BVF) and empirical visibility function (EVF) models. We show that these new models are of low computational complexity and admit a simple analysis. Moreover, because these two new models take the actual structure of empirical networks into account better than the SVF, we are able to show that they are an improvement over the SVF model in approximating the aggregate behaivour of the ABM. We note that the analysis of the SVF model, as well as the BVF and EVF models[1], are useful in guiding our computational analysis of the ABM. For example, in our analysis we extend the concept of the basic reproduction number R_0 from epidemiology to the linear threshold ABM and show how it is related to the initial slope of the BVF and EVF. Finally, in small scale numerical simulations using online and offline proxy network data we provide initial support to the hypothesis that the adoption of online social media may facilitate the spread of political revolutions by effectively changing the connectivity structure of the population in a way that makes linear threshold spreading more effective.

In Chapter 4 we develop a compartmental model for the dynamics of smoking prevalence. Using estimates for smoking prevalence from a number of developed countries, which we compute for a period of roughly one century, we show that our model matches real-world smoking prevalence data well. We are further able to show that the predictions of

[1]These analyses are provided in Chapter 2, i.e. in Sections 2.3.1 and 2.4.2.

our model (calibrated using estimated smoking prevalence data) are consistent with independently collected data, i.e. with Hofstede's individualism index IDV. In particular, we are able to confirm our model's prediction that the level of individualism or collectivism of a society significantly affects the temporal dynamics of smoking prevalence, i.e. we confirm that more collectivistic societies tend to have slower rates of smoking adoption and cessation and that more individualistic societies tend to have faster rates of smoking adoption and cessation. The significance of this effect is illustrated by considering the counterfactual scenario of how the smoking prevalence might have evolved in the United States had the United States been approximately 2% less individualistic, i.e. if the United States had had the same IDV as the United Kingdom.

Finally, in Chapter 5 we develop an agent-based model (ABM) for individuals' BMI dynamics with implications at both the individual and populations levels. We also develop a new BMI data set more abundant than any previously reported, consisting of BMI measurements that are calculated from anonymized records of patients of the Northwestern Medicine system. This results in a data set consisting of over 300,000 distinct individuals whose BMI can be computed in consecutive years (over 1,000,000 year-over-year change in BMI data points in total). We test the main predictions of our model using these data, as well as BMI data from National Health and Nutrition Examination Survey (NHANES) and Behavioral Risk Factor Surveillance System (BRFSS) surveys. At the population level our model provides a mechanistic explanation for the right-skewness observed in empirical BMI distributions, and fits empirical BMI distributions better than either the log-normal or skew-normal distributions. At the individual level, our model is able to reproduce the mean and standard deviation of year-over-year change in individuals' BMI. At both the individual and population levels we provide evidence that social factors play a role in the dynamics of individuals' BMI.

Given the inherent complexity of human beings and their interactions with one another, it can be tempting to develop extremely detailed and complex mathematical models in order to study social phenomena. However, we note that it is often the case that the more complex a model becomes the less useful it is. Indeed, models that are overly complicated are often difficult to interpret and analyze, and are prone to overfitting and spurious results (this is especially the case when the data are sparse and/or noisy). Throughout this thesis we have advocated for the development of the simplest possible mathematical model that is still capable of capturing the essence of the research problem in question. Moreover, we emphasize that the results of this thesis have demonstrated that these simple models can have significant explanatory power. Indeed, in Chapters 2 and 4 we are able to achieve the modelling objectives using simple compartmental models only. While more complicated agent-based models are necessary to model the phenomena considered in Chapters 3 and

5, we emphasize that in both cases the agent-based models are nevertheless kept as simple as possible. Moreover, we also note that in both of Chapters 3 and 5 we provide additional validation for our ABMs by accompanying them with simple population-level models. In the case of Chapter 3 the ABM is accompanied by the simple compartmental model of Chapter 2, which we use to guide the computational analysis of the ABM. In the case of Chapter 5 the ABM for the dynamics of individuals' BMI is accompanied by a partial differential equation for the distribution of BMIs in the general population, which we use to further validate our mathematical model. Specifically, we argue that our ability to reproduce both individual and population level data in Chapter 5 provides much stronger evidence of our model's validity than if we were able to reproduce individual level data only. In closing, we conclude this thesis with the following general remark: whereas advances in scientific computing will continue to improve our ability to develop and simulate complex (e.g. individual-level, or agent-based) mathematical models, simple (e.g. population-level) mathematical models will continue to play a critical role in applied mathematics .

References

[1] D.M. Abrams and S.H. Strogatz. Linguistics: Modelling the dynamics of language death. *Nature*, 424:900, 2003.

[2] D.M. Abrams, H.A. Yaple, and R.J. Wiener. Dynamics of Social Group Competition: Modeling the Decline of Religious Affiliation. *Physical Review Letters*, 107:088701, 2011.

[3] J. Adams and M. White. When the population approach to prevention puts the health of individuals at risk. *International Journal of Epidemiology*, 34(1):40–43, 2004.

[4] H. Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723, 1974.

[5] J.B. Alterman. The Revolution Will Not Be Tweeted. *The Washington Quarterly*, 34(4):103–116, 2011.

[6] G.M. Ames, D.B. George, C.P. Hampson, A.R. Kanarek, C.D. McBee, D.R. Lockwood, J.D. Achter, and C.T. Webb. Using network properties to predict disease dynamics on human contact networks. *Proceedings of the Royal Society of London B*, 278(1724):3544–3550, 2011.

[7] L. Anderson. Demystifying the Arab Spring: Parsing the Differences Between Tunisia, Egypt, and Libya. *Foreign Affairs*, 90(3):2–7, 2011.

[8] T. Andreyeva, R. Sturm, and S. Ringel. Moderate and Severe Obesity Have Large Differences in Health Care Costs. *Obesity Research*, 12(12):1936–1943, 2004.

[9] M.P. Atkinson, A. Gutfraind, and M. Kress. When do armed revolts succeed: lessons from Lanchester theory. *Journal of the Operational Research Society*, 63:1363–1373, 2011.

[10] A. Avery, J. Chase, L. Johansson, S. Litvak, D. Montero, and M. Wydra. America's Changing Attitudes toward Homosexuality, Civil Unions, and Same-Gender Marriage: 1977-2004. *Social Work*, 52(1):71–79, 2007.

[11] D.B. Bahr, R.C. Browning, H.R. Wyatt, and J.O. Hill. Exploiting Social Networks to Mitigate the Obesity Epidemic. *Obesity*, 17(4):723–728, 2009.

[12] S. Bansal, B.T. Grenfell, and L.A. Meyers. When individual behaviour matters: homogeneous and network models in epidemiology. *Journal of The Royal Society Interface*, 4(16):879–891, 2007.

[13] E. Barbeau, N. Krieger, and M.-J. Soobader. Working Class Matters: Socioeconomics Disadvantage, Race/Ethnicity, Gender, and Smoking in NHIS 2000. *American Journal of Public Health*, 94(2):269–278, 2004.

[14] C.T. Bauch. A versatile ODE approximation to a network model for the spread of sexually transmitted diseases. *Journal of Mathematical Biology*, 45(5):375–395, 2002.

[15] C. Beckett. After Tunisia and Egypt: towards a new typology of media and networked political change. Blog, Februrary 2011. Accessed August 27, 2012. URL: http://blogs.lse.ac.uk/polis/2011/02/11/after-tunisia-and-egypt-towards-a-new-typology-of-media-and-networked-political-change/.

[16] J. Benavides, B.C.P. Demianyk, S.N. Mukhi, M. Laskowski, M. Friesen, and R.D. McLeod. Smartphone Technologies for Social Network Data Generation and Infectious Disease Modeling. *Journal of Medical and Biological Engineering*, 32(4):235–244, 2012.

[17] A. Berrington de Gonzalez, P. Hartge, J.R. Cerhan, A.J. Flint, L. Hannan, R.J. MacInnis, S.C. Moore, G.S. Tobias, H. Anton-Culver, L.B. Freeman, W.L. Beeson, S.L. Clipp, D.R. English, A.R. Folsom, D.M. Freedman, G. Giles, N. Hakansson, K.D. Henderson, J. Hoffman-Bolton, J.A. Hoppin, K.L. Koenig, I-Min Lee, M.S. Linet, Y. Park, G. Pocobelli, A. Schatzkin, H.D. S., E. Weiderpass, B.J. Willcox, A. Wolk, A. Zeleniuch-Jacquotte, W.C. Willett, and M.J. Thun. Body-Mass Index and Mortality among 1.46 Million White Adults. *New England Journal of Medicine*, 363(23):2211–2219, 2010.

[18] Patrick Billingsley. *Convergence of probability measures*. John Wiley & Sons, 2013.

[19] G. Blight, S. Pulham, and P. Torpey. Arab spring: an interactive timeline of Middle East protests. The Guardian, January 2012. Accessed Sept 14,

2012. URL: http://www.theguardian.com/world/interactive/2011/mar/22/middle-east-protest-interactive-timeline.

[20] E. Bonabeau. Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences*, 99:7280–7287, 2002.

[21] D.M. Boyd and N.B. Ellison. Social Network Sites: Definition, History, and Scholarship. *Journal of Computer-Mediated Communication*, 13:210–230, 2008.

[22] D. Braha. Global Civil Unrest: Contagion, Self-Organization, and Prediction. *PLoS ONE*, 7(10):e48596, 2012.

[23] A. Burns and B. Eltham. Twitter Free Iran: An Evaluation of Twitter's Role in Public Diplomacy and Information Operations in Iran's 2009 Election Crisis. *Record of the Communications Policy & Research Forum 2009*, pages 298–310, 2009.

[24] B. Butland, S. Jebb, P. Kopelman, K. McPherson, S. Thomas, J. Mardell, and V. Parry. Tackling Obesities: Future Choices - Project report. Technical report, UK Government Office for Science, Foresight Programme, 2007.

[25] M. Cabanac and P. Frankham. Evidence that transient nicotine lowers the body weight set point. *Physiology & Behavior*, 76(4-5):539–542, 2002.

[26] A. Calavó-Armengol and M.O. Jackson. Peer Pressure. *Journal of the European Economic Association*, 8(1):62–89, 2010.

[27] E. Campbell and M. Salathé. Complex social contagion makes networks more vulnerable to disease outbreaks. *Scientific Reports*, 3(1905), 2013.

[28] H. Carlsson and E. van Damme. Equilibrium Selection in Stag Hunt Games. In K.G. Binmore, A.P. Kirman, and P. Tani, editors, *Frontiers of Game Theory*, chapter 12, pages 237–254. MIT Press, 1993.

[29] R.Y. Cavana and M. Tobias. Integrative System Dynamics: Analysis of Policy Options for Tobacco Control in New Zealand. *Systems Research and Behavioral Science*, 25(5):675–694, 2008.

[30] J. Cawley and C. Meyerhoefer. The medical care costs of obesity: An instrumental variables approach. *Journal of Health Economics*, 31(1):219–230, 2012.

[31] Centers for Disease Control and Prevention (CDC). Behavioural Risk Factor Surveillance System Survey Data. Atlanta, GA: U.S. Depaartment of Health and Human Services, Centers for Disease Control and Prevention. URL: http://www.cdc.gov/brfss/annual_data/annual_data.htm, 2015.

[32] Centers for Disease Control and Prevention (CDC). National Center for Health Statistics (NCHS). National Health and Nutrition Survey Data. Hyattsville, MD: U.S. Depaartment of Health and Human Services, Centers for Disease Control and Prevention. URL: http://www.cdc.gov/nchs/nhanes/nhanes_questionnaires.htm, 2015.

[33] D. Centola. An Experimental Study of Homophily in the Adoption of Health Behavior. *Science*, 334(6060):1269–1272, 2011.

[34] D. Centola, V.M. Eguiluz, and M.W. Macy. Cascade dynamics of complex propagation. *Physica A*, 374(1):449–456, 2007.

[35] D. Centola and M.W. Macy. Complex Contagions and the Weakness of Long Ties. *American Journal of Sociology*, 113(3):702–734, 2007.

[36] D. Centola, R. Willer, and M.W. Macy. The Emperor's Dilemma: A Computational Model of Self-Enforcing Norms. *American Journal of Sociology*, 110(4):1009–1040, 2005.

[37] O. Chapa, M.D. Hernandez, Y.J. Wang, and C. Skalski. Do Individualists Complain More than Collectivists? A Four-Country Analysis on Consumer Complaint Behavior. *Journal of International Consumer Marketing*, 26(5):373–390, 2014.

[38] J.Y. Chiao and K.D. Blizinsky. Culture-gene coevolution of individualism-collectivism and the serotonin transporter gene. *Proceedings of the Royal Society B*, 277(1681):529–537, 2010.

[39] N.A. Christakis and J.H. Fowler. The Spread of Obesity in a Large Social Network over 32 Years. *The New England Journal of Medicine*, 357(4):370–379, 2007.

[40] N.A. Christakis and J.H. Fowler. The Collective Dynamics of Smoking in a Large Social Network. *The New England Journal of Medicine*, 358(21):2249–2258, 2008.

[41] E. Cohen-Cole and J.M. Fletcher. Is obesity contagious? Social networks vs. environmental factors in the obesity epidemic. *Journal of Health Economics*, 27(5):1382–1387, 2008.

[42] C.D. Collier. *Cremation as an Emerging Cultural System*. PhD thesis, University of Georgia, 2007.

[43] P. Collier and A. Hoeffler. Greed and grievance in civil war. *Oxford Economic Papers*, 56(4):563–595, 2004.

[44] S. Currarini, M.O. Jackson, and P. Pin. An Economic Model of Friendship: Homophily, Minorities, and Segregation. *Econometrica*, 77(4):1003–1045, 2009.

[45] D.M. Cutler and E.L. Glaeser. Why Do Europeans Smoke More than Americans? In D.A. Wise, editor, *Developments in the Economics of Aging*. University of Chicago Press, Chicago, 2009.

[46] G. Demirel, R. Prizak, P.N. Reddy, and T. Gross. Cyclic dominance in adaptive networks. *The European Physical Journal B*, 84(4):541–548, 2011.

[47] A. Dunn. Unplugging a Nation: State Media Strategy During Egypt's January 25 Uprising. *Fletcher Forum of World Affairs*, 35(2):15–24, 2011.

[48] K. Ejima, K. Aihara, and H. Nishiura. Modeling the obesity epidemic: social contagion and its implications for control. *Theoretical Biology and Medical Modelling*, 10(1):1–13, 2013.

[49] G. Eknoyan. Adolphe Quetelet (1796-1874)–the average man and indices of obesity. *Nephrology Dialysis Transplantation*, 23(1):47–51, 2008.

[50] N. Eltantawy and J.B. Wiest. The Arab Spring— Social Media in the Egyptian Revolution: Reconsidering Resource Mobilization Theory. *International Journal of Communication*, 5:18, 2011.

[51] J.M. Epstein. Why Model? *Journal of Artificial Societies and Social Simulation*, 11(4):12, 2008.

[52] N. Etemadi. An elementary proof of the strong law of large numbers. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 55(1):119–122, 1981.

[53] B. Feldman. Somalia: Amidst the Rubble, a Vibrant Telecommunications Infrastructure. *Review of African Political Economy*, 34(113):565–572, 2007.

[54] L. Feng, P. Monterola, and Y. Hu. The simplified self-consistent probabilities method for percolation and its application to interdependent networks. *New Journal of Physics*, 17(6):063025, 2015.

[55] N.M. Ferguson and G.P. Garnett. More Realistic Models of Sexually Transmitted Disease Transmission Dynamics: Sexual Partnership Networks, Pair Models and Moment Closure. *Sexually Transmitted Diseases*, 27(10):600–609, 2000.

[56] A.E. Field, E.H. Coakley, A. Must, J.L. Spadano, N. Laird, W.H. Diets, Rimm E., and G.A. Colditz. Impact of Overweight on the Risk of Developing Common Chronic Diseases During a 10-year Period. *Archives of Internal Medicine*, 161(13):1581–1586, 2001.

[57] C.L. Fincher, R. Thornhill, D.R. Murray, and M. Schaller. Pathogen prevalence predicts human cross-cultural variability in individualism/collectivism. *Proceedings of the Royal Society B*, 275(1640):1279–1285, 2008.

[58] E.A. Finkelstein, J.G. Trogdon, J.W. Cohen, and W. Dietz. Annual Medical Spending Attributable To Obesity: Payer-And Service-Specific Estimates. *Health Affairs*, 28(5):822–831, 2009.

[59] M.A. Finkelstein. Individualism/collectivism: Implications for the volunteer process. *Social Behaviour and Personality: an international journal*, 38(4):445–452, 2010.

[60] M.M. Finucane, G.A. Stevens, M. Cowan, G. Danaei, J.K. Lin, C.J. Paciorek, G.M. Singh, H.R. Gutierrez, Y. Lu, A.N. Bahalim, F. Farzadfar, L.M. Riley, and M. Ezzati. National, regional and global trends in body mass index since 1980: systematic analysis of health examination surveys and epidemiological studies with 960 country-years and 9.1 million participants. *Lancet*, 377(9765):557–567, 2011.

[61] K.M. Flegal. Trends in Body Weight and Overweight in the U.S. Population. *Nutrition Reviews*, 54(4):S97–S100, 1996.

[62] K.M. Flegal and R.P. Troiano. Changes in the distribution of body mass index of adults and children in the US population. *International Journal of Obesity*, 24(7):807–818, 2000.

[63] B. Forey, J. Hamling, J. Hamling, A. Thorton, and P. Lee. *International Smoking Statistics: A collection of worldwide historical data*. P N Lee Statistics & Computing Ltd, 17 Cedar Road, Sutton, SM2 5DA, UK, web edition, 2013. Accessed October 7, 2013.

[64] B. Forey, J. Hamling, P. Lee, and N. Wald. *International Smoking Statistics: A collection of historical data from 30 economically developed countries*. Oxford University Press, Great Clarendon Street, Oxford, OX2 6DP, UK, 2 edition, 2002.

[65] D. Frankenfield, L. Roth-Yousey, and C. Compher. Comparison of Predictive Equations for Resting Metabolic Rate in Healthy Nonobese and Obese Adults: A Systematic Review. *Journal of the American Dietetic Association*, 105(5):775 – 789, 2005.

[66] S. Frederick, G. Lowenstein, and T. O'Donoghue. Time discounting and time preference: A critical review. *Journal of Economic Literature*, 40(2):351–401, 2002.

[67] S. Gallus, A. Schiaffino, C. La Vecchia, J. Townsend, and E. Fernandes. Price and cigarette consumption in Europe. *Tobacco Control*, 2(114-119), 15.

[68] Crispin Gardiner. *Stochastic methods*. Springer Berlin, 2009.

[69] J.P. Gleeson and D.J. Calahane. Seed size strongly affects cascades on random networks. *Physical Review E*, 75:056103, 2007.

[70] David F. Gleich. *Models and Algorithms for PageRank Sensitivity*. PhD thesis, Stanford University, September 2009. Chapter 7 on MatlabBGL.

[71] J.A. Goldstone. Towards a Fourth Generation of Revolutionary Theory. *Annual Review of Political Science*, 4:139–187, 2001.

[72] S. González-Bailón, J. Borge-Holthoefer, A. Rivero, and Y. Moreno. The Dynamics of Protest Recruitment through an Online Network. *Scientific Reports*, 1:197, 2011.

[73] G. González-Parra, L. Acedo, R-J. Villanueva Micó, and A.J. Arenas. Modeling the social obesity epidemic with stochastic networks. *Physica A*, 389(17):3692–3701, 2010.

[74] J. Gorstein and R.N. Grosse. The Indirect Costs of Obesity to Society. *Pharmacoeconomics*, 5(1):58–61, 1994.

[75] M. Granovetter. Threshold Models of Collective Behavior. *American Journal of Sociology*, 83(6):1420–1443, 1978.

[76] V. Griskevicius, J.M. Tybur, A.W. Delton, and T.E. Robertson. The Influence of Mortality and Socioeconomic Status on Risk and Delayed Rewards: A Life History Theory Approach. *Journal of Personality and Social Psychology*, 100(6):1015–1026, 2011.

[77] T. Gross, C.J. Dommar D'Lima, and B. Blasius. Epidemic Dynamics on an Adaptive Network. *Physical Review Letters*, 96:208701, 2006.

[78] A. Gürsoy and M. Atun. *Euro-Par 2000 Parallel Processing: 6th International Euro-Par Conference Munich, Germany, August 29 – September 1, 2000 Proceedings*, chapter Neighbourhood Preserving Load Balancing: A Self-Organizing Approach, pages 234–241. Springer Berlin Heidelberg, Berlin, Heidelberg, 2000.

[79] C. Hadjichrysanthou, M. Broom, and I.Z. Kiss. Approximating evolutionary dynamics on networks using a Neighbourhood Configuration model. *Journal of Theoretical Biology*, 312:13–21, 2012.

[80] K.D. Hall and P.N. Jordan. Modeling weight-loss maintenance to help prevent body weight regain. *The American Journal of Clinical Nutrition*, 88(6):1495–1503, 2008.

[81] K.D. Hall, G. Sacks, D. Chandramohan, C.C. Chow, Y.C. Wang, S.L. Gortmaker, and B.A. Swinburn. Quantification of the effect of energy imbalance on bodyweight. *The Lancet*, 378(9793):826 – 837, 2011.

[82] C.S. Hendrix. Measuring state capacity: Theoretical and empirical implications for the study of civil conflict. *Journal of Peace Research*, 47(3):273–285, 2010.

[83] J. Henrich and R. McElreath. The evolution of cultural evolution. *Evolutionary Anthropology*, 12(3):123–135, 2003.

[84] H.W. Hethcote. The Mathematics of Infectious Diseases. *SIAM Review*, 42(4):599–653, 2000.

[85] A. Hofheinz. The Internet in the Arab World: Playground for Political Liberalization. *International Politics and Society*, 3(1):78–96, 2005.

[86] G. Hofstede, G.J. Hofstede, and M. Minkov. *Cultures and Organizations: Software of the Mind: Intercultural Cooperation and its Importance for Survival*, chapter I, We, and They, pages 89–134. McGraw Hill, New York, 3 edition, 2010.

[87] J. Holliday. The Struggle for Syria in 2011: An Operational and Regional Analysis. Technical report, Institute for the Study of War, December 2011.

[88] W Hosking, R. Borland, H.H. Yong, G. Fong, M. Zanna, F. Laux, J. Thrasher, W.B. Lee, B. Sirirassamee, and M. Omar. The effects of smoking norms and attitudes on quitting intentions in Malaysia, Thailand and four Western nations: a cross-cultural comparison. *Psychology & Health*, 24(1):95–107, 2009.

[89] P.N. Howard, A. Duffy, D. Freelon, M. Hussain, W. Mari, and M. Mazaid. Opening Closed Regimes: What Was the Role of Social Media During the Arab Spring? *Project on Information Technology & Political Islam*, Research Memo. Seattle, University of Washington, 2011. Accessed May 7, 2012.

[90] Y. Hu, S. Havlin, and H.A. Makse. Conditions for Viral Influence Spreading through Multiplex Correlated Social Networks. *Physical Review X*, 4(2):021031, 2014.

[91] M. Huisman, A.E. Kunst, and J.P. Mackenback. Educational inequalities in smoking among men and women aged 16 years and older in 11 European countries. *Tobacco Control*, 14(2):106–113, 2005.

[92] M. Husain and R. Pollack. How facebook changed the world: The arab spring (parts 1 and 2). BBC Documentary, September 2011.

[93] P.T. James, R. Leach, E. Kalamara, and M. Shayeghi. The Worldwide Obesity Epidemic. *Obesity Research*, 9(S11):228S–233S, 2001.

[94] I.L. Janis. The Role of Social Support in Adherence to Stressful Decisions. *American Psychologist*, 38(2):143–160, 1983.

[95] P. Jha. Avoidable global cancer deaths and total deaths from smoking. *Nature Reviews Cancer*, 9:655–664, 2009.

[96] L. Joossens and M. Raw. The Tobacco Control Scale: a new scale to measure country activity. *Tobacco Control*, 15(3):247–253, 2006.

[97] R.E. Keesey and M.D. Hirvonen. Body Weight Set-Points: Determination and Adjustment. *The Journal of Nutrition*, 127(9):1875S–1883S, 1997.

[98] D. Kempe, J. Kleinberg, and E. Tardos. Maximizing the Spread of Influence through a Social Network. *Theory of Computing*, 11(4):105–147, 2003.

[99] W.O. Kermack and A.G. McKendrick. A Contribution to the Mathematical Theory of Epidemics. *Proceedings of the Royal Society of London A*, 115(772):700–721, 1927.

[100] S. Khamis and K. Vaughn. Cyberactivism in the Egyptian Revolution: How Civic Engagement and Citizen Journalism Tilted the Balance. *Arab Media and Society*, 14(3):1–25, 2011. Accessed May 17, 2012.

[101] A. Khelil, C. Becker, J. Tian, and K. Rothermel. An Epidemic Model for Information Diffusion in MANETs. In *Proceedings of the 5th ACM international workshop on Modeling analysis and simulation of wireless and mobile systems*, pages 1–7, 2002.

[102] H.H. Khondker. Role of the New Media in the Arab Spring. *Globalizations*, 8(5):675–679, 2011.

[103] S. Khorshid. Egypt's counterrevolution. *The New York Times*, 2013. Accessed Dec. 30, 2013.

[104] I.Z. Kiss, K.M. Green, and R.R. Kao. The effect of contact heterogeneity and multiple routes of transmission on final epidemic size. *Mathematical Biosciences*, 203(1):124–136, 2006.

[105] M. Kress. Modelling Armed Conflicts. *Science*, 336(6083):865–869, 2012.

[106] C. Kuhlman, V. Kumar, M. Marathe, S. Swarup, G. Tuli, S. Ravi, and D. Rosenkrantz. A bi-threshold model of complex contagion and its application to the spread of smoking behaviour. In *Proceedings of the workshop on social network mining and analysis (SNA-KDD 2011)*, 2011.

[107] T. Kuran. Now out of Never: The Element of Surprise in the East European Revolution of 1989. *World Politics*, 44(1):7–48, 1991.

[108] J.C. Lang, D.M. Abrams, and H. De Sterck. The influence of societal individualism on a century of tobacco use: modelling the prevalence of smoking. *BMC Public Health*, 15(1):1–13, 2015.

[109] J.C. Lang and H. De Sterck. The Arab Spring: A simple compartmental model for the dynamics of a revolution. *Mathematical Social Sciences*, 69:12–21, 2014.

[110] J.C. Lang and H. De Sterck. A hierarchy of linear threshold models for the spread of political revolutions on social networks. *Journal of Complex Networks*, 2016. DOI: 10.1093/comnet/cnv030.

[111] D.T. Levy, J.E. Bauer, and H. Lee. Simulation Modeling and Tobacco Control: Creating More Robust Public Health Policies. *American Journal of Public Health*, 96(3):494–498, 2006.

[112] S.S. Lim, T. Vos, A.D. Flaxman, G. Danaei, K. Shibuya, and et. al. A comparative risk assessment of burden of disease and injury attributable to 67 risk factors and

risk factor clusters in 21 regions, 1990-2010: a systematic analysis for the Global Burden of Disease Study 2010. *Lancet*, 380(9859):2224–2260, 2012.

[113] P. Lind, L. da Silva, J. Andrade, and H. Herrmann. Spreading gossip in social networks. *Physical Review E*, 76:036117, 2007.

[114] J. Lindquist, J. Ma, P. van den Driessche, and F.H. Willeboordse. Effective degree network disease models. *Journal of Mathematical Biology*, 62(2):143–164, 2011.

[115] A.D. Lopez, N.E. Collishaw, and T. Piha. A descriptive model of the cigarette epidemic in developed countries. *Tobacco Control*, 3(3):242–247, 1994.

[116] G. Lotan, E. Graeff, M. Ananny, D. Gaffney, I. Pearce, and D. Boyd. The Revolutions Were Tweeted: Information Flows during the 2011 Tunisian and Egyptian Revolutions. *International Journal of Communication*, 5:1375–1405, 2011.

[117] C.M. Macal and M.J. North. Tutorial on agent-based modelling and simulation. *Journal of Simulation*, 4(3):151–162, 2010.

[118] J.P. Mackenbach. Health Inequalities: Europe in Profile. An Independent, Expert Report Commissioned by the UK Presidency of the EU. Technical report, Rotterdam, Germany: Erasmus University Medical Center, 2006.

[119] C.D. Mathers and D. Loncar. Projections of global mortality and burden of disease from 2002 to 2030. *PLoS Medicine*, 3(11):e442, 2006.

[120] J. McAuley and J. Leskovec. Learning to discover social circles in ego networks. In *Advances in Neural Information Processing Systems 25*, pages 548–556, 2012.

[121] M. McPherson and J.M. Cook L. Smith-Lovin. Birds of a Feather: Homophily in Social Networks. *Annual Review of Sociology*, 27:415–444, 2001.

[122] K. Menkhaus. Governance without Government in Somalia: Spoilers, State Building, and the Politics of Coping. *International Security*, 31(3):74–106, 2007.

[123] K. Midthjell, O. Krüger, J. Holmen, A. Tverdal, T. Claudi, A. Bjørndal, and P. Magnus. Rapid changes in the prevalence of obesity and known diabetes in an adult Norwegian population: The Nord-Trøndelag Health Surveys: 1984-1986 and 1995-1997. *Diabetes Care*, 22(11):1813–1820, 1999.

[124] J.C. Miller. Complex contagions and hybrid phase transitions. *Journal of Complex Networks*, 2015.

[125] A.H. Mokdad, M.K. Serdula, W.H. Dietz, B.A. Bowman, J.S. Marks, and J.P. Koplan. The Spread of the Obesity Epidemic in the United States, 1991-1998. *Journal of the American Medical Association*, 282(16):1519–1522, 1999.

[126] Y. Moreno, M. Nekovee, and A.F. Pacheco. Dynamics of rumour spreading in complex networks. *Physical Review E*, 69:066130, 2004.

[127] J.M. Moskowitz, Z. Lin, and E.S. Hudes. The impact of workplace smoking ordinances in California on smoking cessation. *American Journal of Public Health*, 90(5):757–761, 2000.

[128] M. Nekovee, Y. Moreno, G. Bianconi, and M. Marsili. Theory of rumour spreading in complex social networks. *Physica A*, 374(1):457–470, 2007.

[129] OECD. OECD.Stat, Demography and Population: Population Statistics: Population and Vital Statistics. *OECD.Stat database*, 2013. Accessed January 20, 2014.

[130] OECD. OECD.Stat, Health: Non-Medical Determinants of Health: Tobacco consumption. *OECD.Stat database*, 2013. Accessed October 24, 2013.

[131] World Health Organization. Global target 7: Halt the rise in diabetes and obesity. In S. Mendis, editor, *Global status report on noncommunicable diseases 2014*, pages 79–93. World Health Organization, 2014.

[132] A. Palloni. Diffusion in sociological analysis. In *Diffusion Processes and Fertility Transition: Selected Perspectives*, pages 66–114. National Academies Press, Washington, DC, 2001.

[133] F. Pampel. *International Differences in Mortality at Older Ages: Dimensions and Sources*, chapter Divergent Patterns of Smoking Across High-Income Nations, pages 132–163. National Academies Press, Washington D.C., 2010.

[134] F.C. Pampel. Cigarette Siffusion and Sex Differences in Smoking. *Journal of Health and Social Behaviour*, 42(4):388–404, 2001.

[135] E Paoletti. Libya: Roots of a Civil Conflict. *Mediterranean Politics*, 16(2):313–319, July 2011.

[136] H.V.D. Parunak, R. Savit, and R.L. Riolo. *Multi-agent systems and agent-based simulation*, chapter Agent-based modeling vs. equation-based modeling: A case study and users' guide., pages 10–25. Springer Berlin Heidelberg, 1998.

[137] R. Pastor-Satorras and A. Vespignani. Epidemic dynamics in finite size scale-free networks. *Physical Review E*, 65:035108(R), 2002.

[138] W. Pearlman. Emotions and the Microfoundations of the Arab Uprisings. *Perspectives on Politics*, 11(2):387–409, 2013.

[139] S. Pei, L. Muchnik, J.S. Andrade, Z. Zheng, and H.A. Makse. Searching for superspreaders of information in real-world social media. *Scientific Reports*, 4:5547, 2014.

[140] A.D. Penman. The changing distribution of body mass index in the Mississippi adult population, 1990-2003: implications for public health policy. Master's thesis, University of Mississippi Medical Center, 2005.

[141] A.D. Penman and W.D. Johnson. The Changing Shape of the Body Mass Index Distribution Curve in the Population: Implications for Public Health Policy to Reduce the Prevalence of Adult Obesity. *Preventing Chronic Disease*, 3(3):A74, 2006.

[142] J.U. Pentecostes. Individualism vs Collectivism: Implications for Health Promotion. *Phillippine Journal of Psychology*, 32(2):397–402, 1999.

[143] T. Pischon, H. Boeing, K. Hoffmann, M. Bergmann, M.B. Schulze, K. Overvad, Y.T. van der Schouw, E. Spencer, K.G.M. Moons, A. Tjønneland, J. Halkjaer, M.K. Jensen, J. Stegger, F. Clavel-Chapelon, M.-C. Boutron-Ruault, V. Chajes, J. Linseisen, R. Kaaks, A. Trichopoulou, D. Trichopoulos, C. Bamia, S. Sieri, D. Palli, R. Tumino, P. Vineis, S. Panico, P.H.M. Peeters, A.M. May, H.B. Bueno-de Mesquita, F.J.B. van Duijnhoven, G. Hallmans, L. Weinehall, J. Manjer, B. Hedblad, E. Lund, A. Agudo, L. Arriola, A. Barricarte, C. Navarro, C. Martinez, J.R. Quirós, T. Key, S. Bingham, K.T. Khaw, P. Boffetta, M. Jenab, P. Ferrari, and E. Riboli. General and Abdominal Adiposity and Risk of Death in Europe. *New England Journal of Medicine*, 359(20):2105–2120, 2008. PMID: 19005195.

[144] J. Pollock. Streetbook: How Egyptian and Tunisian youth hacked the Arab Spring. *Technology Review*, 2011. Accessed Aug. 24, 2012.

[145] J. Poncela-Casanovas, B. Spring, D. McClary, A.C. Moller, R. Mukogo, C.A. Pellegrini, M.J. Coons, M. Davidson, S. Mukherjee, and L.A.N. Amaral. Social embeddedness in an online weight management programme is linked to greater weight loss. *Journal of The Royal Society Interface*, 12:20140686, 2015.

[146] U.S. Census Bureau Population Estimates Program, Population Division. Historical National Population Estimates: July 1, 1900 to July 1, 1999. Online, April 2000.

[147] U.S. Census Bureau Population Estimates Program, Population Division. Monthly Intercensal Resident Population Estimates for the United States: April 1, 2000 to July 1, 2010. Online, September 2011.

[148] D. Power, T. Schoenherr, and D. Samson. The cultural characteristics of individualism/collectivism: A comparative study of implications for investment in operations between emerging Asian and industrialized Western countries. *Journal of Operations Management*, 28(3):206–222, 2010.

[149] H. Rahmandad and J. Sterman. Heterogeneity and Network Structure in the Dynamics of Diffusion: Comparing Agent-Based and Differential Equation Models. *Management Science*, 54(5):998–1014, 2008.

[150] M. Ramos, J. Shao, S.D.S. Reis, C. Anteneodo, J.S. Andrade, S. Havlin, and H.A. Makse. How does public opinion become extreme? *Scientific Reports*, 5:10032, 2015.

[151] L.F. Richardson. *Arms and insecurity: A mathematical study of the causes and origins of war*. Boxwood Press, Pittsburgh, Pennsylvania, 1960.

[152] R. Rifai. Timeline: Tunisia's uprising. Al Jazeera, January 2011. Accessed Sept 14, 2012.

[153] D.M. Romero, B. Meeder, and J. Kleinberg. Differences in the Mechanics of Information Diffusion Across Topics: Idioms, Political Hashtags, and Complex Contagion on Twitter. In *Proceedings of the 20th International Conference on World Wide Web*, WWW '11, pages 695–704, New York, NY, USA, 2011. ACM.

[154] G. Rose. Sick individuals and sick populations. *International Journal of Epidemiology*, 30(3):427–432, 2001.

[155] G. Rose and S. Day. The population mean predicts the number of deviant individuals. *BMJ*, 301(6759):1031–1034, 1990.

[156] D.C. Rowe, L. Chassin, C.C. Presson, D. Edwards, and S.J. Sherman. An "Epidemic" Model of Adolescent Cigarette Smoking. *Journal of Applied Social Psychology*, 22(4):261–285, 1992.

[157] A. Said. The Paradox of Transition to "Democracy" under Military Rule. *Social Research: An International Quarterly*, 79(2):397–434, 2012.

[158] M. Salathé, M. Kazandjieva, J.W. Lee, P. Levis, M.W. Feldman, and J.H. Jones. A high-resolution human contact network for infectious disease transmission. *Proceedings of the National Academy of Sciences*, 107(51):22020–22025, 2010.

[159] C.L. Schneider. Violence and State Repression. *Swiss Political Science Review*, 17(4):480–484, 2011.

[160] M. Shah, P.J. Hannan, and R.W. Jeffery. Secular trend in body mass index in the adult population of three communities from the upper mid-western part of the USA: the Minnesota Heart Health Program. *International Journal of Obesity*, 15(8):499–503, 1991.

[161] O. Sharomi and A. B. Gumel. Curtailing smoking dynamics: A mathematical modeling approach. *Applied Mathematics and Computation*, 195(2):475–499, 2008.

[162] C. Shirky. The Political Power of Social Media - Technology, the Public Sphere, and Political Change. *Foreign Affairs*, 90(1):28–41, 2011.

[163] J.R. Speakman, D.A. Levitsky, D.B. Allison, M.S. Bray, J.M. de Castro, D.J. Clegg, J.C. Clapham, A.G. Dulloo, L. Gruer, S. Haw, J. Hebebrand, M.M. Hetherington, S. Higgs, S.A. Jebb, Ruth J. F. Loos, S. Luckman, A. Luke, V. Mohammed-Ali, S. O'Rahilly, M. Pereira, L. Perusse, T.N. Robinson, B. Rolls, M.E. Symonds, and M.S. Westerterp-Plantenga. Set points, settling points and some alternative models: theoretical options to understand how genes and environments combine to regulate body adiposity. *Disease Models and Mechanisms*, 4(6):733–745, 2011.

[164] K. Staub, F.J. Rühli, U. Woitek, and C. Pfister. BMI distribution/social stratification in Swiss conscripts form 1875 to present. *European Journal of Clinical Nutrition*, 64:335–340, 2010.

[165] E. Stepanova. The Role of Information Communication Technologies in the "Arab Spring": Implications Beyond the Region. *PONSARS Eurasia Policy Memo*, 2011. Accessed June 1, 2012.

[166] T. Sun, M. Horn, and D. Merritt. Values and lifestyles of individuals and collectivists: a study on Chinese, Japanese, British, and US consumers. *Journal of Consumer Marketing*, 21(5):318–331, 2004.

[167] C. Swencionis and S.L. Rendell. The psychology of obesity. *Abdominal Imaging*, 37(5):733–737, 2012.

[168] B. Swinburn and G. Egger. The Runaway Weight Gain Train: Too Many Accelerators, Not Enough Brakes. *BMJ*, 329(7468):736–739, 2004.

[169] B. Swinburn, G. Sacks, S.K. Lo, K.R. Westerterp, E.C. Rush, M. Rosenbaum, A. Luke, D.A. Schoeller, J.P. DeLany, N.F. Butte, and E. Ravussin. Estimating the changes in energy flux that characterize the rise in obesity prevalence. *American Journal of Clinical Nutrition*, 89(6):1723–1728, 2009.

[170] B. Swinburn, G. Sacks, and E. Ravussin. Increased food energy supply is more than sufficient to explain the US epidemic of obesity. *American Journal of Clinical Nutrition*, 90(6):1453–1456, 2009.

[171] S. Thomas, D. Fayter, K. Misso, D. Oglivie, M. Petticrew, A. Sowden, M. Whitehead, and G. Worthy. Population tobacco control interventions and their effects on social inequalities in smoking: systematic review. *Tobacco Control*, 17(4):230–237, 2008.

[172] M. Tomasello. The ultra-social animal. *European Journal of Social Psychology*, 44(3):187–194, 2014.

[173] M. Tomasello, A.P. Melis, C. Tennie, E. Wyman, and E. Herrmann. Two Key Steps in the Evolution of Human Cooperation: The Interdependence Hypothesis. *Current Anthropology*, 53(6):673–692, 2012.

[174] H.C. Triandis, R. Bontempo, and M.J. Villareal. Individualism and Collectivism: Cross-Cultural Perspectives on Self-Ingroup Relationships. *Journal of Personality and Social Psychology*, 54(2):323–338, 1988.

[175] J.G. Trogdon, E.A. Finkelstein, T. Hylands, P.S. Dellea, and S.J. Kamal-Bahl. Indirect costs of obesity: a review of the current literature. *Obesity Reviews*, 9(5):489–500, 2008.

[176] J.G. Trogdon, J. Nonnemaker, and J. Pais. Peer effects in adolescent overweight. *Journal of Health Economics*, 27(5):1388–1399, 2008.

[177] Z. Tufekci and C. Wilson. Social Media and the Decision to Participate in Political Protest: Observations From Tahrir Square. *Journal of Communication*, 62(2):363–379, 2012.

[178] P.J. Turnbaugh, R.E. Ley, M.A. Mahowald, V. Magrini, E.R. Mardis, and J.I. Gordonl. An obesity-associated guy microbiome with increased capacity for energy harvest. *Nature*, 444:1027–1031, 2006.

[179] J.B. Unger, L.A. Rorhbach, T.B. Cruz, L. Baezconde-Garbanati, K.A. Howard, P.H. Palmer, and C.A. Johnson. Ethnic variation in peer influences on adolescent smoking. *Nicotine Tob. Res.*, 3(2):167–176, 2001.

[180] US Department of Health and Human Services. The health consequences of smoking–50 years of progress: A report of the surgeon general, 2014.

[181] US Department of Health, Education, and Welfare. Smoking and health: Report of the advisory committee to the surgeon general of the public health service, 1964.

[182] T.W. Valente, K. Fujimoto, C.P. Chou, and D. Spruijt-Metz. Adolescent Affiliations and Adiposity: A Social Network Analysis of Friendships and Obesity. *Journal of Adolescent Health*, 45(2):202–204, 2009.

[183] G.A.K. van Voorn and B.W. Kooi. Smoking epidemic eradication in a eco-epidemiological dynamical model. *Ecological Complexity*, 14:180–189, 2013.

[184] P. Vanhems, A. Barrat, C. Cattuto, J.-F. Pinton, N. Khanafer, C. Régis, B. Kim, B. Comte, and N. Voirin. Estimating Potential Infection Transmission Routes in Hospital Wards Using Wearable Proximity Sensors. *PLoS ONE*, 8(9):1–9, 2013.

[185] R.J. Villanueva, A.J. Arenas, and G. González-Parra. A Nonstandard Dynamically Consistent Numerical Scheme Applied to Obesity Dynamics. *Journal of Applied Mathematics*, 2008:640154, 2008.

[186] K.E. Warner. Cigarette smoking in the 1970's: the impact of the antismoking campaign on consumption. *Science*, 211(4483):729–731, 1981.

[187] D.J. Watts. A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences*, 99(9):5766–5771, 2002.

[188] A. Wedeman. Enemies of the State: Mass Incidents and Subversion in China. In *American Political Science Association 2009 Toronto Meeting Paper*, 2009.

[189] L.R. Yound and M. Nestle. The Contribution of Expanding Portion Sizes to the US Obesity Epidemic. *American Journal of Public Health*, 92(2):246–249, 2002.

[190] M.A. Zárate, M. Shaw, J.A. Marquez, and D. Biagas. Cultural inertia: The effects of cultural change on intergroup relations and the self-concept. *Journal of Experimental Social Psychology*, 48(3):634–645, 2012.

[191] L. Zhao, Q. Wang, J. Cheng, Y. Chen, J. Wang, and W. Huang. Rumor spreading model with consideration of forgetting mechanism: A case of online blogging LiveJournal. *Physica A*, 390(13):2619–2625, 2011.

[192] X. Zhuo, B. Wellman, and J. Yu. Egypt: The First Internet Revolt? *Peace Magazine*, 27:6–10, 2011.

[193] D.M. Zulman, S. Vijan, G.S. Omenn, and R.A. Hayward. The Relative Merits of Population-Based and Targeted Prevention Strategies. *The Milbank Quarterly*, 86(4):557–580, 2008.

Appendices

Appendix A

Glossary

Abbreviations and Symbols for Chapter 1

Nash Equilibrium In a multiplayer game a Nash Equilibrium is a set of strategies for each player such that (a) each player knows the strategies of the other players, and (b) given the strategies of the other players no player can improve their outcome by changing strategy. 2

Abbreviations and Symbols for Chapter 2

α **(Protesters' visibility)** Determines the location of the jump in the step-visibility function $v(\cdot;\alpha)$, see visibility threshold $1-\alpha$. 15

β **(Police capacity)** Determines the location of the jump in the policing term $p(\cdot;\beta)$. 16

BinCDF$(x;n,p)$ Cumulative binomial distribution function evaluated at x with n trials each with success probability p. 19

c_1 **(Protesters' enthusiasm)** Determines the rate at which individuals join the revolution. 15

c_2 **(Policing efficiency)** Determines the rate at which individuals are removed from the revolution. 16

$c^* = c_1/(c_1 + c_2)$ Ratio of protesters' enthusiasm to sum of protesters' enthusiasm and policing efficiency. 22

$d = d(r)$ **(Decay term)** Models the decay of the fraction of protesters in the population. 15

E **(Edges)** Set of edges representing social interactions between individuals in the network $G = G(V, E)$. 19

$g = g(r)$ **(Growth term)** Models the growth of the fraction of protesters in the population. 15

$G = G(V, E)$ **(Network)** Social network representing individuals (i.e. nodes V) connected by social interactions (i.e. edges E). 19

$\mathbb{I}_{\{\cdot\}}$ **(Indicator Function)** $\mathbb{I}_{\{X\}} = 1$ if X is true, $\mathbb{I}_{\{X\}} = 0$ otherwise. 15

$p(r; \beta)$ **(Policing term)** A step function that shuts off the decay term when the fraction of protesters is above the policing capacity β. 16

ϕ **(Average degree)** Average degree of individual $i \in V$ in the network $G = G(V, E)$. 19

$r = r(t)$ **(Fraction of protesters)** Fraction of protesters and/or revolutionaries in the population at time t. 15

$r = 0$ **(Total state control)** Equilibrium of total state control. 16

$r = 1$ **(Realized revolution)** Equilibrium of the realized revolution. 16

$r = c^*$ **(Civil unrest)** Equilibrium of civil unrest. 34

Region I Parameter regime $\alpha + \beta = 1$. 20

Region II (Failed State) Parameter regime $\alpha + \beta < 1$. 14, 20, 33

Region III Parameter regime $\alpha + \beta > 1$. 20

Region III0 (Stable police state) Parameter regime $\alpha + \beta < 1$ and $c^* \leq 1 - \alpha$. 14, 24, 34

Region III1 (Unstable police state) Parameter regime, $\alpha+\beta<1$ and $c^*\geq\beta$. 14, 24, 34

Region IIIe (Meta-stable police state) Parameter regime $\alpha+\beta<1$ and $1-\alpha<c^*<\beta$. 14, 24, 34

θ_i **(Threshold)** Threshold parameter for individual i; individual i will consider joining the revolution only if the fraction of his/her neighbours already participating in the revolution exceeds θ_i. 18

θ **(Average threshold)** Average threshold, i.e. $\theta=\langle\theta_i\rangle$. 19

$v(r;\alpha)$ Visibility term; a step function which shuts off the growth term when the fraction of protesters is below the visibility threshold $1-\alpha$. 15

$1-\alpha$ **(Visibility threshold)** The visibility function $v(r;\alpha)$ is a step function whose transition occurs at the visibility threshold $1-\alpha$, i.e. $v(r;\alpha)=0$ if $r\leq 1-\alpha$ and $v(r;\alpha)=1$ otherwise. 15

V **(Nodes)** Set of nodes (i.e. individuals) in the social network $G=G(V,E)$. 19

Abbreviations and Symbols for Chapter 3

α **(Protesters' visibility)** Determines the location of the jump in the step-visibility function $v_s(\cdot;\alpha)$. 55

ABM Linear threshold agent-based model. 43

β **(Police capacity)** Determines the location of the jump in the policing term $p(\cdot;\beta)$. 52

$v_b(r;\theta,\rho_k)$ **(BVF)** Binomial visibility function. 44, 55

c_1 **(Protesters' enthusiasm)** Determines the rate at which individuals join the revolution. 51

c_2 Policing efficiency parameter; determines the rate at which individuals are removed from the revolution. 52

$c^* = c_1/(c_1 + c_2)$ Ratio of protesters' enthusiasm to sum of protesters' enthusiasm and policing efficiency. 71

Clustering coefficient (Local) The local clustering coefficient for node i with degree d_i is the number of wedges, i.e. $d_i(d_i - 1)/2$, divided by the number of triangles involving node i (node i is in a triangle with nodes j and k if (i, j), (j, k) and (k, i) are edges in the network). 47

Cumulative degree distribution The cumulative degree distribution is the function $F(k) = \sum_{j=1}^{k} \rho_j$, where ρ_j is the degree distribution. 47

$d(t) = d(t|t_0)$ **(Decay term)** Models the decay of the fraction of protesters in the population. 53

DA Degree approximation. 44, 66

E **(Edges)** Set of edges representing social interactions between individuals in the network $G = G(V, E)$. 44

$v_e(r; \theta)$ **(EVF)** Empirical visibility function. 44, 57

$g(t) = g(t|t_0)$ **(Growth term)** Models the growth of the fraction of protesters in the population. 53

$G = G(V, E)$ **(Network)** Social network representing individuals (i.e. nodes V) connected by social interactions (i.e. edges E). 44

G_F **(Facebook subnetwork)** Network constructed from the ego-networks of ten individual Facebook members. 47

G_P **(Physical Contact network)** Network generated by distributing wireless sensors to students, teachers, and staff at a U.S. highschool during a one day period. 47

$\mathbb{I}_{\{\cdot\}}$ **(Indicator Function)** $\mathbb{I}_{\{X\}} = 1$ if X is true, $\mathbb{I}_{\{X\}} = 0$ otherwise. 53

$\bar{n}_k = \sum_l \rho_{k,l} r^{(l)}$ Fraction of nodes $n \in N_k$ in state 1. 67

N_k Set of nodes that have at least one neighbour of degree k. 66

$\nu(r_a(t_0); \theta)$ **(Visibility function)** Visibility function for the linear threshold ABM. 53

ODE Ordinary differential equation. 44

$r_a(t) = r_a(t|t_0)$ **(Fraction of protesters - ABM)** The fraction of nodes in the linear threshold ABM model that are expected to be in state 1 (i.e. active) at time t, conditioned on information at time t_0. 52

$r^{(k)}(t) = r^{(k)}(t|t_0)$ **(Fraction of protesters by degree- DA)** Fraction of population with degree k that are expected to be in state 1 (i.e. active) at time t, conditioned on information at time t_0. 67

$r_d(t) = r_d(t|t_0)$ **(Fraction of protesters - DA)** Fraction of population that are expected to be in state 1 (i.e. active) at time t, conditioned on information at time t_0. 67

R_0 Basic reproductive number. 46, 70

$\rho_{k,j}$ Degree distribution of N_k. 67

ρ_k **(Degree distribution)** The degree distribution ρ_k is the fraction of nodes in the network with degree k, i.e. the fraction of nodes in the network with k neighbours. 47, 51, 55

Region III0 (Stable police state) SVF model parameter regime $\alpha + \beta < 1$ and $c^* \leq 1 - \alpha$. 58

Region III1 (Unstable police state) SVF model parameter regime, $\alpha + \beta < 1$ and $c^* \geq \beta$. 58

Region IIIe (Meta-stable police state) SVF model parameter regime $\alpha + \beta < 1$ and $1 - \alpha < c^* < \beta$. 58

$s_v(t)$ **(State)** State of individual v at time t, i.e. $s_v(t) = 1$ if individual v is active, $s_v(t) = 0$ if individual v is inactive. 51

Sparsity The sparsity of a networks is the number of edges divided by the number of possible edges, i.e. the sparsity is the number of edges divided by $N(N-1)/2$ where N is the number of nodes in the network. 47

$v_s(r; \alpha)$ **(SVF)** Step visibility function. 45, 54

θ_v **(Linear threshold)** Individual v may change states from inactive to active only once the fraction of his/her neighbours that are already active exceeds the linear threshold θ_v. 43

V **(Nodes)** Set of nodes (i.e. individuals) in the social network $G = G(V, E)$. 44

V_k Set of nodes with degree k. 66

Abbreviations and Symbols for Chapter 4

a **(Relative conformity)** Determines the balance between individual utility and social utility. 83

b Timescale of the ordinary differential equation describing the dynamics of smoking prevalence. 83

$c(t)$ **(Tobacco consumption)** Cigarette consumption in grams per person per day at time t. 86

δ **(Information discounting)** Discounting factor for new information on the health effects of smoking. 84

IDV A measure of collectivism versus individualism developed by Hofstede [86]. 80

$n(t)$ **(Proxy for smoking knowledge)** Cumulative number of scholarly articles on the health effects of smoking. 84

OECD Organization for Economic Co-operation and Development. 80

s_x **(Slope)** Average rate of increase in smoking prevalence. 80, 96

Social utility from smoking The utility an individual derives indirectly from smoking through social interactions with other smokers. 80

t_{max} **(Peak year)** Peak year in tobacco consumption. 80

Total utility from smoking The total utility from smoking is a weighted product of the individual utility from smoking and the social utility from smoking. 80

u_x **(Individual utility from smoking)** The utility an individual derives directly from the act of smoking. 80, 83

u_y **(Individual utility from not smoking)** The utility an individual derives directly from the act of non-smoking. 83

u_0 **(Individual utility from smoking - no knowledge)** The individual utility from smoking in the absence of any knowledge of the health effects of smoking. 84

u_∞ **(Individual utility from smoking - perfect knowledge)** The individual utility from smoking with perfect knowledge of the health effects of smoking. 84

$x = x(t)$ **(Prevalence of smoking)** Fraction of the population that smokes at time t. 83

Abbreviations and Symbols for Chapter 5

$a_N(x_i; \vec{x})$ **(Obesity dynamics - advection term)** Deterministic term of the Langevin equation that describes the dynamics of individuals' BMI; for use with finite population size (N). 111

$a(x)$ **(Obesity dynamics - advection term)** Deterministic term of the Langevin equation that describes the dynamics of individuals' BMI; alternatively the advection term in the advection-diffusion (Fokker-Planck) equation. 112, 113

A_{ij} **(Weighting function)** Weighting function that specifies the strength of the social influence of individual j on individual i. 111, 112

advection-diffusion equation A partial differential equation that describes the aggregate behaviour of particles subject to advection and diffusion phenomena; also known as a Fokker-Planck equation. 113

AIC Akaike Information Criterion. 123

$b(x)$ **(Obesity dynamics - diffusion term)** Deterministic component of the stochastic term of the Langevin equation that describes the dynamics of individuals' BMI; alternatively the diffusion term in the advection-diffusion (Fokker-Planck) equation. 112, 113

BMI Body mass index; weight (in kg) divided by height squared (in meters^2). 106

BRFSS Behavioral Risk Factor Surveillance System. 115

$\eta(t)$ **(White noise)** Gaussian white noise. 112

$f_{\log}(x;\mu,\sigma)$ Log-normal distribution function. 117, 118

$f_{skew}(x;\xi,\omega,\alpha)$ Skew-normal distribution function. 118

k_I Proportionality constant for the gradient of the individual utility that, together with k_S, sets the relative importance of individual versus social utility in the advection term $a(x)$. 111

k_S Proportionality constant for the gradient of the social utility that, together with k_I, sets the relative importance of individual versus social utility in the advection term $a(x)$. 111

k_b Proportionality constant for the temporal variance in individuals' BMI. 114

$k_0 = k_I/k_b$ Ratio of k_I to k_b. 117

Langevin equation A stochastic differential equation that can be used to describe the individual behaviour of particles subject to advection and diffusion phenomena. 112

NHANES National Health and Nutrition Examination Survey. 115

NU Medical records for patients of the Northwestern Medicine system of hospitals and clinics. 115

$p(x,t)$ The probability density function for BMI x and time t. 113

$p_{eq}(x)$ Equilibrium distribution function for the social model. 115, 117

$p_{eq}^{(0)}(x)$ Equilibrium distribution function for the non-social model of individuals' BMI dynamics. 117

$\phi_{\mu,\sigma}(x)$ The probability density function for a normal random variable with mean μ and standard deviation σ evaluated at x. 112

$\Phi_{\mu,\sigma}(x)$ The cumulative density function for a normal random variable with mean μ and standard deviation σ evaluated at x. 118

σ Width of Gaussian function used to specify weighting function A_{ij}. 112

$S(k_I, x^\star, k_S, \sigma)$ Objective function for estimation of k_I, $x^\star$, k_S, and σ at the individual-level. 119

$\hat{S} = \min_{(k_I, x^\star, k_S, \sigma)} S(k_I, x^\star, k_S, \sigma)$ Error resulting from fitting social model to empirical micro-level BMI data. 119

$\hat{S}\Big|_{k_S=0} = \min_{(k_I, x^\star)} S(k_I, x^\star, 0, -)$ Error resulting from fitting nonsocial model to empirical micro-level BMI data. 119

$u_I(x) = u_I(x; x^\star)$ **(Individual utility)** The utility an individual derives from an action in the absence of any social interactions. 110

$u_S(x) = u_S(x; x^\star)$ **(Social utility)** The utility an individual derives from an action in the absence of any social interactions. 110

$x^\star$ **(Setpoint)** An individual's natural equilibrium BMI. 109, 110

$x = x(t)$ The BMI of an individual at time t. 110

Appendix B

Supplementary Materials for Chapter 2: The Arab Spring

B.1 Examples of sigmoidal functions

The family of sigmoidal functions $\mathbb{S}$ was defined in Section 2.4 to be the set of twice piecewise differentiable functions $s : [0, 1] \to [0, 1]$ satisfying

(i) $\exists s_0, s_1 \in [0, 1]$ such that

- $\forall r \in (s_0, s_1) : s(r)$ is twice differentiable,
- $\forall r \leq s_0 : s(r) = 0$, and
- $\forall r \geq s_1 : s(r) = 1$,

(ii) $\forall r \in (s_0, s_1) : s'(r) > 0$,

(iii) $\exists \xi \in (s_0, s_1)$ such that

- $\forall r \in [s_0, \xi] : s''(r) \geq 0$ and
- $\forall r \in [\xi, s_1] : s''(r) \leq 0$,

(iv) $\forall r \in (s_0, s_1) : s''(r)s(r) \leq [s'(r)]^2$.

We now consider three examples of functions that belong to $\mathbb{S}$.

Example 1: By construction, the function

$$s(r) = \begin{cases} 0 & \text{if } x \leq s_0 \\ \frac{r-s_0}{s_1-s_0} & \text{if } x \in (s_0, s_1) \\ 1 & \text{otherwise} \end{cases},$$

satisfies conditions (i)-(iii). This function also satisfies condition (iv), since

$$s''(r)s(r) = 0 \leq (s_1 - s_0)^{-2} = [s'(r)]^2.$$

Example 2: By construction, the function

$$s(r) = \begin{cases} 0 & \text{if } x \leq s_0 \\ \frac{1}{2}\sin\left[\frac{\pi}{2}\left(\frac{2r-s_0-s_1}{s_1-s_0}\right)\right] + \frac{1}{2} & \text{if } x \in (s_0, s_1) \\ 1 & \text{otherwise} \end{cases},$$

satisfies conditions (i)-(iii). For convenience, we denote $\Theta(r) = \frac{\pi}{2}\left(\frac{2r-s_0-s_1}{s_1-s_0}\right)$. We now verify that $s(r)$ satisfies condition (iv), since

$$\begin{aligned} s''(r)s(r) - [s'(r)]^2 &= -\frac{1}{4}[\Theta'(r)]^2\left(\sin^2[\Theta(r)] + \sin[\Theta(r)] + \cos^2[\Theta(r)]\right) \\ &= -\frac{1}{4}[\Theta'(r)]^2(1 + \sin[\Theta(r)]) \\ &\leq 0. \end{aligned}$$

Example 3: By construction, the function

$$s(r) = \begin{cases} 0 & \text{if } x \leq s_0 \\ \frac{1}{2}\tanh[\log(r-s_0) - \log(s_1-r)] + \frac{1}{2} & \text{if } x \in (s_0, s_1) \\ 1 & \text{otherwise} \end{cases},$$

satisfies conditions (i)-(iii). For convenience, we denote $\Theta(r) = \log(r-s_0) - \log(s_1-r)$.

We now compute $s''(r)s(r) - [s'(r)]^2$.

$$\begin{aligned}
s''(r)s(r) - [s'(r)]^2 &= \left(-\operatorname{sech}^2 [\Theta(r)] \tanh [\Theta(r)] [\Theta'(r)]^2 + \frac{1}{2} \operatorname{sech}^2 [\Theta(r)] \Theta''(r) \right) \\
&\quad \times \left(\frac{1}{2} \tanh [\Theta(r)] + \frac{1}{2} \right) - \frac{1}{4} \operatorname{sech}^4 [\Theta(r)] [\Theta'(r)]^2 \\
&= -\frac{1}{4} [\Theta'(r)]^2 \operatorname{sech}^2 [\Theta(r)] \\
&\quad \times \underbrace{\left((\tanh [\Theta(r)] + 1)^2 - \frac{2r - s_0 - s_1}{s_1 - s_0} (\tanh [\Theta(r)] + 1) \right)}_{\geq 0 \text{ if } r \leq \frac{s_0+s_1}{2}} \qquad \text{(B.1)} \\
&= -\frac{1}{4} [\Theta'(r)]^2 \operatorname{sech}^2 [\Theta(r)] (\tanh [\Theta(r)] + 1) \\
&\quad \underbrace{\times \left(\tanh [\Theta(r)] + 2 \frac{s_1 - r}{s_1 - s_0} \right)}_{\geq 0 \text{ if } r \geq \frac{s_0+s_1}{2}} . \qquad \text{(B.2)}
\end{aligned}$$

Because $-\frac{2r-s_0-s_1}{s_1-s_0} \geq 0 \iff r \leq \frac{s_0+s_1}{2}$, the right hand side of Eq. (B.1) is non-positive whenever $r \leq \frac{s_0+s_1}{2}$. Because $\tanh [\Theta(r)] \geq 0 \iff r \geq \frac{s_1+s_0}{2}$, the right hand side of Eq. (B.2) is non-positive whenever $r \geq \frac{s_0+s_1}{2}$. It follows that

$$\forall r \in (s_0, s_1) : s''(r)s(r) - [s'(r)]^2 \leq 0,$$

i.e. $s(r)$ satisfies condition (iv).

Appendix C

Supplementary Materials for Chapter 3: A Hierarchy of Linear Threshold Models

C.1 Physical Contact Network Data

In this appendix, we briefly discuss the network extraction protocol for the physical contact network G_P from [158].

The network presented in [158] was constructed by distributing wireless sensors to students, teachers, and staff at a U.S. high school during a one day period from approximately 08:00 to 16:30. When two wireless sensors are in proximity of one another, i.e. when they are less than approximately 3m apart, they register an interaction with a temporal resolution of 20s. Data are given for each separate interaction in comma separated value (CSV) format with three columns: identification number (ID) of first wireless sensor, ID of second wireless sensor, duration of interaction (measured in 20s increments). A weighted undirected network $G_{P,0}$ is formed by connecting each pair of individuals with an edge whose weight is given by the total amount of time they spent in proximity to one another. In order to admit a comparison with the unweighted undirected Facebook subnetwork, we de-weight network $G_{P,0}$ by discarding all edges whose weight is less than the minimum duration w_P, and by weighting all remaining edges equally. We select the largest connected component of the resulting network, and for convenience we refer to this network as the *physical contact network,* G_P.

The choice of the minimum duration w_P is important since it determines, for example, the average degree of the physical contact network G_P. We choose w_P such that the average degree of G_P is close to the average degree of G_F, for the following reason. In the simulations carried out in Chapter 3 to compare propagation under the linear threshold ABM on G_P and G_F, we seed the networks with a fixed percentage of active nodes and use the same threshold θ on both networks. We want to calibrate the average degree of G_P to the average degree of G_F, such that, on average, nodes in G_P have the same chance as nodes in G_F to satisfy the linear threshold criterion and see the revolution. After this calibration, differences in propagation between G_P and G_F (for the same θ) are only due to the differences in network structure that go beyond the average degree. This motivates us to choose $w_P = 34$ sensor measurements, equivalent to 11 minutes and 20 seconds, in order to match the average degrees of G_P and the Facebook subnetwork as closely as possible.

C.2 Gillespie's Algorithm

Following [6], we numerically simulate the ABM by implementing Gillespie's Algorithm. In this appendix we give a brief overview of this algorithm. We first introduce some notation. As above, let r_a denote the fraction of individuals in the population that are expected to be active in the revolution at time t. For each node $v \in V$ we let $\gamma_v(t) = 1$ if v can see the revolution at time t, i.e. if v satisfies Eq. (3.1), and let $\gamma_v(t) = 0$ otherwise. Furthermore, we let $\xi_{1,v}$ and $\xi_{2,v}$ denote the first arrival times of independent Poisson processes with rates c_1 and c_2, respectively.

Gillespie's Algorithm is based on the fact that the sum of two independent Poisson variables is also a Poisson variable with rate equal to the sum of the rates of the original processes. It follows that

$$\sum_{v \in V} \left[\xi_{1,v}\ (1 - s_v(t))\ \gamma_v(t) + \xi_{2,v}\ s_v(t)\ p(r_a(t); \beta)\right]$$

is a Poisson process with rate

$$\Lambda = \sum_{v \in V} \left[c_1\ (1 -\ s_v(t))\ \gamma_v(t) + c_2\ s_v(t)\ p(r_a(t); \beta)\right].$$

The first arrival time of this process, therefore, is an exponential random variable τ with rate Λ. At time $t + \tau$ the state of exactly one of the nodes will change. Moreover, since $\xi_{1,v}$ and $\xi_{2,v}$ are independent, the probability that the state of node v will change is

$$\mathbb{P}_v = \frac{c_1\ (1 -\ s_v(t))\ \gamma_v(t) + c_2\ s_v(t)\ p(r_a(t); \beta)}{\Lambda}$$

The Gillespie Algorithm then proceeds iteratively in three steps.

1. Find the time τ of the next event by drawing τ from an exponential distribution with rate Λ.
2. Determine which node changes state by drawing one node from $V = \{v_i\}_{i=1}^{N}$, where node $v \in V$ is drawn with probability $\mathbb{P}_v$.
3. Update $t \leftarrow t + \tau$ and re-calculate $s_v(t)$, $r(t)$, $\gamma_v(t)$.

We include the implementation of this algorithm in Matlab below.

ABM_Gil.m (Matlab m-file)

```
% ABM simulation that uses the Gillespie's Algorithm
%
% Input:
%   L:              cell array, cell i is a row vector listing
%                       neighbours of node i
%   itn:            number of iterations to record
%   dt:             size of each iteration
%   theta:          threshold (scalar or Nx1 vector)
%   c1, c2:         protesters' enthusiasm, policing efficiency
%   beta:           police capacity
%   ic:             initial condition (fraction chosen unif. at random)
%
% Output:
%   r:              r(t) for t = 0:dt:itn*dt
%   T_final:        the time at which the dynamics end (or inf if
%                       dynamics do not end before itn*dt)

function [r, T_final] = ABM_Gil(L, itn, dt, theta, c1, c2, beta, ic)
N = length(L);              % Number of individuals
T_final = itn*dt;           % Placeholder for T_final

%% Set up initial condition
r = zeros(itn+1, 1);        % Record r at equal timepoints given in time
r(1) = floor(ic*N)/N;

time_temp = 0;              % Keep track of time inside the loop
time_next = 1;              % *dt = Next time to record r
r_temp = r(1);              % Keep track of r inside loop

% Select floor(ic*N) individuals uniformly at random to be active
A_num = floor(ic*N);        % Number of active nodes

A_idx = zeros(N,1);         % Indices of active nodes padded with 0's
A_idx(1:A_num) = datasample(1:N, A_num, 'Replace', false);

A = zeros(N,1);             % A(i) = 1 if i is active (=0 o/w)
```

ABM_Gil.m (Matlab m-file - continued)

```
A( A_idx ) = 1;

% Fraction of neighbours that are active
neigh_activity = zeros(N,1);

for i=1:N
    % calculate fraction of neighbours that are active
    neigh_activity(i) = mean( A(L{i}) );
end

% Nodes who are considering joining the revolution
C = (1 - A).*(neigh_activity >= theta);

C_num = sum(C);             % Number of considering
C_idx = [find(C);...
    zeros(N-C_num, 1)];     % Indices of considering nodes (pad with 0)

% Calculate rate (Lambda)
rate = c1*C_num + c2*A_num*(r_temp < beta);

% While the time elapsed is less that T_final
% While 0 < r_temp < 1 (r = 0 and r = 1 are absorbing states)
% While rate > 0 (this condition implies 0 < r_temp < 1)
while time_temp < T_final && rate > 0
    % Step 1: Choose time of next event and event index

    % Choose tau ~ Exponential(rate)
    tau = 1/rate*log(1/rand);
    time_temp = time_temp + tau;

    % Step 2: Determine which node changes state
    % Step 3: Update

    % Determine whether C->A (inactive to active) or A->I/C (active to
    % inactive)
    if (r_temp >= beta) || (rand < c1*C_num/(c1*C_num + c2*A_num))
```

ABM_Gil.m (Matlab m-file - continued)

```
        % If r(itn-1) >= beta then only C->A is possible
        % If rand~U(0,1) < c1*C_num/(...) then choose C->A
        % Choose one of C_idx at random to change from C to A.
        mu = datasample(1:C_num,1);
        CtoA = C_idx(mu);

        % Update status of chosen node
        A(CtoA) = 1;
        A_num = A_num + 1;
        A_idx(A_num) = CtoA;

        C(CtoA) = 0;
        C_num = C_num - 1;
        C_idx = [C_idx(1:mu-1); C_idx(mu+1:end); 0];

        neigh_change = L{CtoA};
    else
        % Otherwise choose A->I/C
        % Choose one of A_idx at random to change from A to I/C
        mu = datasample(1:A_num,1);
        AtoI = A_idx(mu);

        % Update status of chosen node
        % Since we do not know whether A->I or A->C we update C with
        % the update of neighbours (below)
        A(AtoI) = 0;
        A_num = A_num-1;
        A_idx = [A_idx(1:mu-1);A_idx(mu+1:end);0];

        neigh_change = [AtoI, L{AtoI}];
    end
```

ABM_Gil.m (Matlab m-file - continued)

```
    % Step 3: (Continued) update neigh_activity, C, C_num, C_idx
    for i = neigh_change
        neigh_i = L{i};

        neigh_activity(i) = mean(A(neigh_i));

        if length(theta) > 1
            C(i) = (1-A(i))*(neigh_activity(i)>=theta(i));
        else
            C(i) = (1-A(i))*(neigh_activity(i)>=theta);
        end

        % if C(i)==1 then add/remove i to C_idx and C_num, if necessary
        mu = find(C_idx==i, 1);

        if (C(i)==1) && isempty(mu)
            C_num = C_num+1;
            C_idx(C_num) = i;
        elseif (C(i)==0) && ~isempty(mu)
            C_num = C_num-1;
            C_idx = [C_idx(1:mu-1);C_idx(mu+1:end);0];
        end
    end

    r_temp = mean(A);

    % Step 4: Record, if necessary
    while (time_next<itn+1) && (time_next*dt < time_temp)
        time_next = time_next+1;
        r(time_next) = r_temp;
    end

    % Recalculate rate for next step
    rate = c1*C_num + c2*A_num*(r_temp < beta);
end
```

ABM_Gil.m (Matlab m-file - continued)

```
% If the while loop ends due to an absorbing state or bottle neck (i.e.
% r_temp=0, r_temp=1, or rate=0) then fill out the rest of the vector r
% with the last value of r_temp.
r(find(r==0,1):end) = r_temp;

if time_temp < T_final
    T_final = time_temp;
else
    T_final = inf;
end
```

C.3 Optimization of visibility parameter (α) for the step-visibility function (SVF) model

In this section we present the Matlab code used to compute the optimal α for the SVF model for a given θ of the ABM, see Section 3.4.3. The file dependencies are as follows:

- *fit_alpha.m* $\rightarrow$ *ABM_Gil.m*, *StepOBJ.m*
- *StepOBJ.m* $\rightarrow$ *StepSoln.m*

In words, *fit_alpha.m* is the root m-file that provides the Matlab code for computing the optimal α for the SVF model for a given θ of the ABM. The m-file *fit_alpha.m* uses the file *ABM_Gil.m* to simulate the ABM via Gillespie's Algorithm (see Appendix C.2). The optimal α is computed using the Matlab function *fminsearch* with the objective function defined in *StepOBJ.m*. Finally, the file *StepOBJ.m* uses the file *StepSoln.m* which computes the solution to the SVF ODE as given in Section 2.3.2.

fit_alpha.m (Matlab m-file)

```
% fit_alpha.m: Find the alpha that best fits the output of the SVF 1-D ODE
% model to the simulated ABM process
%
% In order for this procedure to work, parameters must be in Region III0
% and initial condition must be in (1-alpha, beta).
%
% Input:
%   - X: Nx1 cell array, cell i lists neighbours of node i
%   - params: [c1, c2, beta, r0, itn, dt, rep]
%       - c1, c2, beta, r0 are parameters for the SVF ODE
%       - itn + 1 is the number of times which the ABM is recorded
%       - dt is the amount of time between ABM recorded values
%       - rep is the number of realizations of the ABM
%   - theta: linear thresholds (scalar or Nx1 vector)
%
% Global variables:
%   - rep: number of realizations of ABM
%   - r0: initial condition for which to solve SVF ODE
%   - beta, c_1, c_2: parameters for SVF ODE
%   - t: times at which to solve SVF ODE
%   - r: 2-dimensional array containing ABM realizations.
%       - Rows are different times (t)
%       - Columns are different realizations (rep).
%
% Output:
%   - alpha: optimal alpha
%   - R: First rep columns ABM realizations. Last column sol'n to SVF ODE
%   - error_alpha:[alpha values,  L2 error between ABM
%       simulations and SVF ODE]
%
% m-file dependencies:
%   - ABM_Gil.m
%   - StepOBJ.m
```

fit_alpha.m (Matlab m-file - continued)

```
function [alpha, R, error_alpha] = fit_alpha(X, params, theta)
global rep r0 beta c1 c2 t r

c1 = params(1); c2 = params(2); beta = params(3); r0 = params(4);
itn = params(5); dt = params(6); rep = params(7);

t = (0:dt:dt*itn)';

r = zeros(itn+1, rep);

%% Step 1: Simulate ABM

for j = 1:rep
    r(:,j) = ABM_Gil(X, itn, dt, theta, c1, c2, beta, r0);
end

R = r;

%% Step 2: Fit alpha to simulation data

alpha = fminsearch(@StepOBJ, 1-theta);

[~, Y_temp] = StepOBJ(alpha);

R = [R, Y_temp];

%% Step 3: Find error

error_alpha = (.70:0.001:.9)';
error_alpha = [error_alpha, zeros(length(error_alpha), 2)];

for j = 1:length(error_alpha)
    error_alpha(j,2) = StepOBJ(error_alpha(j,1));
end
```

StepOBJ.m (Matlab m-file)

```
% Compute the L2-error for between the solution to the SVF ODE and the ABM
% realizations.
%
% Input:
%   - x = alpha
%
% Global variables:
%   - rep: number of realizations of ABM
%   - r0: initial condition for which to solve SVF ODE
%   - beta, c_1, c_2: parameters for SVF ODE
%   - t: times at which to solve SVF ODE
%   - r: 2-dimensional array containing ABM realizations.
%       - Rows are different times (t)
%       - Columns are different realizations (rep).
%
% Output:
%   - S: L-2 error
%   - Y: vector giving SVF solution

function [S,Y] = StepOBJ(x)
global rep r0 beta c1 c2 t r

Y = StepSoln(x, [c1, c2, beta, r0], t);

S = kron(Y, ones(1,rep));
S = ( S - r ).^2;
S = sum(sum(S));
```

StepSoln.m (Matlab m-file)

```
% Calculate the solution to SVF ODE
%
% Input:
%   - x = alpha
%   - params = [c1, c2, beta, r0]
%    - r0: initial condition for which to solve ODE
%    - beta, c_1, c_2: parameters for ODE
%   - t: times at which to solve ODE
%
% Output:
%   - Y: solution to ODE

function Y = StepSoln(x, params, t)
% make sure that t is a column vector
t = reshape(t, length(t), 1);

c1 = params(1);
c2 = params(2);
beta = params(3);
r0 = params(4);

if (1-x <= beta) && (c1/(c1+c2) <= 1-x) && (r0 < beta) && (1-x < r0)
    % Region III0/I with r0 in (1-alpha, beta)
    talpha = 1/(c1+c2)*log( (c1-r0*(c1+c2)) / (c1-(1-x)*(c1+c2)) );

    idx = find(t > talpha,1);

    if isempty(idx)
        idx = length(t) + 1;
    end

    Y = (c1 - (c1-r0*(c1+c2)) * exp( -(c1+c2)*t(1:(idx-1)) ) ) / (c1+c2);
    Y = [Y; (1-x)*exp( -c2*(t(idx:end) - talpha) )];
```

StepSoln.m (Matlab m-file - continued)

```
elseif (1-x <= beta && r0 >= beta) || (beta < 1-x && r0 > 1-x)
    % Region I, III0, IIIe, III1 with r0 >= beta, or
    % Region II with r0 > 1-alpha
    Y = 1-(1-r0)*exp(-c1*t);

elseif (1-x <= beta && r0 <= 1-x) || (beta < 1-x && r0 < beta)
    % Region I, III0, IIIe, III1, with r0 <= 1-alpha, or
    % Region II with r0 < beta
    Y = r0*exp(-c2*t);

elseif c1/(c1+c2) < beta && c1/(c1+c2) > 1-x && r0 < beta && 1-x < r0
    % Region IIIe with r0 in (1-alpha, beta)
    Y = c1/(c1+c2) - (c1/(c1+c2) -  r0)*exp(-(c1+c2)*t);

elseif 1-x <= beta && beta <= c1/(c1+c2) && r0 < beta && 1-x < r0
    % Region I/III1 with r0 in (1-alpha, beta)
    tbeta = -1/(c1+c2)*log( (c1/(c1+c2) - beta)/(c1/(c1+c2) - r0) );

    idx = find(t>tbeta,1);

    if isempty(idx)
        idx = length(t) + 1;
    end

    Y = c1/(c1+c2) - (c1/(c1+c2) - r0)*exp(-(c1+c2)*t(1:idx-1));
    Y = [Y; 1 - (1-beta)*exp(-c1*(t(idx:end)-tbeta))];

elseif beta < 1-x && r0 >= beta && r0 <= 1-x
    % Region II with r0 in[beta, 1-alpha]
    Y = r0*ones(length(t),1);

end
```

Appendix D

Supplementary Materials for Chapter 4: Modelling the Prevalence of Smoking

D.1 Smoking Prevalence (x) and Cigarette Consumption (c) Data

The following is a link to Additional file 2 of [108], which is a comma separated values (CSV) file containing four columns: country number as it appears in Table 4.1, year (t), measurement ($x(t)$ or $c(t)$), and type of measurement (0 indicates a smoking prevalence measurement, while 1 indicates a cigarette consumption measurement).

D.2 Proxy Data (n): articles published on the health effects of smoking

The following is a link to Additional file 3 of [108], which is a comma separated values (CSV) file containing three columns: year (t), number of articles published in year t, and cumulative number of articles published up to and including year t ($n(t)$).

D.3 Matlab Code for Modelling the Prevalence of Smoking

The following is a link to Additional file 4 of [108], which is ZIP file containing Matlab data files and simulation code used in preparation of Chapter 4.

D.4 US Population Data for Smoking Counterfactual Scenario

The following is a link to Additional file 5 of [108], which is comma separated values (CSV) file containing two columns: year (t) and population ($N_{pop}(t)$).

Appendix E

Supplementary Materials for Chapter 5: Modelling the Dynamics of the Body Mass Index (BMI)

E.1 Data

E.1.1 National Health and Nutrition Examination Survey (NHANES)

NHANES data are available from the NHANES website

http://www.cdc.gov/nchs/nhanes/nhanes_questionnaires.htm.

Directly measured BMI measurements are given by the variable BMXBMI. Self-reported BMI measurements are calculated from the variables WHD010 (self-reported height at time of interview) and WHD020 (self-reported weight at time of interview). Self-reported year-over-year change in BMI are calculated from self-reported BMI and from variables WHD010 and WHD050 (self-reported weight one year prior to interview).

Data were downloaded directly from the Centers for Disease Control and Prevention (CDC) website as ".XPT" files (in SAS format) and imported into Matlab. The variable BMXBMI is found in data files with names starting with "BMX", the variables WHD010, WHD020, and WHD050 are found in data files with names starting with "WHQ". We also

make use of the variable RIDAGEYR (age at time of interview), which is found in data files with names starting with "DEMO", and the SEQN variable (interviewee identification numbers), which is found in all data files. File names are completed by adding the suffix ".XPT" for survey year 1999-2000, "_B.XPT" for survey year 2001-2002, "_C.XPT" for survey year 2003-2004, etc...

E.1.2 Northwestern Medicine (NU) Medical Records

The NU data set consists of medical records from the Northwestern Medical system of hospitals and clinics, i.e., patients of the following Chicago area locations:

- Bucktown (1776 N. Milwaukee Avenue, Chicago, IL, 60647),
- Deerfield (350 S. Waukegan Road Suites 100, 150 and 200, Deerfield, IL, 60015),
- Delano Court (in the Roosevelt Collection, 1135 S. Delano Court Suite A201, Chicago, IL, 60605),
- Evanston (1704 Maple Avenue Suites 100 and 200, Evanston, IL, 60201),
- Glenview (2701 Patriot Boulevard, Glenview, IL, 60026),
- Grayslake (1475 E. Belvidere Road, Pavilion C Suite 385, Grayslake, IL, 60030),
- Highland Park (600 Central Avenue Suite 333, Highland Park, IL, 60035),
- Lake Forest Hospital (660 N. Westmoreland Rd., Lake Forest, IL, 60045),
- Libertyville (1800 Hollister Drive Suite 102, Libertyville, IL, 60048),
- Lakeview (1333 W. Belmont Avenue Suites 100 and 200, Chicago, IL, 60657),
- Loop 1 (20 S. Clark Street Suite 1100, Chicago, IL, 60603),
- Loop 2 (111 W. Washington St. Suite 1801, Chicago, IL, 60602),
- Northwestern Memorial Hospital (251 E. Huron St., Chicago, IL, 60611),
- River North (635 N. Dearborn Street Suite 100, Chicago, IL, 60654),
- Sauganash (4801 W Peterson, Suite 406, Chicago, IL, 60646),

- Skokie (10024 Skokie Blvd Suite 304, Skokie, IL, 60077), and
- SoNo (South of North Avenue, 1460 N. Halsted Street Suites 203, 502, and 504, Chicago, IL, 60642).

We note that the NU data set may contain multiple measurements per individual per year. In that case the BMI for individual i in year t is calculated using the average weight of individual i in year t and the average height of individual i taken over all years.

E.1.3 Behavioral Risk Factor Surveillance System (BRFSS)

BRFSS data are available from the BRFSS website

http://www.cdc.gov/brfss/annual_data/annual_data.htm.

BRFSS surveys record BMI measurements in variable _BMI for survey years 1984-1999, _BMI2 for survey years 2000-2002, _BMI3 for survey year 2003, _BMI4 for survey years 2004-2010, and _BMI5 for survey years 2011 onwards. Data were downloaded directly from the CDC website as ".XPT" files (in SAS format), converted to CSV format, and then imported into Matlab. File names for BRFSS survey data for years 1978–2010 start with "CDBRFS", while file names for BRFSS survey data for years 2011–2013 start with "LLCP". File names are completed by adding the suffix "87.XPT' for year 1987, "88.XPT" for year 1988, etc...

E.2 Additional Details for Methods

E.2.1 Properties of $p_{eq}^{(0)}(x; k_0, x^\star)$ (Eq. (5.16))

We note that for any population the BMI distribution must be strictly contained in the interval $[0, \infty)$. This implies that $p_{eq}(0) = 0$ and that $\lim_{x\to\infty} p_{eq}(x) = 0$. Assuming that $b(0) = 0$ (which holds for our model, see Eq. (5.15)), it follows that integrating Eq. (5.12) yields

$$0 = -p_{eq}(x)a(x) + \frac{1}{2}\frac{d}{dx}\left[p_{eq}(x)b^2(x)\right] + \underbrace{p_{eq}(0)a(0) - \frac{1}{2}\frac{d}{dx}\left[p_{eq}(x)b^2(x)\right]_{x=0}}_{=0}, \tag{E.1}$$

which has the solution

$$p_{eq}(x) = \xi \exp\left(2\int_0^x \frac{a(\tilde{x}) - b(\tilde{x})b'(\tilde{x})}{b^2(\tilde{x})} d\tilde{x}\right) \ , \tag{E.2}$$

where ξ is a normalization constant such that $\int_0^\infty p_{eq}(x)dx = 1$. When $a(x) = k_I(x^\star - x)$ and $b(x) = \sqrt{k_b}\, x$, i.e., when we assume $k_S = 0$ in Eqs. (5.10) and (5.15), we can re-arrange Eq. (E.1) to yield

$$\frac{dp_{eq}^{(0)}}{dx}(x) = 2\frac{k_I(x^\star - x) - k_b x}{k_b x^2} p_{eq}^{(0)}(x) \ ,$$

which implies that $p_{eq}^{(0)}(x)$ is a single peaked probability distribution whose mode[1] is given by the expression $x^\star \frac{k_I}{k_b}/(\frac{k_I}{k_b} + 1)$. We can also re-arrange Eq. (E.1) to yield

$$x p_{eq}^{(0)}(x) = -\frac{k_b}{k_I}\frac{d}{dx}\left[\frac{x^2}{2} p_{eq}^{(0)}(x)\right] + x^\star p_{eq}^{(0)}(x) \ ,$$

which implies that

$$\langle x \rangle = \int_0^\infty x p_{eq}^{(0)}(x)dx = -\frac{k_b}{k_I}\underbrace{\int_0^\infty \frac{d}{dx}\left[\frac{1}{2}x^2 p_{eq}^{(0)}(x)\right] dx}_{=0} + x^\star \underbrace{\int_0^\infty p_{eq}^{(0)}(x)dx}_{=1} = x^\star \ .$$

Multiplying Eq. (E.1) by x and re-arranging yields

$$x^2 p_{eq}^{(0)}(x) = -\frac{1}{2}\frac{k_b}{k_I}\frac{d}{dx}\left[x^3 p_{eq}^{(0)}(x)\right] + \frac{1}{2}\frac{k_b}{k_I} x^2 p_{eq}^{(0)}(x) + x x^\star p_{eq}^{(0)}(x) \ ,$$

which implies

$$\begin{aligned}
\left\langle x^2 \right\rangle &= \int_0^\infty x^2 p_{eq}^{(0)}(x)dx \\
&= -\frac{1}{2}\frac{k_b}{k_I}\int_0^\infty \frac{d}{dx}\left[x^3 p_{eq}^{(0)}(x)\right] dx + \frac{1}{2}\frac{k_b}{k_I}\int_0^\infty x^2 p_{eq}^{(0)}(x)dx + x^\star \int_0^\infty x p_{eq}^{(0)}(x)dx \\
&= \frac{1}{2}\frac{k_b}{k_I}\left\langle x^2 \right\rangle + x^{\star 2} \ .
\end{aligned}$$

Re-arranging now yields $\langle x^2 \rangle = 2x^{\star 2}\frac{k_I}{k_b}/(2\frac{k_I}{k_b} - 1)$. We note that we require $\langle x^2 \rangle \geq 0$, i.e. that $2k_I/k_b > 1$. We also note that this condition is satisfied by all empirical BMI

[1] The mode of a continuous random variable with probability density function $f(x)$ is $\text{argmax}_x f(x)$.

distributions in the NHANES, NU, and BRFSS data sets. Similarly, multiplying Eq. (E.1) by x^2, re-arranging, integrating, and solving for $\langle x^3 \rangle$, yields $\langle x^3 \rangle = \frac{k_I}{k_b}/(\frac{k_I}{k_b} - 1)\, \langle x^2 \rangle\, x^\star$.

The mean, variance, skewness, and mode skewness can now be computed using the following relations to the mode and the first three moments.

$$\begin{aligned} \text{mean} &= \langle x \rangle \\ \text{variance} &= \langle x^2 \rangle - \langle x \rangle^2 \\ \text{skewness} &= \frac{\langle x^3 \rangle - 3\, \langle x \rangle\, (\text{variance}) - \langle x \rangle^3}{(\text{variance})^{\frac{3}{2}}}, \text{ and} \\ \text{mode skewness} &= \frac{\text{mean} - \text{mode}}{(\text{variance})^{\frac{1}{2}}} \ . \end{aligned}$$

E.2.2 Solving Eq. (5.12) numerically for $p_{eq}^{(0)}(x)$ and $p_{eq}^{(n)}(x)$

For $a(x)$ and $b(x)$ as in Eqs. (5.10) and (5.15), respectively, the solution to Eq. (5.12) is given implicitly by Eq. (E.2), i.e., Eq. (E.2) becomes

$$p_{eq}(x) \propto p_{eq}^{(0)}(x; k_i/k_b, x^\star) \exp\left(2\frac{k_S}{k_b} \int_0^x \int_0^\infty \frac{\phi_{\tilde{x},\sigma}(\hat{x})(\hat{x} - \tilde{x}) \left(1 - \frac{1}{2}\frac{(\hat{x}-\tilde{x})^2}{\sigma^2}\right) p_{eq}(\hat{x})}{\tilde{x}^2} d\hat{x} d\tilde{x} \right) \ .$$

In order to solve for $p_{eq}(x)$ we implement the following iterative scheme.

$$\begin{aligned} p_{eq}^{(n+1)}(x) &= p_{eq}^{(n+1)}(x; k_I/k_b, x^\star, k_S/k_b, \sigma) \\ &\propto p_{eq}^{(0)}(x; k_I/k_b, x^\star) \exp\left(2\frac{k_S}{k_b} \int_0^x \int_0^\infty \frac{\phi_{\tilde{x},\sigma}(\hat{x})(\hat{x} - \tilde{x}) \left[1 - \frac{1}{2}\frac{(\hat{x}-\tilde{x})^2}{\sigma^2}\right] p_{eq}^{(n)}(\hat{x})}{\tilde{x}^2} d\hat{x} d\tilde{x} \right) \ . \end{aligned} \tag{E.3}$$

Let $m = 90$, $\Delta z = 1$, and $\forall i = 1, 2, \ldots, m : z_i = 10.5 + (i-1)\Delta z$. We set

$$p_{eq}^{(0)}(z_i) = \frac{z_i^{-2(k_I/k_b+1)} \exp\left(-2\frac{k_I}{k_b}\frac{x^\star}{z_i}\right)}{\sum_{j=1}^m z_j^{-2(k_I/k_b+1)} \exp\left(-2\frac{k_I}{k_b}\frac{x^\star}{z_j}\right) \Delta z} \ .$$

We then set

$$p_{eq}^{(n+1)}(z_i) = p_{eq}^{(0)}(z_i)\times$$
$$\exp\left(2\frac{k_S}{k_b}\sum_{k=1}^{i}\sum_{j=1}^{m}\frac{1}{2}\left(1+\mathbb{I}_{\{k<i\}}\right)\frac{\phi_{z_k,\sigma}(z_j)(z_j-z_k)\left(1-\frac{1}{2}\frac{(z_j-z_k)^2}{\sigma^2}\right)p_{eq}^{(n)}(z_j)}{z_k^2}\Delta z^2\right).$$

where $\mathbb{I}_{\{X\}} = 1$ if X is true and $\mathbb{I}_{\{X\}} = 0$ otherwise, and where we terminate the iterative process once

$$\left\|\frac{p_{eq}^{(n+1)} - p_{eq}^{(n)}}{m}\right\|_2 = \sqrt{\sum_{i=1}^{m}\left(\frac{p_{eq}^{(n+1)}(z_i) - p_{eq}^{(n)}(z_i)}{m}\right)^2} < 10^{-12}\ .$$

E.2.3 Fitting distribution functions to empirical BMI data

Suppose that $f(x;\theta)$ is a probability density function with parameters θ. We fit $f(x;\theta)$ to empirical BMI data measurements $\{x_i\}_{i=1}^N$ using the principle of maximum likelihood parameter estimation. Specifically, we set

$$\hat{\theta} = \underset{\theta}{\operatorname{argmax}}\prod_{i=1}^{N} f(x_i;\theta),$$

where $\mathcal{L}_f(\theta|x) = \prod_{i=1}^N f(x_i;\theta)$ is called the likelihood function. In Matlab we perform this optimization using the Matlab function *fminsearch* to solve the equivalent optimization problem

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}} -\log(\mathcal{L}_f(\theta|x)) = \underset{\theta}{\operatorname{argmin}} -\sum_{i=1}^{N}\log\left[f(x_i|\theta)\right]$$

We note that we compute a separate set of parameters for each year of BMI data.

Akaike Information Criterion (AIC)

The Akaike Information Criterion (AIC) is a generalization of the principle of maximum likelihood parameter estimation [4]. Suppose that we have data $\vec{x}$ that is distributed according to an unknown distribution function $g(x)$. Further suppose that we wish to approximate $g(x)$ by the distribution function $f(\cdot|\theta)$ with parameters θ. We can estimate

the maximum likelihood estimate (MLE) of θ by maximizing the average log-likelihood function $S_N(f(\cdot|\theta))$, i.e.

$$\hat{\theta} = \underset{\theta}{\operatorname{argmax}} \frac{1}{N} \log\left(\mathcal{L}_f(\theta|x)\right) = \underset{\theta}{\operatorname{argmax}} \underbrace{\frac{1}{N} \sum_{i=1}^{N} \log(f(x_i|\theta))}_{=S_N(f(\cdot|\theta))}.$$

It can be shown that the mean log-likelihood function $S_N(f(\cdot|\theta))$ converges with probability 1 to

$$S(g; f(\cdot|\theta)) = \int g(x) \log(f(x|\theta)) dx.$$

From this quantity we define

$$I(g; f(\cdot|\theta)) = S(g; g) - S(g; f(\cdot|\theta)),$$

which is called the Kullback-Leibler mean information for the discrimination between $g(x)$ and $f(\cdot|\theta)$, and which can be shown to be non-negative with $I(g; f(\cdot|\theta)) = 0 \iff f(x|\theta) = g(x)$ almost everywhere. Roughly speaking, $I(g; f(\cdot|\theta))$ can be interpreted as the amount of information lost when $f(\cdot|\theta)$ is used to approximate $g(x)$. This quantity induces a natural model selection criteria, i.e. we select the model that minimizes $I(g; f(\cdot|\theta))$.

The key observation for the establishment of the AIC criterion is that the quantity $I(g; f(\cdot|\theta))$ can be approximated as follows. Suppose that the true model is $g(x) = f(x|\theta_0)$ for some $\theta_0 \in \Theta$ and suppose that $\hat{\theta}$ is the MLE for the model restricted to some subspace $\Theta' \subset \Theta$, i.e.

$$\hat{\theta} = \underset{\theta \in \Theta' \subset \Theta}{\operatorname{argmax}} \mathcal{L}_f(\theta|x).$$

Then it can be shown that

$$\begin{aligned}
\mathbb{E}\left[2N\ I(g; f(\cdot|\hat{\theta}))\right] &= \mathbb{E}\left[2N\ I(f(\cdot|\theta_0); f(\cdot|\hat{\theta}))\right] \\
&= c + 2k - 2\sum_{i=1}^{N} \log\left(f(x_i|\hat{\theta})\right) \\
&= c + \underbrace{2k - 2\log\left(\mathcal{L}_f(\hat{\theta}|x)\right)}_{=AIC(f(\cdot|\hat{\theta}))},
\end{aligned}$$

where c is a constant, k is the dimension of Θ', and where $AIC(f(\cdot|\hat{\theta})) = 2k - \mathcal{L}_f(\hat{\theta}|x)$ is the Akaike Information Criterion (AIC). It follows that minimizing $I(p; f(\cdot|\theta))$ is equivalent to minimizing the AIC.

We now define the relative likelihood ratio for model selection. Suppose that we compute the AIC two different functional forms resulting in AIC values $AIC_1 = 2k_1 - 2\log(\mathcal{L}_1)$ and $AIC_2 = 2k_2 - 2\log(\mathcal{L}_2)$ with $AIC_1 < AIC_2$. Then the relative likelihood ratio

$$r = \exp\left(\frac{AIC_1 - AIC_2}{2}\right) = \exp(k_1 - k_2)\frac{\mathcal{L}_2}{\mathcal{L}_1}$$

is interpreted as follows: model 1 is r times as likely to be the "best" approximation to the true distribution than model 2. In this case, "best" is in the context of minimizing the AIC (i.e. minimizing the loss of information when using models 1 or 2 to approximate the "true" distribution $g(x)$).

E.2.4 Estimating parameters $(k_I, x^\star, k_S, \sigma)$ from individual level data

Consider individual i from survey year $t = t_1$ with BMI measurements at times t_1 and $t_2 = t_1 + \Delta t$, i.e. with $x_i(t_1)$ and $x_i(t_2)$. We denote the change in BMI by $\Delta x_i(t_t) = x_i(t_2) - x_i(t_1)$. For $\epsilon > 0$ we define

$$\hat{a}(x_i(t);\epsilon) = \frac{\sum_{j:|x_j(t)-x_i(t)|<\epsilon} \frac{\Delta x_j(t)}{\Delta t}}{N(x_i(t),\epsilon)}, \text{ and} \tag{E.4a}$$

$$\hat{b}(x_i(t);\epsilon) = \sqrt{\frac{\sum_{j:|x_j(t)-x_i(t)|<\epsilon} \frac{\Delta x_j^2(t)}{\Delta t}}{N(x_i(t),\epsilon)} - \Delta t\left[\frac{\sum_{j:|x_j(t)-x_i(t)|<\epsilon} \frac{\Delta x_j(t)}{\Delta t}}{N(x_i(t),\epsilon)}\right]^2}, \tag{E.4b}$$

respectively, where $N(x_i(t),\epsilon)$ is the number of individuals j with $|x_j(t) - x_i(t)| < \epsilon$, i.e.,

$$N(x_i(t),\epsilon) = \sum_{j:|x_j(t)-x_i(t)|<\epsilon} 1\ .$$

To estimate $\sqrt{k_b}$ we compute $\hat{b}(x_i(t);\epsilon)$ from BMI data and regress it on $x_i(t)$. Solving for the remaining parameters now breaks down into two cases, i.e. $k_S > 0$ and $k_S = 0$.

Assume that $k_S > 0$ in Eq. (5.10), i.e. that social interactions between individuals do have an effect on the dynamics of individuals' BMI. In this case, in order to estimate the remaining parameters, i.e. $(k_I, x^\star, k_S, \sigma)$, we define the objective function

$$S(k_I, x^\star, k_S, \sigma;\epsilon) = \sqrt{\frac{\sum_{i,t}\left[\hat{a}(x_i(t);\epsilon) - a(x_i(t))\right]^2}{\sum_{i,t} 1}}$$

and solve the optimization problem

$$\left(\hat{k}_I, \hat{x}^\star, \hat{k}_S, \hat{\sigma}\right) = \underset{(k_I, x^\star, k_S, \sigma)}{\operatorname{argmin}} \; S(k_I, x^\star, k_S, \sigma; \epsilon), \tag{E.5}$$

where we have suppressed the dependence of $(\hat{k}_I, \hat{x}^\star, \hat{k}_S, \hat{\sigma})$ on ϵ for convenience of notation. Recall from Eq. (5.10) that

$$a(x_i(t)) = k_I(x^\star - x_i(t)) + k_S \frac{du_S}{dx_i}(x_i(t)),$$

where

$$\frac{du_S}{dx_i}(x_i(t)) = \int_0^\infty \phi_{x_i(t),\sigma}(x_j(t))(x_j(t) - x_i(t)) \left(1 - \frac{1}{2}\frac{(x_j(t) - x_i(t))^2}{\sigma^2}\right) p(x_j(t), t) dx_j.$$

Observe that, for fixed σ, the objective function $S(k_I, x^\star, k_S, \sigma; \epsilon)$ is the objective function for the linear regression of $\hat{a}(x_i(t); \epsilon)$ on $-x_i(t)$, $\frac{du_S}{dx_i}(x_i(t); \sigma)$, and a constant. It follows, therefore, that there is a unique $(\tilde{k}_I(\sigma), \tilde{x}^\star(\sigma), \tilde{k}_S(\sigma))$ that solves

$$\left(\tilde{k}_I(\sigma), \tilde{x}^\star(\sigma), \tilde{k}_S(\sigma)\right) = \underset{(k_I, x^\star, k_S)}{\operatorname{argmin}} \; S(k_I, x^\star, k_S, \sigma; \epsilon) \;, \tag{E.6}$$

and that this solution can be computed using linear regression. Solving the optimization problem in Eq. (E.5) is now reduced to a one dimensional problem, i.e., we solve

$$\hat{\sigma} = \underset{\sigma}{\operatorname{argmin}} \; S\left(\tilde{k}_I(\sigma), \tilde{x}^\star(\sigma), \tilde{k}_S(\sigma), \sigma\right)$$

and set $(\hat{k}_I, \hat{x}^\star, \hat{k}_S) = (\tilde{k}_I(\hat{\sigma}), \tilde{x}^\star(\hat{\sigma}), \tilde{k}_S(\hat{\sigma}))$.

Assume that $k_S = 0$ in Eq. (5.10), i.e. that social interactions do not have an effect on the dynamics of individuals' BMI. In this case, in order to estimate the remaining parameters, i.e. $(k_I, x^\star)$ (with $k_S = 0$ and σ undetermined), we regress $\hat{a}(x_i(t); \epsilon)$ on $-x_i(t)$ and a constant.

We note that the methodology presented in this section can only be applied to NHANES and NU BMI data, because these are the only data sets that have information on how individuals' BMI changes over time. We are able to compute $\Delta x_i(t)$ for individuals i in the NHANES data set with self-reported weights WHD010 (current, i.e. at time t_2) and WHD050 (one year prior to survey, i.e. at time t_1), and with self-reported height WHD020.

For convenience we set $t = 1999$ for the 1999-2000 NHANES survey, $t = 2001$ for the 2001-2002 NHANES survey, etc... We note that for NHANES data $\Delta t = 1$. For NU data we also consider $\Delta t = 1$, i.e., we consider individuals for whom we can compute BMI in two consecutive years. NU data for individuals in consecutive years exists for years

$$t \in \{1996, \ldots, 2013\}.$$

As above, if multiple weight measurements are present in year t we calculate the BMI for that year using the average weight in year t, whereas if multiple heights measurements are present then we calculate BMI using the average height (where the average is taken over all years). We note that for both data sets we use $\epsilon = 1$ to compute Eqs. (E.4a) and (E.4b).

All computations are performed in Matlab. Regressions are performed using the Matlab function *regress*. Optimization are performed using the Matlab function *fminsearch*.

E.2.5 Simulating Stochastic Differential Equation

We simulate Eq. (5.9) using the following forward Euler method with $\Delta t = 0.01$:

1. Initialize initial condition $x_i(t_1)$ from the data and set at time $t = t_1$.
2. For each individual i compute $a(x_i(t))$ and $b(x_i(t))$.
3. For each individual generate a random variable $\eta_i \sim \mathcal{N}(0, 1)$.
4. Set $x_i(t + \Delta t) = x_i(t) + a(x_i)\Delta t + b(x_i)\eta_i\sqrt{\Delta t}$ and set $t = t + \Delta t$.
5. Repeat steps 2-4 until $t = t_1 + 1 = t_2$.

E.3 Results

E.3.1 Population-level Model Implications

Empirical BMI distributions and fitted models

The following three figures, i.e. Figs. E.1-E.3, display the results from fitting empirical NHANES, NU, and BRFSS data to the nonsocial model distribution $p_{eq}^{(0)}(x)$, the social model distribution $p_{eq}(x)$, the log-normal distribution $f_{\log}(x)$, and the skew-normal distribution $f_{\text{skew}}(x)$.

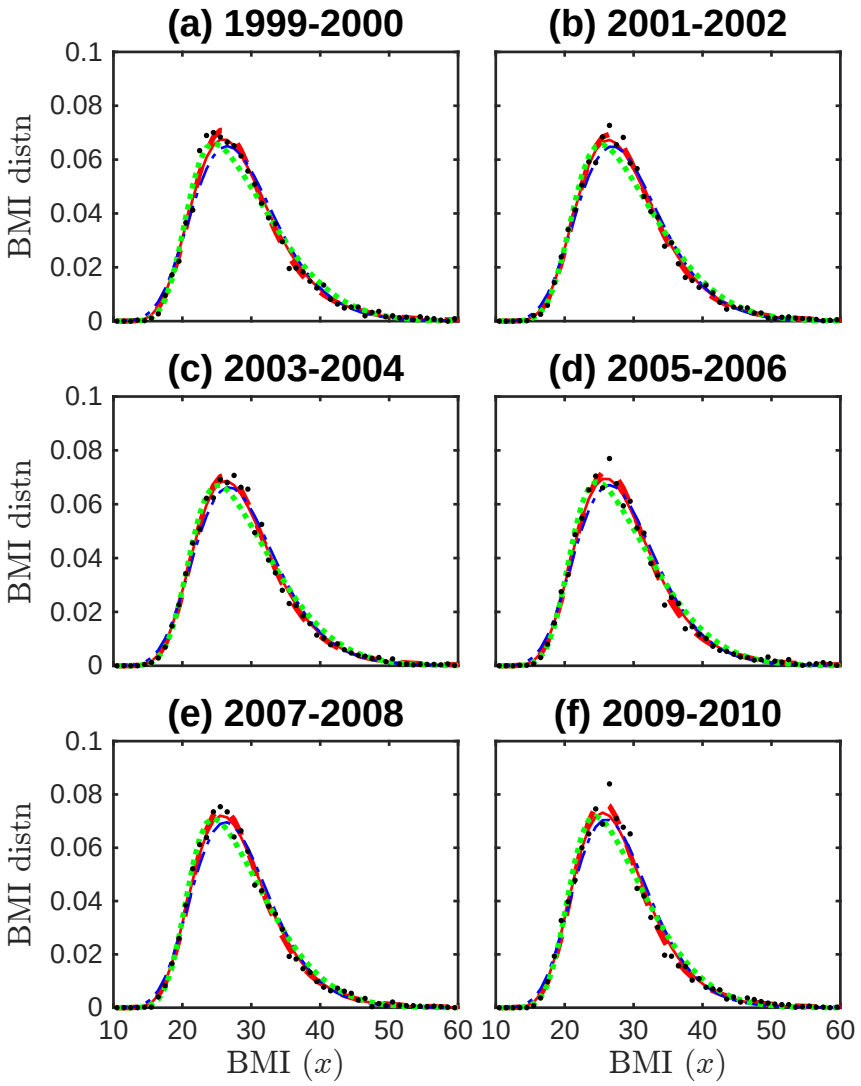

Figure E.1: Results from fitting NHANES BMI distributions (solid black dots) to four probability distribution functions: model without social utility (Eq. (5.16)—solid red), model with social utility (Eq. (E.3)—dashed red), log-normal null distribution (Eq. (5.17)—dash-dotted blue), and skew normal null distribution (Eq. (5.18)—dotted green). NHANES BMI distributions are calculated from NHANES survey years: (a) 1999-2000, (b) 2001-2002, (c) 2003-2004, (d) 2005-2006, (e) 2007-2008, and (f) 2009-2010. Results from 2011-2012 are shown in Fig. 5.1.

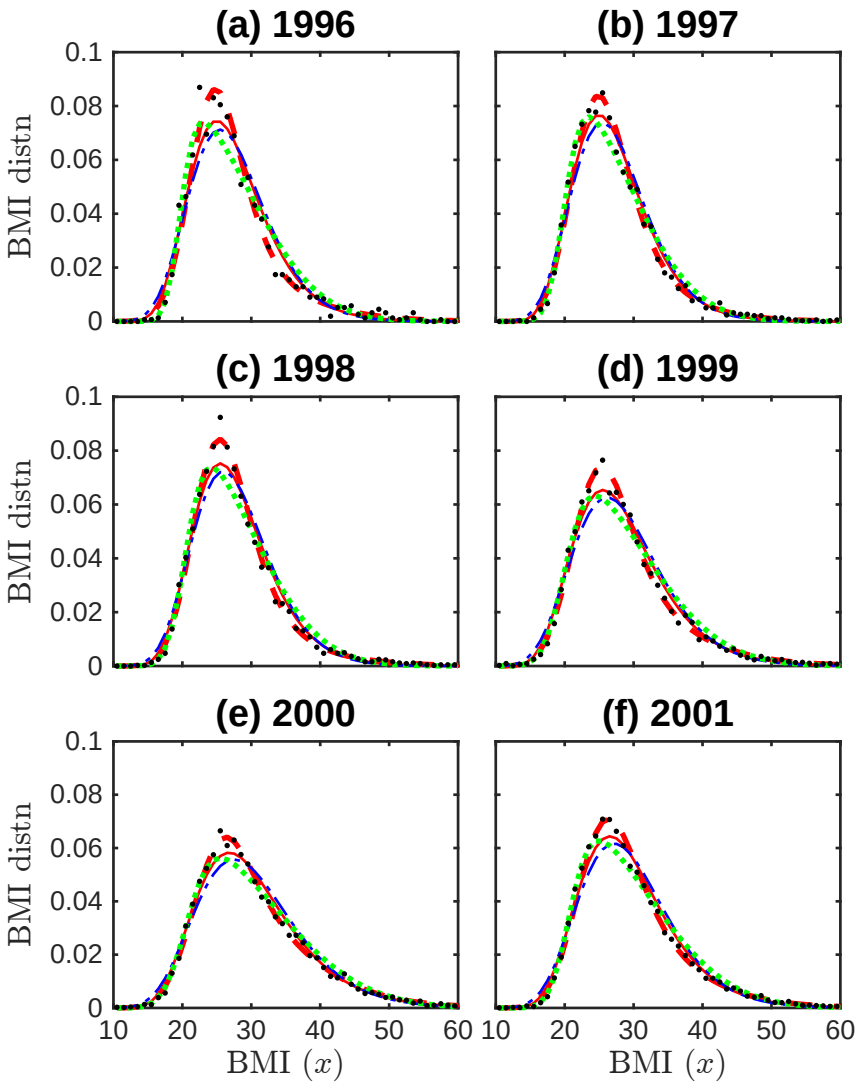

Figure E.2: Results from fitting NU BMI distributions (solid black dots) to four probability distribution functions: model without social utility (Eq. (5.16)—solid red), model with social utility (Eq. (E.3)—dashed red), log-normal null distribution (Eq. (5.17)—dash-dotted blue), and skew normal null distribution (Eq. (5.18)—dotted green). NU distributions are calculated for years (a) 1996, (b) 1997, ..., (o) 2010, (p) 2012, (q)2013, and (r) 2014. Results from 2011 are shown in Fig. 5.1.

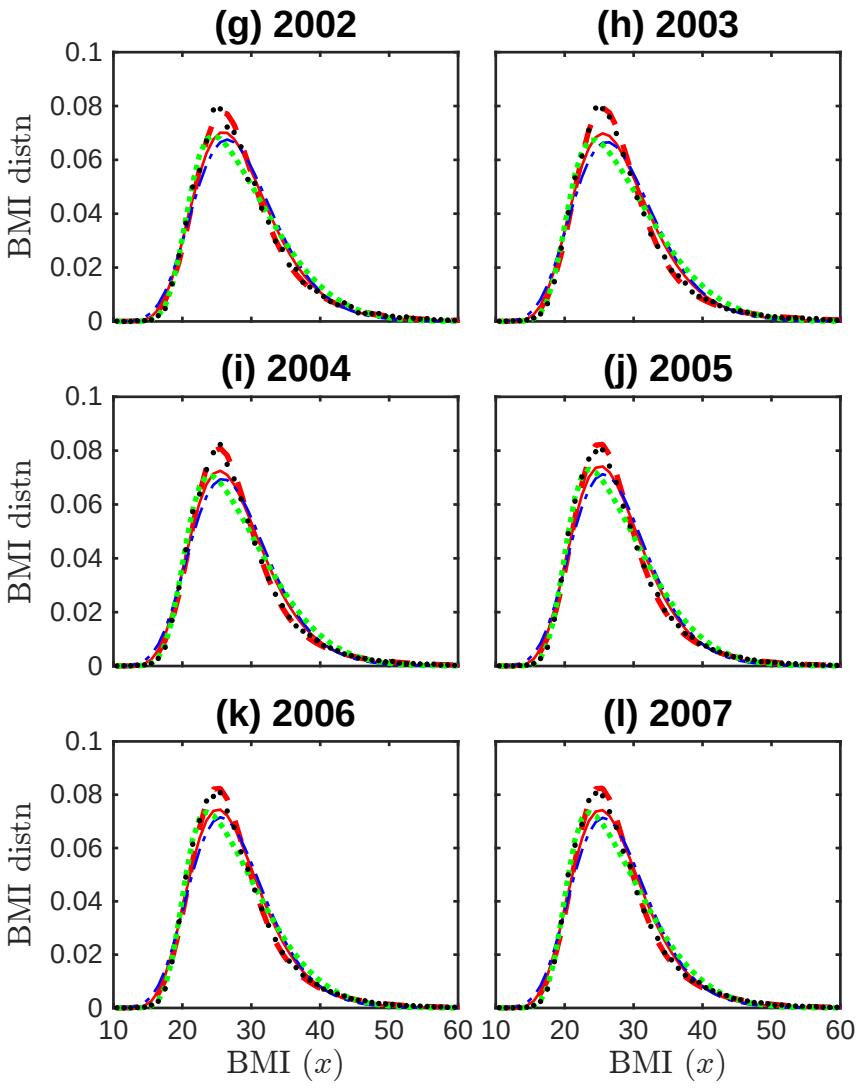

Figure E.2: Results from fitting NU BMI distributions (solid black dots) to four probability distribution functions: model without social utility (Eq. (5.16)—solid red), model with social utility (Eq. (E.3)—dashed red), log-normal null distribution (Eq. (5.17)—dash-dotted blue), and skew normal null distribution (Eq. (5.18)—dotted green). NU distributions are calculated for years (a) 1996, (b) 1997, ..., (o) 2010, (p) 2012, (q)2013, and (r) 2014. Results from 2011 are shown in Fig. 5.1.

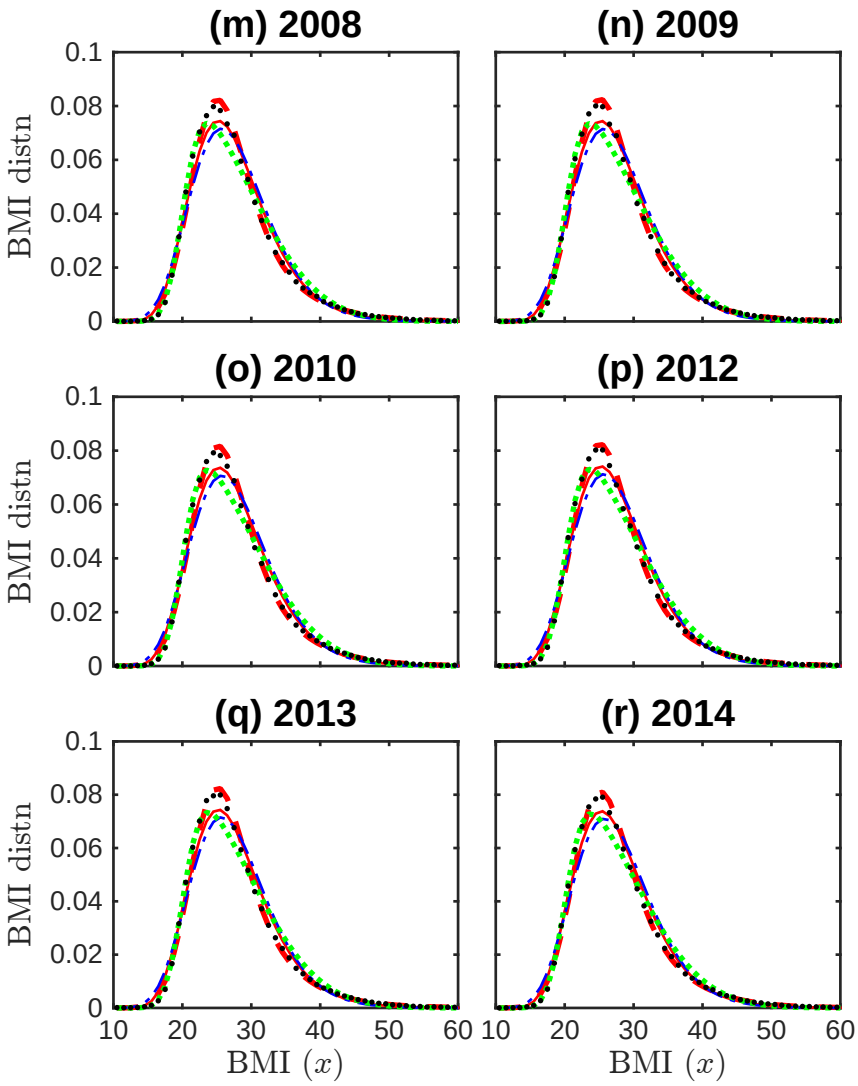

Figure E.2: Results from fitting NU BMI distributions (solid black dots) to four probability distribution functions: model without social utility (Eq. (5.16)—solid red), model with social utility (Eq. (E.3)—dashed red), log-normal null distribution (Eq. (5.17)—dash-dotted blue), and skew normal null distribution (Eq. (5.18)—dotted green). NU distributions are calculated for years (a) 1996, (b) 1997, ..., (o) 2010, (p) 2012, (q)2013, and (r) 2014. Results from 2011 are shown in Fig. 5.1.

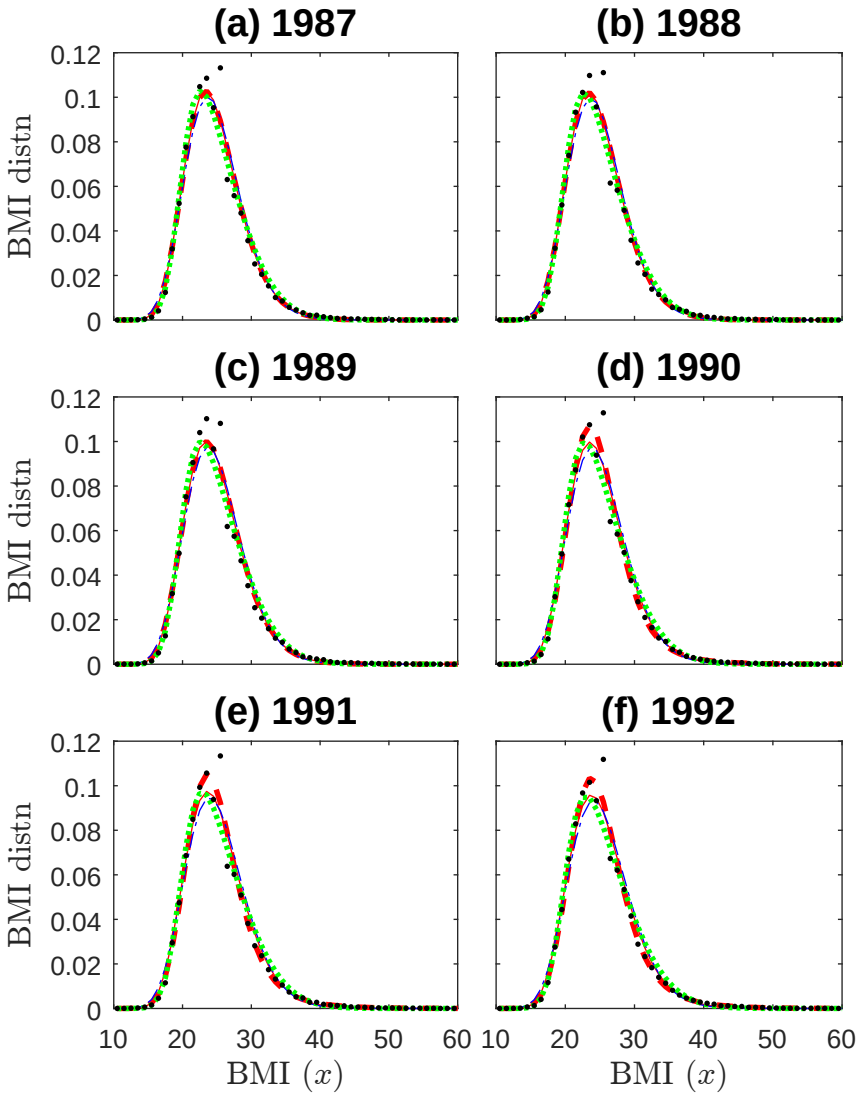

Figure E.3: Results from fitting BRFSS BMI distributions (solid black dots) to four probability distribution functions: model without social utility (Eq. (5.16)—solid red), model with social utility (Eq. (E.3)—dashed red), log-normal null distribution (Eq. (5.17)—dash-dotted blue), and skew normal null distribution (Eq. (5.18)—dotted green). BRFSS distributions are calculated for years (a) 1987, (b) 1988, ..., (x) 2010, (y) 2012, and (z) 2013.Results from 2011 are shown in Fig. 5.1.

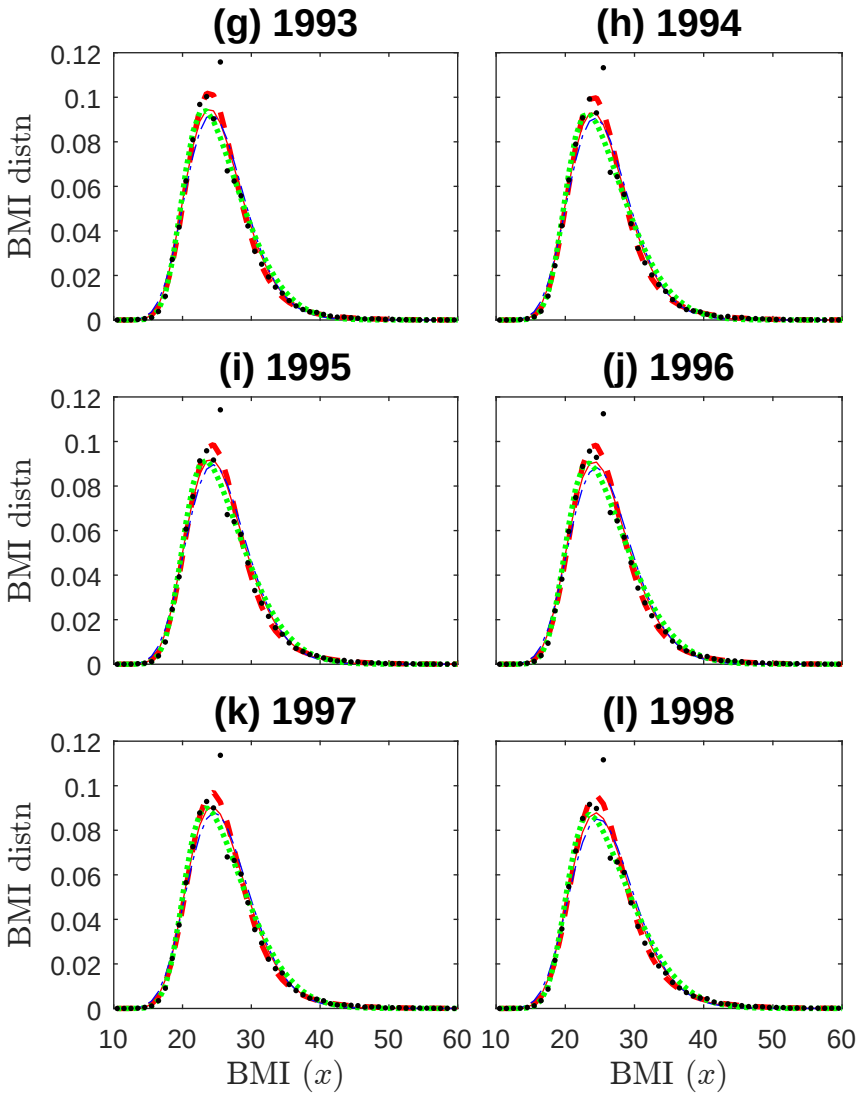

Figure E.3: Results from fitting BRFSS BMI distributions (solid black dots) to four probability distribution functions: model without social utility (Eq. (5.16)—solid red), model with social utility (Eq. (E.3)—dashed red), log-normal null distribution (Eq. (5.17)—dash-dotted blue), and skew normal null distribution (Eq. (5.18)—dotted green). BRFSS distributions are calculated for years (a) 1987, (b) 1988, ..., (x) 2010, (y) 2012, and (z) 2013.Results from 2011 are shown in Fig. 5.1.

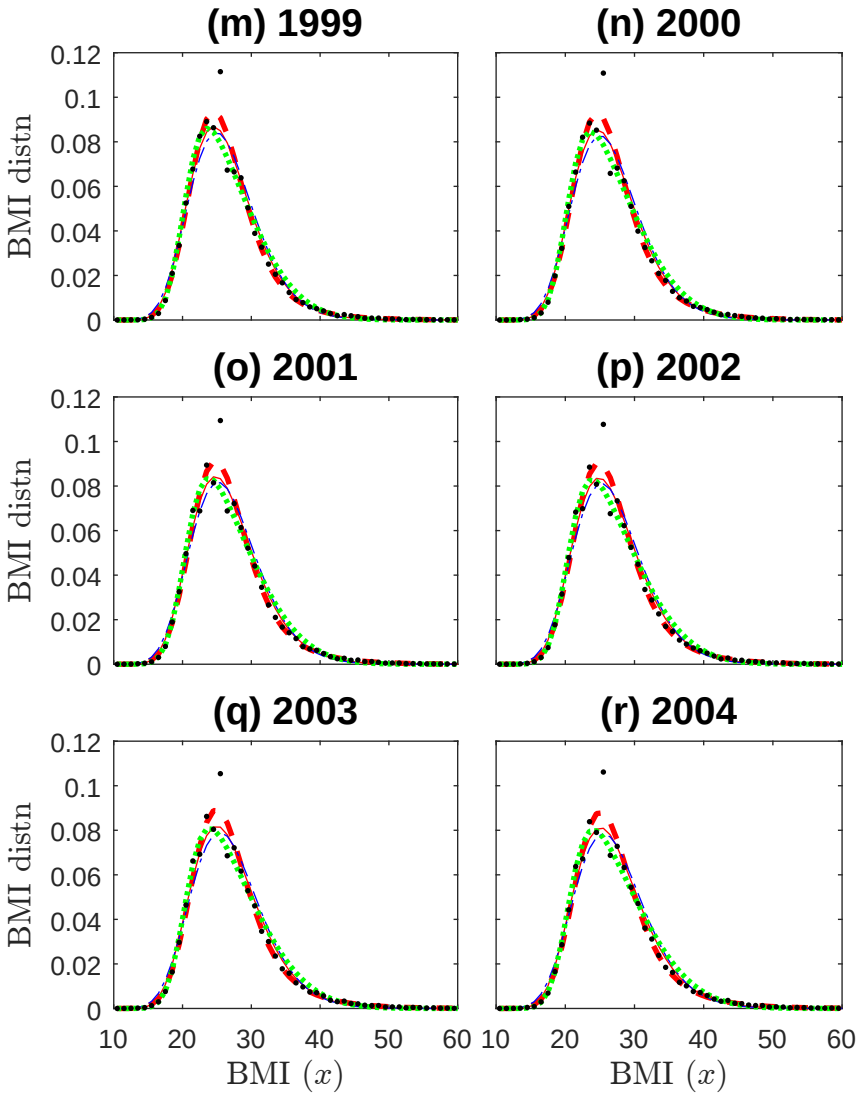

Figure E.3: Results from fitting BRFSS BMI distributions (solid black dots) to four probability distribution functions: model without social utility (Eq. (5.16)—solid red), model with social utility (Eq. (E.3)—dashed red), log-normal null distribution (Eq. (5.17)—dash-dotted blue), and skew normal null distribution (Eq. (5.18)—dotted green). BRFSS distributions are calculated for years (a) 1987, (b) 1988, ..., (x) 2010, (y) 2012, and (z) 2013.Results from 2011 are shown in Fig. 5.1.

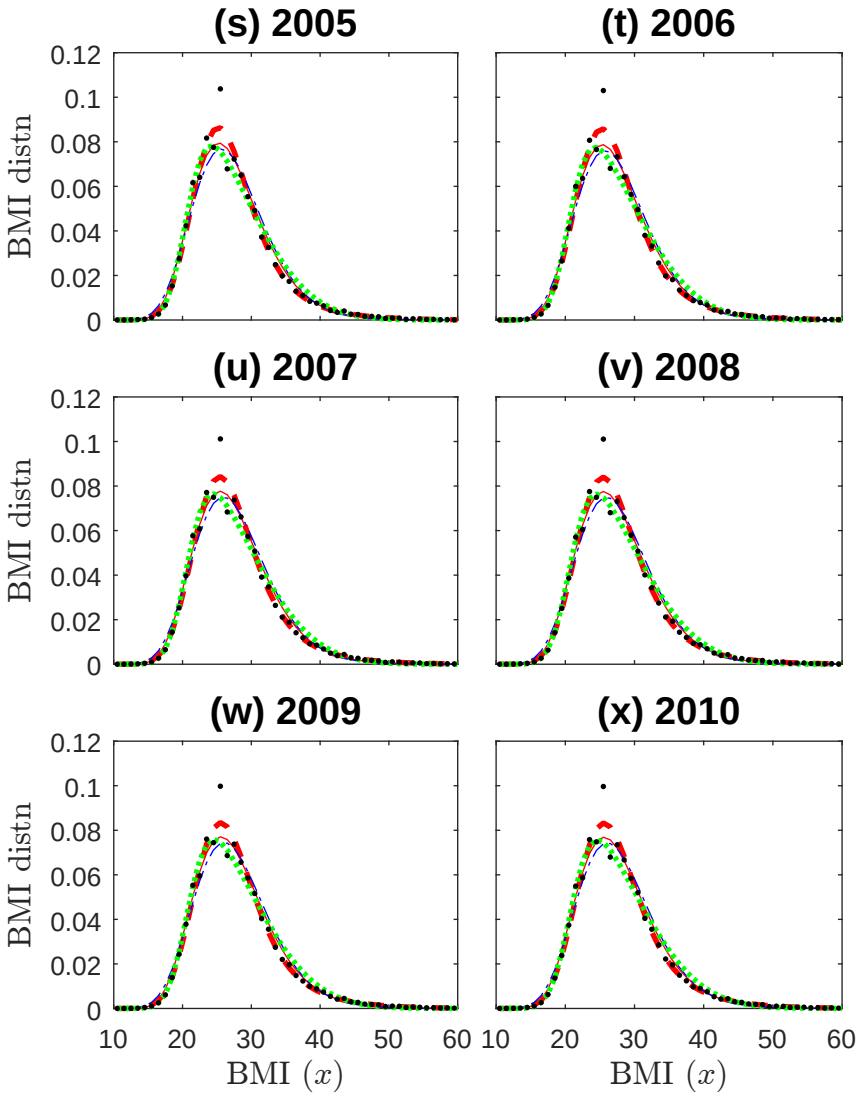

Figure E.3: Results from fitting BRFSS BMI distributions (solid black dots) to four probability distribution functions: model without social utility (Eq. (5.16)—solid red), model with social utility (Eq. (E.3)—dashed red), log-normal null distribution (Eq. (5.17)—dash-dotted blue), and skew normal null distribution (Eq. (5.18)—dotted green). BRFSS distributions are calculated for years (a) 1987, (b) 1988, ..., (x) 2010, (y) 2012, and (z) 2013.Results from 2011 are shown in Fig. 5.1.

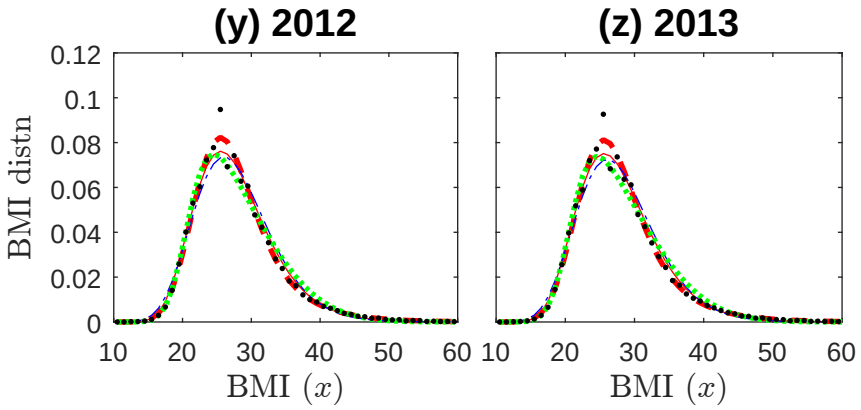

Figure E.3: Results from fitting BRFSS BMI distributions (solid black dots) to four probability distribution functions: model without social utility (Eq. (5.16)—solid red), model with social utility (Eq. (E.3)—dashed red), log-normal null distribution (Eq. (5.17)—dash-dotted blue), and skew normal null distribution (Eq. (5.18)—dotted green). BRFSS distributions are calculated for years (a) 1987, (b) 1988, ..., (x) 2010, (y) 2012, and (z) 2013.Results from 2011 are shown in Fig. 5.1.

Akaike Information Criterion

The following three tables, i.e. Tables E.1-E.3, display the results from computing the relative likelihood ratio for fitting empirical NHANES, NU, and BRFSS data to the nonsocial model distribution $p_{eq}^{(0)}(x)$, the social model distribution $p_{eq}(x)$, the log-normal distribution $f_{\log}(x)$, and the skew-normal distribution $f_{\text{skew}}(x)$.

Table E.1: Relative likelihood of non-social $p_{eq}^{(0)}(x)$, social $p_{eq}(x)$, log-normal $f_{\log}(x)$, and skew-normal $f_{\text{skew}}(x)$ models for NHANES empirical BMI distribution

	Relative Likelihood $\exp[(AIC_{\min} - AIC_i)/2]$			
Year	$p_{eq}^{(0)}(x)$	$p_{eq}(x)$	$f_{\log}(x)$	$f_{\text{skew}}(x)$
1999	2.8×10^{-5}	1	2.5×10^{-34}	7.5×10^{-23}
2001	2.9×10^{-4}	1	3.4×10^{-30}	8.5×10^{-32}
2003	3.1×10^{-5}	1	2.7×10^{-29}	1.2×10^{-27}
2005	1.7×10^{-5}	1	3.0×10^{-29}	2.1×10^{-19}
2007	1.0×10^{-2}	1	1.3×10^{-25}	7.7×10^{-13}
2009	1.3×10^{-5}	1	3.0×10^{-29}	6.2×10^{-21}
2011	1.3×10^{-2}	1	1.4×10^{-21}	1.1×10^{-9}

Table E.2: Relative likelihood of non-social $p_{eq}^{(0)}(x)$, social $p_{eq}(x)$, log-normal $f_{\log}(x)$, and skew-normal $f_{\text{skew}}(x)$ models for NU empirical BMI distribution

	Relative Likelihood $\exp[(AIC_{\min} - AIC_i)/2]$			
Year	$p_{eq}^{(0)}(x)$	$p_{eq}(x)$	$f_{\log}(x)$	$f_{\text{skew}}(x)$
1996	1.9×10^{-10}	1	1.7×10^{-24}	5.3×10^{-9}
1997	2.1×10^{-10}	1	7.5×10^{-35}	2.9×10^{-13}
1998	3.8×10^{-33}	1	5.0×10^{-56}	8.4×10^{-43}
1999	3.8×10^{-33}	1	6.0×10^{-84}	4.2×10^{-75}
2000	1.4×10^{-19}	1	1.4×10^{-77}	5.0×10^{-56}
2001	1.9×10^{-35}	1	1.8×10^{-136}	8.0×10^{-78}
2002	1.8×10^{-64}	1	9.9×10^{-202}	2.6×10^{-114}
2003 – 2014	$< 10^{-300}$	1	$< 10^{-300}$	$< 10^{-300}$

Table E.3: Relative likelihood of non-social $p_{eq}^{(0)}(x)$, social $p_{eq}(x)$, log-normal $f_{\log}(x)$, and skew-normal $f_{\text{skew}}(x)$ models for BRFSS empirical BMI distribution

	Relative Likelihood $\exp[(AIC_{\min} - AIC_i)/2]$			
Year	$p_{eq}^{(0)}(x)$	$p_{eq}(x)$	$f_{\log}(x)$	$f_{\text{skew}}(x)$
1987	9.8×10^{-138}	1	$< 10^{-300}$	2.5×10^{-228}
1988	4.1×10^{-135}	1	$< 10^{-300}$	5.9×10^{-242}
1989	4.3×10^{-243}	1	$< 10^{-300}$	$< 10^{-300}$
1990	2.3×10^{-241}	1	$< 10^{-300}$	$< 10^{-300}$
1991	8.6×10^{-279}	1	$< 10^{-300}$	$< 10^{-300}$
1992	1.0×10^{-299}	1	$< 10^{-300}$	$< 10^{-300}$
1993	$< 10^{-300}$	1	$< 10^{-300}$	$< 10^{-300}$
1994	1	0.14	$< 10^{-300}$	2.1×10^{-290}
1995	1	0.14	$< 10^{-300}$	$< 10^{-300}$
1996	1	0.14	$< 10^{-300}$	$< 10^{-300}$
1997 – 2013	$< 10^{-300}$	1	$< 10^{-300}$	$< 10^{-300}$

Nonsocial model fitted parameters

Below we plot the parameter estimates resulting from fitting $p_{eq}^{(0)}$ to empirical BMI data. It is interesting to note that not only do we observe in Fig. E.4 that the parameters k_I/k_b and $x^\star$ are indeed highly negatively correlated, but the relationship between k_I/k_b and $x^\star$ seems to be approximately linear. Indeed, it appears that the relationship between k_I/k_b and $x^\star$ could described by the same line of best fit in all three independently collected data sets. We emphasize that this relationship is not a consequence of the way in which we specified our mathematical model, but a feature of the data that our model has uncovered and that will need additional research in order to properly explain.

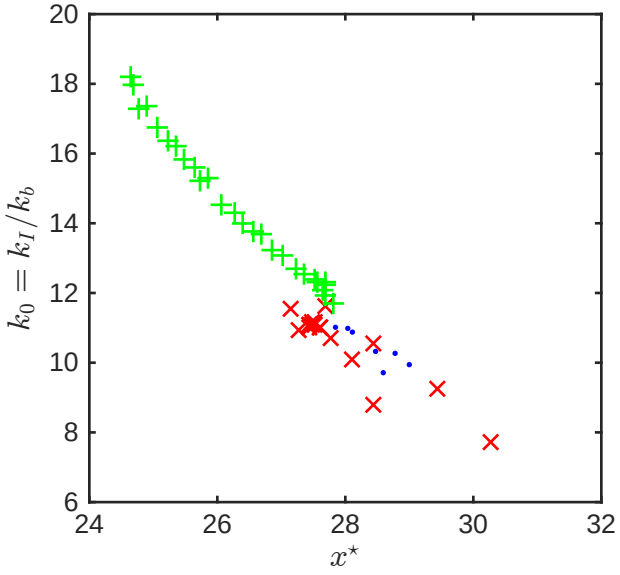

Figure E.4: Relationship between fitted parameters k_I/k_b and $x^\star$ of $p_{eq}^{(0)}(x; k_I/k_b, x^\star)$. Values of parameters k_I/k_b and $x^\star$ resulting from fitting $p_{eq}^{(0)}(x)$ (given in Eq. (5.16)) to empirical BMI distributions for each year of available data. (Blue dots) NHANES surveys; (red exes) NU data; (green crosses) BRFSS surveys. The parameters k_I/k_b and $x^\star$ are highly negatively correlated.

E.3.2 Individual-level Model Implications

The following two figures, i.e. Figs. E.5-E.6, display the results from computing the average and standard deviation in year-over-year change in individuals' BMI.

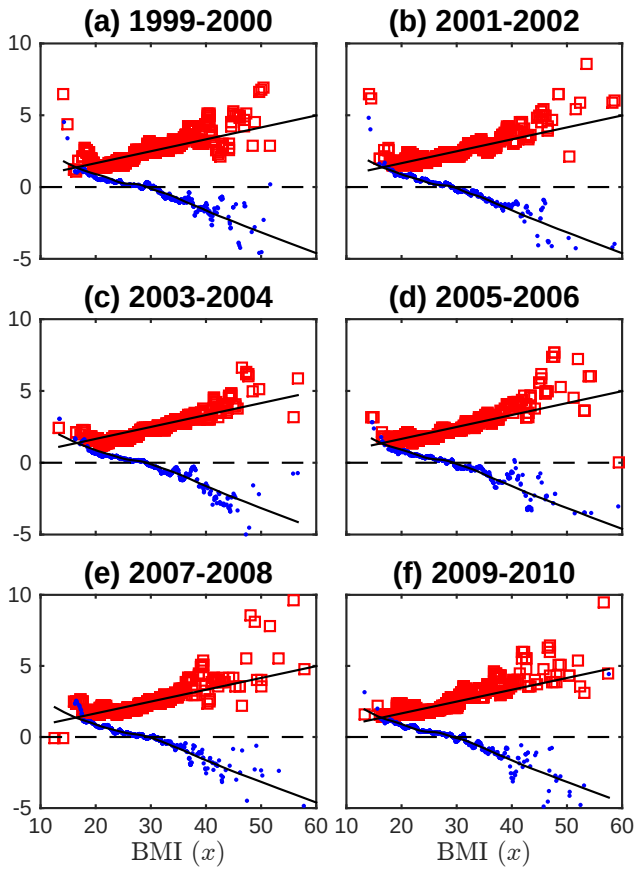

Figure E.5: Average and standard deviation in year-over-year change in individuals' BMI. The line of best fit for average year-over-year change in individuals' BMI (blue dots) is given by the regression specified in Eq. (5.19) (solid black line). The line of best fit for standard deviation for year-over-year in individuals' BMI (red squares) is given by the regression specified in Eq. (5.20) (solid black line). The line $y = 0$ is given by a dashed black line. NHANES survey years: (a) 1999-2000, (b) 2001-2002, (c) 2003-2004, (d) 2005-2006, (e) 2007-2008, and (f) 2009-2010. Results from 2011-2012 are shown in Fig. 5.3. Note: to improve visualization of data we display only 1,000 data points for each of $\hat{a}\,(x)$ and $\hat{b}\,(x)$, selected uniformly at random.

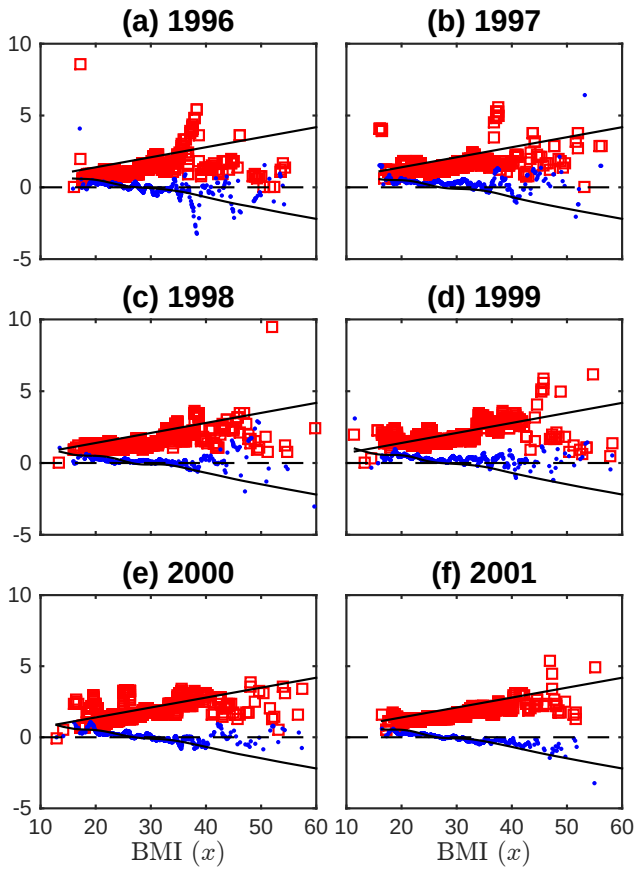

Figure E.6: Average and standard deviation in year-over-year change in individuals' BMI. The line of best fit for average year-over-year change in individuals' BMI (blue dots) is given by the regression specified in Eq. (5.19) (solid black line). The line of best fit for standard deviation for year-over-year in individuals' BMI (red squares) is given by the regression specified in Eq. (5.20) (solid black line). The line $y = 0$ is given by a dashed black line. NU data for years (a) 1996, (b) 1997, ..., (o) 2010, (p) 2011, (q)2012, and (r) 2013. Note: to improve visualization of data we display only 1,000 data points for each of $\hat{a}(x)$ and $\hat{b}(x)$, selected uniformly at random.

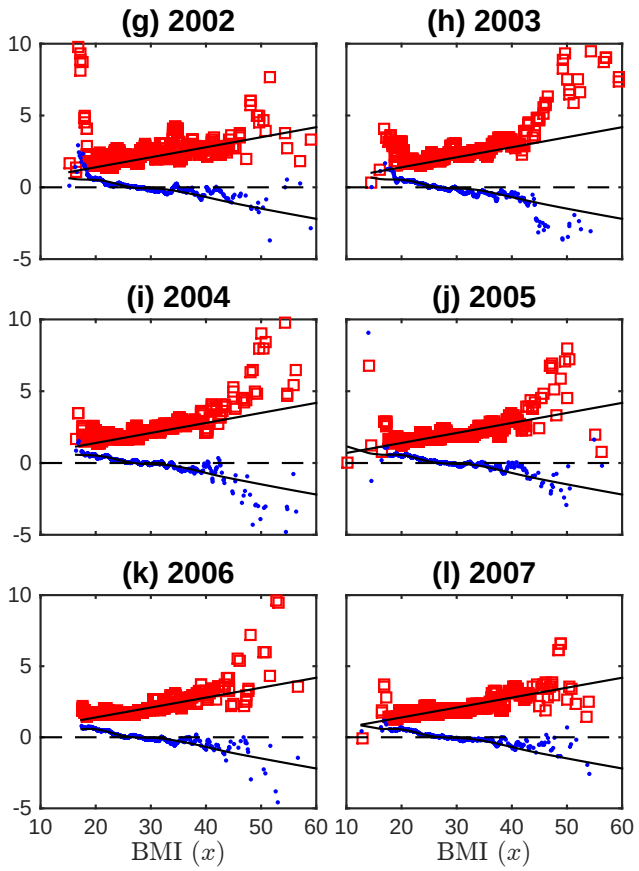

Figure E.6: Average and standard deviation in year-over-year change in individuals' BMI. The line of best fit for average year-over-year change in individuals' BMI (blue dots) is given by the regression specified in Eq. (5.19) (solid black line). The line of best fit for standard deviation for year-over-year in individuals' BMI (red squares) is given by the regression specified in Eq. (5.20) (solid black line). The line $y = 0$ is given by a dashed black line. NU data for years (a) 1996, (b) 1997, ..., (o) 2010, (p) 2011, (q)2012, and (r) 2013. Note: to improve visualization of data we display only 1,000 data points for each of $\hat{a}(x)$ and $\hat{b}(x)$, selected uniformly at random.

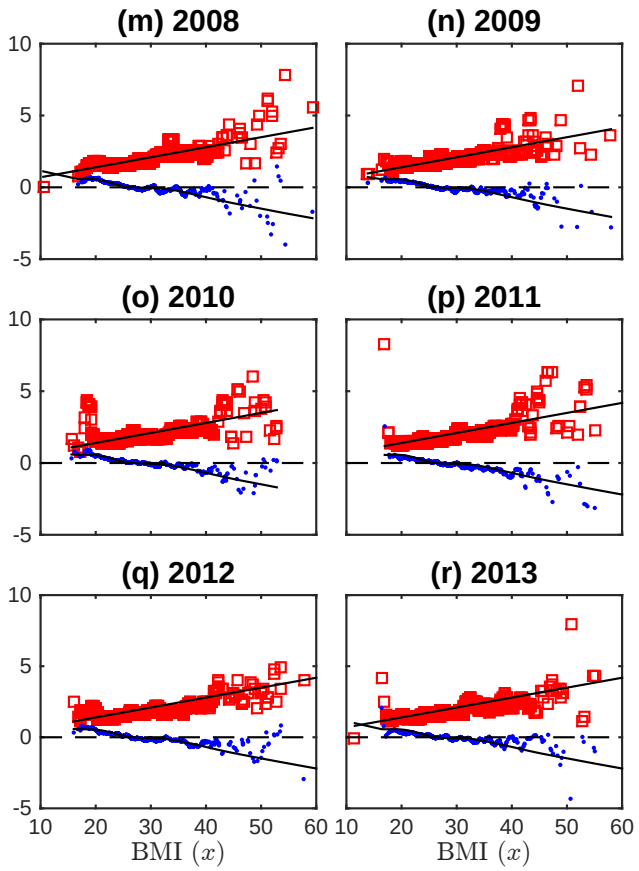

Figure E.6: Average and standard deviation in year-over-year change in individuals' BMI. The line of best fit for average year-over-year change in individuals' BMI (blue dots) is given by the regression specified in Eq. (5.19) (solid black line). The line of best fit for standard deviation for year-over-year in individuals' BMI (red squares) is given by the regression specified in Eq. (5.20) (solid black line). The line $y = 0$ is given by a dashed black line. NU data for years (a) 1996, (b) 1997, ..., (o) 2010, (p) 2011, (q)2012, and (r) 2013. Note: to improve visualization of data we display only 1,000 data points for each of $\hat{a}(x)$ and $\hat{b}(x)$, selected uniformly at random.