Dissemination models: applications in epidemiology, rumour spreading, wealth distribution, and opinion formation

Chan, K.M.D.

Summary

Dissemination models: applications in epidemiology, rumour spreading, wealth distribution and opinion formation

Over the last century, societies have become increasingly connected. This has brought substantial benefits, including online communication and low-cost transport by air and sea. However, greater connectivity also entails risks: pathogens and misinformation can spread more rapidly, potentially intensifying epidemic outbreaks and contributing to polarisation. To predict such dynamics and to design effective intervention strategies, it is crucial to understand the mechanisms that drive these spreading processes. This thesis contributes to such understanding through mathematical modelling. In particular, we examine how individual-level interaction rules and population structure translate into population-level outcomes.

This thesis studies dissemination models across epidemiology, rumour spreading, wealth distribution, and opinion formation, with an emphasis on long-run behaviour and qualitative regimes. Across these settings, we ask when large-scale spread occurs, which long-run behaviours are possible (e.g., extinction, coexistence and oscillations), and how modelling assumptions, such as interaction structure, heterogeneity, and the definition of the diffusing quantity, shape key outcomes. We address these questions using a combination of dynamical-systems methods, matrix/spectral analysis, and Markov-process techniques, supported by numerical computations.

In Chapter 2, we formulate a general class of deterministic compartmental models that allow freely chosen state transitions and state-dependent transition rates. This general set-up places canonical models in epidemiology (SIR) and rumour spreading (Daley-Kendall/Maki-Thompson), as well as their compartmental extensions, within a single unified framework. Consequently, our theoretical results apply directly across this range of models. We show that when the transition graph is acyclic (that is, transitions have a clear direction), the models are asymptotically stable. Under this acyclicity assumption, and with an additional weaker-than-acyclic condition on the associated dependency graph (which describes which contacts drive which transitions), we derive a classification rule. Starting from the sink node, the rule can be applied iteratively to determine which states will be unoccupied in the long run (vanishing states) and which will not. This classification in turn motivates the hypothesis that vanishing states decay at an exponential rate. When the conditions are violated, however, we demonstrate that vanishing states may decay more slowly than exponentially, and that the prescribed rule need not hold.

In Chapter 3, motivated by wealth distribution dynamics, we introduce a highly general stochastic dissemination model, which also generalises to applications such as opinion formation and file storage. The model is driven by two types of events: (i) units of wealth arrive externally to individuals, and (ii) units of wealth are redistributed among agents. To capture changing external conditions, we allow the background regime to be Markov-modulated. Despite the generality of this set-up, we show that the time-dependent joint probability generating function satisfies a system of coupled differential equations. Moreover, for the means and (co-)variances, this system reduces to *linear* differential equations. Through examples in wealth distribution, opinion formation, and file storage systems, we demonstrate how these equations provide analytical descriptions of the dynamics, which we validate with numerical computations.

In Chapter 4, we study group-structured SIR models. Throughout the chapter, we interpret the groups as age classes, although the framework also applies to other partitions, such as by occupation, behaviour, or geography. Although group-structured SIR models have been studied extensively, our focus is on a specific individual-level to population-level relation: how do interaction structures adapt to shifts in population structure, such as ageing, and how do these adaptations affect epidemic measures? To address this question, we propose a family of relationships between population structure and interaction structure, typically assuming reciprocal contacts. For most relationships, we derive analytical sensitivity results linking changes in population structure to qualitative changes in the basic reproduction number and the final size. In particular, under the assumption that older individuals are required to adjust their number of contacts in response to the contact capacity of younger family members and care professionals, we show that ageing leads to less severe epidemic outbreaks.

In the final chapter, Chapter 5, we explore the consequences of defining opinion relative to the population. This is motivated by the observation that the meaning of an individual being left-politically oriented changes if the population contains no right-oriented individuals. With this motivation in mind, we study *relative* opinions in the canonical DeGroot model and generalise it (in the spirit of Altafini) to allow individuals to repel or attract others' opinions. Such models typically exhibit consensus formation in the long run and, under additional modifications such as bounding the opinion space, may exhibit polarisation. We show that polarisation can also arise by incorporating *only* a relative notion of opinions, and we further show that long-run oscillations can arise as well. Moreover, we show that the model remains analytically tractable: the standard toolkit of matrix and spectral methods continues to apply to this set-up. To the best of our knowledge, this combination of long-run behaviours, including polarisation, consensus formation, and oscillations, together with this level of analytical tractability, has not been established before.

Overall, this thesis shows that the long-run behaviour of dissemination processes is often highly sensitive to details at the individual level. Across multiple application areas, we demonstrate that changes in interaction structure, population heterogeneity, or even the definition of the diffusing quantity can lead to qualitatively different outcomes. The key message is that context-specific assumptions are not a mere technicality; they are frequently decisive for behaviour at the population level.