



Universiteit
Leiden
The Netherlands

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

Chan, K.M.D.

Citation

Chan, K. M. D. (2026, September 16). *Dissemination models: applications in epidemiology, rumour spreading, wealth distribution, and opinion formation*. Retrieved from <https://hdl.handle.net/1887/4309746>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/4309746>

Note: To cite this publication please use the final published version (if applicable).

Chapter 1

Introduction

1.1 Spreading processes

A connected world

Over the last century, a wave of technological advances has shrunk the effective distance between people and reshaped society into a much more connected world. Where a letter from New York to Amsterdam once took weeks to arrive, a message now crosses the Atlantic almost instantaneously. A similar shift has taken place in transport: air travel has become increasingly commodified and global logistics by ship and plane have expanded significantly.

These changes have transformed how we interact and, crucially, how information and viruses spread. Much of communication now occurs digitally rather than face-to-face: people can easily find online communities with like-minded people, and react unfiltered while remaining anonymous. In parallel, vast quantities of goods, and the pathogens that may accompany them, move across continents day and night.

This increased openness has clear benefits, but it comes with risks: polarisation can intensify, and epidemic outbreaks can spread more rapidly. With this in mind, understanding spreading processes is even more essential; not only to predict them, but also to design interventions that prevent unwanted consequences. Throughout this thesis, we use *spreading* and *dissemination* interchangeably.

A toy example: the croissant rumour

Before we dive deeper into dissemination models, we clarify what we mean by a spreading process, see also Figure 1.1 for an illustration. Consider a simple example. On a Monday morning in your office, someone tells you: ‘There are free croissants in the kitchen.’ This piece of information (the diffusion object) moves through a population of colleagues (the hosts) via everyday communication channels such as hallway conversations and Teams messages (the interaction mechanism).

In this example, dissemination is driven by repeated contacts. Once someone has heard the message, they may pass it on in later interactions, so that the set of informed individuals grows over time. In this way, a simple local rule — information can be transmitted during a contact — can generate a spreading process in the population.

This example highlights the basic ingredients that recur throughout this thesis: a diffusion object, a population structure, and an interaction structure that specifies who can interact with whom and how often. Any model studied here makes explicit assumptions about these components, and uses them to describe how dissemination unfolds over time.




Figure 1.1: Toy example (croissant rumour). Shaded nodes have heard the message. Solid edges indicate hallway chats; dashed edges indicate Teams online contacts. Arrows show one possible sequence of transmissions.

Empirical and mathematical approaches

As mentioned earlier, a deeper understanding of spreading processes is essential; while also satisfying a genuine curiosity. Because dissemination arises in many settings, from sociology to epidemiology and beyond, research on it is scattered across disciplines. This diversity is reflected not only in application areas but also in methodology: mathematical studies seek mechanistic understanding through analytical and numerical methods, whereas empirical studies primarily draw on data, with the two approaches typically supporting one another.

To illustrate the value of both approaches, we consider a spreading process in an election-related situation. During an election, a politically relevant information may be spreading, and differences in receptiveness can shift what voters come to believe and, in turn, influence election outcomes. This can motivate a researcher to investigate how receptiveness among voters may influence election outcomes.

Addressing this study empirically provides among others a data-based approximation of voter characteristics. While a mathematical approach can provide insight in how receptiveness can influence the dissemination process in isolation, and allows for the exploration of what-if scenarios. The two approaches are complementary; the mathematical theory provides a structural grounding for the empirical work, while the empirical results can be used as a modelling assumption.

We also emphasise another benefit of the mathematical approach. Mathematical modelling forces us to state clearly what the essential mechanisms are. To answer a question, we typically seek the simplest model that retains the ingredients that matter. Deciding what can be ignored is seldom obvious at first, so models are often simplified through trial and error, a process that can yield new insight into the question at hand. A fitting metaphor is Picasso's *Le Taureau*: across successive sketches, detail is removed while the defining structure of the bull remains, much as in the construction of the simplest model

(Figure 1.2). Once the model is in place, mathematical tools allow us to analyse these mechanisms systematically and to compare the resulting dynamics across different fields of application. In short, the mathematical approach not only provides answers; it also clarifies the problem and can reveal common structure across disciplines.




Figure 1.2: A metaphor for the process of mathematical modelling: progressive abstraction retains only key features. Illustration inspired by Picasso's *Le Taureau* (1945–46), generated with ChatGPT.

Key questions and analytical perspectives

To clarify what mathematical modelling is used for, we highlight the typical questions that dissemination models are designed to answer. These questions can be grouped into three themes. First, there are *threshold questions*: under which conditions does a dissemination process lead to a substantial outbreak rather than an early fade-out, and how do these conditions depend on contact patterns and differences in intrinsic host traits? Second, there are *regime and long-run questions*: which qualitative behaviours are possible (e.g. extinction, coexistence, or oscillations), what determines their stability, and how sensitive are they to modelling assumptions? Third, there are *outcome questions*: how do modelling ingredients translate into population-level measures, such as outbreak size, the transmission burden carried by subgroups, or the emergence of polarisation? While these three themes cover much of the core agenda, other important questions concern the role of stochasticity and finite-population variability, inference from partial observations, and the design of interventions and control strategies.

Addressing these questions requires tools that bridge microscopic rules to macroscopic behaviour. In deterministic settings, analysis often proceeds by identifying equilibria and conserved quantities, deriving threshold conditions that separate qualitatively different regimes, and studying how outcomes change as assumptions or parameters are varied (dynamical-systems methods, matrix and spectral analysis). In stochastic settings, the focus is typically on finite-population uncertainty, for example by quantifying extinction probabilities, fluctuations around average behaviour, and the likelihood of rare but influential events (Markov-process techniques, branching approximations, scaling limits). When heterogeneity is naturally described by a continuous trait (such as age or spatial location), it is often important to track how population-level quantities vary over that trait through time using continuum descriptions (PDE and integro-differential equation methods). Finally, numerical computation and simulation are widely used to explore regimes beyond analytical tractability and to complement theoretical predictions; when models are confronted with data or used for intervention design, statistical and optimisation-based methods also play an important role (simulation, statistical inference, optimisation/control).

Outline of the thesis

In this thesis, using a *mathematical approach*, we develop and analyse models for dissemination processes across several application areas, including epidemiology, rumour spreading,

wealth distribution, and opinion formation. A central theme is how individual-level interaction rules and population structure shape population-level outcomes, in particular long-run behaviour and qualitative regimes.

Chapter 2 bridges disciplines by introducing a general class of dissemination models that unifies canonical models from epidemiology and rumour spreading, and many of their extensions. Within this framework, structural conditions are identified in terms of the underlying ‘graph’ that determines the long-term behaviour of the system’s dynamics.

Chapter 3 moves to a probabilistic setting and discusses a broad class of dissemination models with applications in, among others, wealth distribution and file-storage systems. We show that, although the model allows intrinsic randomness and substantial generality, it remains analytically tractable.

Chapter 4 returns to epidemiology and investigates how demographic change, such as ageing, can affect epidemic severity. Using a classical epidemiological framework (the multi-type SIR framework), we analyse how changes in population structure translate into changes in the interaction structure and, in turn, how these changes impact standard outbreak-severity measures.

Finally, Chapter 5 considers opinion formation. Instead of considering opinions as a stand-alone object, we approach them as a measure relative to the population. This perspective is, among others, motivated by the idea that political positions are interpreted in relation to the population (what is an individual with a ‘left’ ideology if the population does not contain individuals with a ‘right’ orientation?). Under this alternative approach, we formulate and analyse a model in which interactions can generate both attraction and repulsion, and show that adopting a relative notion of opinion can reveal dynamics that remained hidden under a non-relative approach.

The remainder of this chapter provides background and context for the main part of the thesis. We present a classification of dissemination models along four modelling dimensions: interaction structure, population structure, model representation, and diffusion object, which is used throughout to position both canonical models and the models considered in our research. We conclude with a chapter-by-chapter summary of contributions within this classification.

1.2 A classification of dissemination models

As introduced earlier, we present and discuss here a classification of dissemination models along four dimensions: interaction structure, population structure, model representation, and diffusion object. To guide the reader, we try to follow the same structure for each dimension: we first explain its meaning, then discuss a number of illustrative models, and finally reflect on the impact of the associated modelling choices. In selecting examples, we have aimed to cover the full range of the classification; however, to keep the discussion compact and accessible, several canonical models are used repeatedly to illustrate multiple dimensions.

1.2.1 Interaction structure

Uniform mixing vs. network-based interactions

By *interaction structure* we mean the pattern of contacts between hosts (e.g. individuals or mosquitoes): who meets whom, and how often. Its importance for dissemination is immediate, since some form of contact is necessary for transmission. These contact patterns, especially in large populations, can be highly complex and, if represented in full detail, quickly become mathematically intractable. As mentioned in Section 1.1, we therefore make well-considered modelling assumptions in order to arrive at a more tractable context. Next, we discuss several standard simplifications.

Perhaps the simplest and most widely adopted simplification is that of *uniform mixing*. Here, each individual is assumed to be equally likely to interact with any other, as illustrated in Figure 1.3a. In practice, uniform mixing can arise in settings where interactions are close to random; one example is a phone interaction between a customer and a call-centre agent. If customers are served on a first-come-first-served basis, then agent-customer interactions are approximately random.



Figure 1.3: Illustration of uniform mixing and network-based interactions. In uniform mixing, interaction partners are effectively random. In network-based interactions, possible interactions are defined by an explicit graph structure.

In most real-world situations, uniform mixing is an oversimplification; however, the assumption allows for the formulation of baseline models. These models, which can be interpreted as the simplest description of a dissemination process, provide a tractable starting point for mathematical analysis and serve as a foundation for subsequent work that relies on more realistic assumptions.

The interaction structure can also be simplified through milder assumptions, such as *network-based* interactions (see Figure 1.3b). Network structure is important, for example, in a community-structured population, where individuals interact more frequently with others within their community than with those outside it. Another typical situation where network structure plays an important role is in settings with fewer individuals: in a small company, if a rumour spreads, the question of who speaks to whom is central.

Uniform mixing: averaging and canonical models

Notationally, an interaction structure can be represented by a matrix whose (i, j) entry gives the per-unit-time rate at which individual- j interacts with individual- i . Under uniform mixing, where individuals are treated identically, the entries of this matrix are all

equal to the same constant. By contrast, in a network-based model with N individuals, one may in principle allow up to N^2 distinct interaction rates, reflecting population heterogeneity. This highlights the extent to which uniform mixing simplifies the interaction structure.

Reducing the interaction description to effectively a single contact rate can be interpreted as individuals *experiencing population averages*: conditional on making a contact, the type of the contacted host is distributed proportionally to the current host-type fractions.


We illustrate this with the canonical deterministic SIR model from epidemiology [90], which partitions the population into susceptible (S), infectious (I), and recovered (R) hosts. Under the uniform mixing assumption, an infectious individual contacts a susceptible with probability S/N . Consequently, since there are I infectious individuals, the rate of new infections at any given time is proportional to $I(S/N)$. This averaging principle is precisely what makes uniform mixing-based models tractable: it allows the dynamics to be expressed in terms of compartment sizes, leading to a low-dimensional system of ordinary differential equations.

Network-based interactions: local structure and canonical models


We now move to *network-based* interaction structure. Here, possible contacts are described by an explicit graph: nodes represent hosts, and edges represent potential interactions between them. To understand these graphs better, we introduce some basic notions. The *degree* of a node is the number of its direct neighbours. *Clustering* captures the extent to which the neighbours of a node are also connected to one another, reflecting the tendency of nodes to form tightly connected groups. Finally, the *distance* between two nodes is the number of edges in a shortest path connecting them; large distances indicate that the diffusion object must pass through many intermediaries to travel from one node to the other. Having established this terminology, we are now ready to discuss two important network-based structures: *small-world* networks [165] and *scale-free* networks [1, 15].

A typical attribute of a small-world network is that it couples high clustering with short distances. This combination arises when the network has strong local connectivity while also having some long-range edges. The presence of these long-range edges dramatically reduce distances between nodes (see Figure 1.4a). Heuristically, this captures the popular ‘six degrees of separation’ phenomenon: even in large populations, individuals are at most separated by six intermediary contacts [156]. Small-world networks are therefore a natural approximation of social networks [164]; clusters represent friend groups, while long-range edges provide shortcuts between communities.

Another important network structure are the scale-free networks; characterised by a highly skewed degree distribution (approximately following a power law). In such networks, most nodes have few edges, while a small number of nodes act as highly connected hubs, playing a disproportionate role in dissemination dynamics (see Figure 1.4b). Examples that have been reported as approximately scale-free networks include the World Wide Web [2] (nodes are websites and edges are hyperlinks) and scientific citation networks [136] (nodes are papers and citations are edges).



(a) Small-world network (highlighted nodes participate in long-range links).



(b) Scale-free network (highlighted nodes are hubs).

Figure 1.4: Small-world network and scale-free network.

Impact: interaction structure in the voter model

Next, we briefly illustrate the impact of the aforementioned interaction structures: uniform mixing, small-world and scale-free networks on the *voter* model [39, 82, 103]. Originating in opinion dynamics, the classical voter model describes hosts that carry a binary opinion: at random times, an individual adopts the opinion of an interaction partner, as determined by the underlying graph.

As a baseline, consider the voter model under uniform mixing. Under this assumption, the opinions eventually reach consensus: all hosts end up holding either opinion 0 or opinion 1 [104]. In contrast, on a small-world network, the mix of predominantly local connections and a few long-range edges can prevent complete ordering: opinions can coexist for a long time, with the time to consensus growing with the population size [29]. A different phenomenon arises on a scale-free network. Here, the voter model still reaches consensus; however, the eventual consensus opinion is largely determined by the initial opinions of the hubs [148].

These examples illustrate why interaction structure is one of the key modelling ingredients in the study of dissemination processes.

So far, we have discussed who meets whom, and how often. A second important modelling choice concerns *population structure*, as discussed in the next section.

1.2.2 Population structure

Homogeneous versus heterogeneous populations

An important feature of hosts is that, in general, they are not identical (see Figure 1.5): differences in intrinsic traits may affect their spreading behaviour. In epidemiology, for example, some hosts may be more susceptible to infection due to a weaker immune system. In rumour spreading, certain hosts may be less inclined to spread rumours on principle. Such differences between hosts can significantly affect dissemination dynamics. Populations in which hosts are assumed to differ in intrinsic traits are called *heterogeneous*, whereas populations in which hosts are assumed to have identical traits are referred to as *homogeneous*. Note that we treat ‘with whom’ an individual can have contact, and at what rate, as a feature of the *interaction structure*, rather than a population structure. This convention is a choice, and alternative classifications are defensible.



Figure 1.5: Population structure. (a) Homogeneous population: identical intrinsic traits. (b) Heterogeneous population: multiple host types, indicated by different symbols.

Truly homogeneous populations are rarely encountered in practice; however, examples of approximately homogeneous populations exist. For instance, students living in a residence hall who have similar ages and routines may reasonably be treated as homogeneous in an epidemiological context. Even when homogeneity is an oversimplification, it serves (just as uniform mixing in the interaction classification) as a starting point, supporting the formulation of baseline models.

Moving beyond homogeneity brings us to a variety of intrinsic host characteristics that can differ among hosts. An example in epidemiology is ‘partial immunity’, which follows from earlier infection or vaccination. In innovation diffusion, a varying trait among hosts is ‘willingness to adopt’, while in financial contagion ‘resilience’ (such as a liquidity buffer) is strongly dependent on the financial entity (such as banks, firms or households). We now examine various canonical models in which these intrinsic traits are assumed to be heterogeneous.

Illustrating population structure through canonical models

We briefly touch upon several canonical models from different fields to illustrate a range of population structures. Here, we describe the models and defer discussion of how population structure affects dissemination dynamics to the subsection below.

We begin by introducing a pair of baseline models in rumour spreading, the Daley-Kendall (DK) and Maki-Thompson (MT) models [44, 110]. These models assume a *homogeneous* host population. They let each host be in one of the following states: ignorant, spreader, and stifler. Under uniform mixing, ignorants turn into spreaders upon contact with a spreader. The distinguishing feature of these models, compared with the earlier mentioned SIR framework, is that spreaders stop spreading only through interaction: a spreader becomes a stifler after meeting another spreader or stifler, which reflects loss of interest once the rumour is perceived to be widely known. Another baseline model that assumes homogeneity is the canonical SIR model in epidemiology, which we described in Chapter 1.2.1.

Turning to heterogeneity-based models, we discuss an age-structured extension of the deterministic SIR model introduced earlier [6, 51]. Building on the standard SIR mechanism, the population is partitioned into age groups (a continuous age structure is also common), and hosts’ transition parameters are allowed to depend on age.

Further, in innovation diffusion, the canonical adoption threshold model (here we consider

the Granovetter version [73]) equips each potential adopter with an intrinsic threshold: adoption occurs only when the overall adoption rate, interpreted as the level of social exposure, exceeds that individual threshold. This setup provides a mechanism through which individual-level willingness to adopt drives the spread of ‘innovation’.

Finally, another canonical model we consider is the payment-clearing model of Eisenberg and Noe [59]. This financial contagion model describes entities’ ability to meet their obligations in full. Entities may differ in their assets (e.g. external assets or liquidity buffers) and liabilities. In this setting, the default of one entity can trigger a cascade of defaults.

Impact: population structure in canonical models

Here, we return to the models discussed in the previous subsection and shed light on how their underlying population structure influences the resulting dynamics.

Under the deterministic SIR framework, heterogeneity in comparison to homogeneity [24, 54], can reveal phenomena such as concentrated infections in hosts with certain intrinsic traits [168], and prolonged outbreak durations [87]. Similarly, in the DK and MT models, letting hosts differ in the extent of scepticism or enthusiasm, can reveal that a subset of highly receptive spreaders can account for most transmission.

The influence of heterogeneity is especially clear in the adoption threshold model (in innovation diffusion). If all hosts shared a single adoption threshold, adoption would be essentially an all-or-nothing mechanism. Specific distributions of thresholds, on the contrary, can generate gradual take-off, tipping phenomena, and cascade-like adoption.

Finally, in the Eisenberg-Noe payment-clearing framework, heterogeneity in balance-sheet resilience is central to whether cascades can occur at all. A model constructed only around an ‘average’ entity may suggest stability of a financial system, whereas a heterogeneous model can reveal systemic tail risk: a small set of financial entities may default under stress, transmit losses to counterparties, and thereby induce further defaults.

Thus far, we have varied who interacts with whom and how hosts differ intrinsically. We now turn to a third modelling dimension: *model representation*. The same dissemination mechanism may be modelled deterministically or stochastically, and studied either analytically or through simulation-based approaches.

1.2.3 Model representation

In this section, we discuss two important choices of model representation: deterministic versus stochastic formulations, and the intended mode of analysis, namely analytical versus computational approaches.

Deterministic and stochastic modelling

A central modelling choice is how to account for the stochastic nature of a dissemination process. In practice, randomness can enter at a variety of levels; the following sources serve as illustrations.

One source is the interaction structure: even under uniform mixing, where each host has the same probability of contacting any other host, there remains a positive probability of rare events, such as contacting the same host multiple times consecutively. Another source of randomness is the timing of contacts: it is possible, though unlikely, that no contact occurs during the next hour even if the *average* inter-contact time is ten minutes. Stochasticity can also enter through intrinsic traits: for instance, an individual who is *on average* stubborn may occasionally act in a non-stubborn manner. Another source is randomness in the diffusion object: the rumour being spread may usually remain close to its original form, but in rare circumstances — for example, when an individual is in an exceptionally good mood — it may shift to a more sensational variant. Finally, randomness may enter a spreading process through the *outcome* of a contact and through random durations, such as the time an individual remains infectious.

Regarding stochasticity, a modeller faces a trade-off between a more detailed stochastic description and, at the other extreme, a more simplified deterministic model. In settings where randomness plays a key role in the dissemination mechanism, a stochastic formulation is essential: contacts occur at random times, interaction partners are selected probabilistically (even under uniform mixing), and the outcome of a contact may be uncertain. Here, randomness is not an add-on feature but a crucial part of the mechanism one aims to model.

At the same time, adopting a deterministic description greatly simplifies the model and therefore facilitates the formulation of baseline models. Analogous to assumptions such as uniform mixing and homogeneous population, a deterministic representation replaces these random, event-level ingredients by their average counterparts, yielding a more tractable description.

Connecting stochastic and deterministic models. Building on this trade-off, there is a natural bridge from stochastic to deterministic representations once we move from modelling individual events to population-level proportions. A common stochastic formulation is a continuous-time Markov chain in which events (infection, adoption, default) occur at random times, often driven by independent Poisson ‘clocks’. As the population size N increases, the number of events over any fixed time interval typically grows, and random fluctuations become relatively small compared with the overall scale. In this sense, the deterministic approximation can be viewed as a law-of-large-numbers limit: the stochastic dynamics concentrate around a non-random trajectory as $N \rightarrow \infty$.

This intuition can be formalised by the classical limit results of Kurtz [95, 96]: after rescaling the state by N so that we track proportions rather than counts, and letting $N \rightarrow \infty$, the stochastic process converges to a deterministic trajectory given by an ODE. One can also go a step back and capture finite- N variability around this limit: under appropriate conditions, the centred and rescaled fluctuations admit a Gaussian approximation (by Central limit theorem), providing a more refined description. The deterministic limit describes the typical, average evolution, but it does not by itself capture finite-population variability or rare events that remain visible in the stochastic model, such as early extinction in an epidemic context.

Deterministic vs. stochastic modelling in canonical models With this stochastic-deterministic link in mind, we revisit several canonical dissemination models through both lenses. Important deterministic examples are the classical SIR model, as well as deterministic variants

of the DK and MT rumour models, which were introduced in Chapters 1.2.1 and 1.2.2. Their stochastic counterparts are important as well, not only because they capture finite-population variability, but also because they are classical illustrations of the bridge from stochastic to deterministic descriptions, as discussed earlier. In particular, the stochastic SIR model [7] can be formulated as a continuous-time Markov chain in which infection events occur when infectious and susceptible individuals meet, while recoveries occur at Poisson rates for infectious individuals. The stochastic versions of the DK and MT models are formulated in a closely analogous way. It is worth noting that, historically, the DK and MT models were introduced in their stochastic form. Another canonical example with inherent stochastic feature is the voter model (introduced in Chapter 1.2.1), here, stochasticity arises in the selection of interaction partners and the resulting opinion updates.

Consequences of stochastic vs. deterministic modelling We revisit the canonical models in the previous paragraph and discuss the consequences of including stochasticity in their formulation. First, in a stochastic SIR model, even if an infectious individual infects on average two new susceptibles, a material probability exists that the first infected individual recovers before contacting anyone else; an ‘outbreak’ never occurs. This early-extinction effect cannot be replicated in a deterministic model. Similarly, consider a stochastic DK or MT rumour model on a scale-free network (see Chapter 1.2.1) with hubs that have far above-average contact rates. In this model, by randomness a rumour may never reach any hub, leading to a much lower dispersion speed of the rumour than a deterministic approximation would suggest. Finally, a deterministic voter model amounts to tracking average opinion densities, which suggests coexistence of opinions 0 and 1. In the original voter model, by contrast, consensus is always reached: all hosts end up holding opinion 0 or opinion 1; a qualitative feature that a deterministic description cannot reproduce.

Analytical and computational approaches

Another axis in the model representation classification concerns the intended approach of analysis: at one extreme, a model is formulated to be studied analytically, while at the other extreme it is formulated to be investigated through computation and simulation. In general, an analytical approach offers greater mathematical tractability, but typically requires stronger simplifying assumptions. A computation-based approach, by contrast, can typically incorporate substantially more details (such as highly heterogeneous agents and explicit network-based interaction structures) at the cost of reduced analytical tractability and fewer closed-form results. In practice, both approaches are valuable: they often provide complementary insights into the same question, but from different angles.

In the following, we discuss both approaches separately, including canonical models that adopt each approach and examples that illustrate the impact of choosing one approach over the other.

Analytical approach: canonical examples Within an analytical approach, common techniques to study deterministic models are ODE theory and matrix theory, whereas stochastic models are commonly studied through the theory of Markov processes (e.g. continuous-time Markov chains) and related probabilistic tools. Canonical models that are set up for analytical study include many of the baseline models discussed earlier in the Chapter 1.2.1 and 1.2.2, such as the SIR model, DK and MT models, the voter model and the adoption threshold model. The impact of these models is clear: they serve as foundations

in their respective fields and successfully capture key dynamics in a simple, transparent formulation.

Computational-based approach and (canonical) agent-based models On the computation side, typical analysis techniques include numerical integration of ODEs, agent-based simulation, and event-driven stochastic simulation (e.g. simulating an underlying jump process by sampling successive event times and event types). While there are many dissemination models that are primarily studied through computation, we focus here on canonical models in the class of agent-based models (ABMs). In brief, ABMs specify traits and updating rules at the individual level, and then study the resulting population-level behaviour by simulation. We restrict attention to ABMs both to keep the discussion compact and because, in our view, they provide a clear illustration of what the computational/simulation-based approach can encapsulate.

Building on this, let us discuss the agent-based variant of the earlier described (see Chapter 1.2.1) SIR model. In this variant, the disease state is host-specific; contacts are simulated explicitly, and transmission and recovery occur as individual-level events. This setup naturally supports additional extensions such as the inclusion of individual disease histories, time-varying behaviour, and explicit locations or settings (e.g. households and workplaces). A second important, classical, ABM is Axelrod’s cultural dissemination model [11]. There, each agent is represented by a vector of cultural features; agents interact locally with a probability that increases with cultural similarity. Given an interaction, the host copies one trait from the contacted neighbour. The final example we consider is the Granovetter-type adoption threshold model (see Chapter 1.2.2) but approached as an agent-based model [32, 166]. In this setting, each agent is still assigned an intrinsic adoption threshold and adopts when social exposure exceeds it, however, here, each agent is kept track of individually.

Regarding impact of agent-based modelling, a general advantage of the approach is that it can include detailed assumptions on individual-level, and can reveal dynamics that remain hidden when the model is formulated through population-level equations. For example, in the agent-based SIR model, individual-level simulation can produce phenomena such as stochastic fade-out, clustered outbreaks driven by local contact structure, and ‘superspreading’ dynamics; effects that are typically averaged out in deterministic ODE descriptions. Further, in Axelrod’s cultural dissemination model, individual-level modelling makes it possible for early random encounters to steer the system towards entirely different macroscopic configurations, even under identical parameter choices. Likewise, in a Granovetter-type adoption threshold innovation diffusion model, an agent-based description allows, for example, the emergence of multi-wave adoption when awareness spreads locally and gradually; a feature that is typically collapsed into averaged response curves in analytical treatments.

Without going into detail, note that another axis in the modelling representation classification concerns the time description: models may be formulated in continuous time (with events occurring at random or deterministic rates) or in discrete time (with updates applied at fixed time steps).

1.2.4 Diffusion object

We now turn to the state space of the diffusion object (see Figure 1.6). For some diffusion objects the state space is naturally discrete, for instance, an individual either carries a virus or does not. For other objects, such as opinions, a binary description can be too coarse; a better representation of the state space may be a continuous interval.¹




Figure 1.6: A variety of diffusion-object state spaces. Left: a discrete state space (three possible states, as in the SIR model). Middle: a continuous, bounded state space (here taken to be values in $[0, 1]$, as in opinion formation). Right: a two-dimensional trait space. For illustration, an individual's current state is indicated by a shaded node.

Once one has decided between a discrete or continuous state space, a further choice is whether the diffusion object should be modelled as bounded or unbounded. For example, if the diffusion object is wealth, it is reasonable to assume that there is no hard upper bound: individuals or institutions can, in principle, always become wealthier. For ideological diffusion (e.g. state space being the progressive-conservative axis), one may similarly ask whether there are limits; can agents always become ‘even more’ progressive or conservative? A question we do not aim to answer here, but which affects modelling choices. Finally, the diffusion object need not be one-dimensional: what spreads may involve multiple components, such as belief in a rumour X alongside an opinion Y .

Against this background, we highlight the range of diffusion objects through a few canonical models. In opinion formation, a canonical model that tracks a continuous and bounded opinion space is the DeGroot model [48]. There, opinions take values in an interval such as $[0, 1]$; each host updates by adopting a weighted average of the opinions of the individuals it interacts with. In contrast, models such as the voter model and the SIR model (see Chapter 1.2.1) use diffusion objects that are discrete and bounded, with individuals switching between finitely many states. Unbounded state spaces arise for example in wealth distribution models [56]. Using a birth-death-type formulation, wealth units arrive to agents and are lost or transferred through interactions; over time, some agents may end up with nothing, while others accumulate a large share of the total wealth. Finally, the standard example of a diffusion object being multidimensional, is Axelrod’s model of cultural dissemination: here each agent carries cultural features such as music taste and religion.

¹When an additional continuous structure is tracked (such as age or spatial location), the finite set of coupled ODEs in compartment models (such as the SIR) is naturally replaced by continuum descriptions (in SIR setting: $S(a, t), I(a, t), R(a, t)$), which are often formulated as PDEs or integro-differential equations (as discussed in Chapter 1.1).

1.3 Contributions

Now that we have established a classification of dissemination models, we position the model discussed in each chapter within this framework and describe the contributions of our work in relation to existing literature.

Chapter 2: Graph conditions and asymptotic behaviour

In this chapter, we develop analytical theory for a class of deterministic compartmental models. The class allows for general interaction and population structures, and includes baseline models such as the SIR model and the DK/MT rumour models, as well as many of their extensions. It therefore unifies and generalises compartmental modelling across several fields. Our main result is a graph-theoretic condition that characterises the long-run behaviour of models in this class. To obtain this result, we consider not only standard transition graphs, but also introduce a *dependency graph*, which captures which contacts generate which types of transitions. We illustrate the analytical results with examples, including a pathogen-rumour setting and a multi-rumour setting. This chapter draws on methods from dynamical systems and graph theory. Due to the general set-up, the diffusion object can be interpreted as discrete, bounded, and possibly multi-dimensional.

Source note. This chapter is based on the preprint ‘Asymptotic behavior for a general class of spreading models’ [34], currently under peer review.

Chapter 3: Stochastic dissemination

Chapter 3 shifts focus from a class of deterministic models to a class of stochastic models. This general family of models allows for network-based interaction structures and heterogeneous populations. We introduce the model through a wealth distribution interpretation, where the diffusion object is unbounded (units of wealth), but it can also be interpreted as a model for other processes, such as opinion formation or file storage. This interpretive flexibility follows from the general set-up, which incorporates three key mechanisms: wealth can arrive from external sources, it can be redistributed stochastically through interaction events, and the overall dynamics may be Markov-modulated to reflect changing environments or regimes.

The contribution is primarily analytical and draws on Markov-process techniques, in particular probability generating function (PGF) analysis. We show that the joint transient distribution of agents’ wealth can be characterised by a system of coupled differential equations, and that these equations simplify substantially when considering moments: the corresponding equations for the means and (co-)variances become linear. Given the generality of the underlying stochastic model, it was unexpected that such tractable characterisations remain possible.

Source note. This chapter draws on the article ‘A versatile stochastic dissemination model’ [36].

Chapter 4: Demography, contacts, and severity

In Chapter 4, we study how demographic change such as ageing can affect epidemic outbreaks within a deterministic multi-type SIR framework. Here, demographic change is treated as a change in heterogeneous population structure that induces changes in the network-based interaction structure. We investigate analytically, using dynamical-systems methods and matrix and spectral analysis, how changes in population composition translate into shifts in contact patterns and how these, in turn, affect standard outbreak-severity measures such as the basic reproduction number and the final size. Our contribution is threefold. First, we set up a framework in which the contact matrix is an explicit function of demography. This provides an analytically tractable mapping from population structure to contact structure and, in turn, to epidemic severity measures. Second, we introduce and organise several families of demography-dependent contact patterns, capturing multiple mechanisms by which contact rates may adjust when population size or composition changes. Third, building on this framework and catalogue, we derive qualitative sensitivity results describing how demographic perturbations affect the severity of epidemic outbreaks.

Source note. This chapter draws on the unpublished 2016 manuscript *The impact of population ageing on the severity of an epidemic outbreak*, co-authored by K.M.D. Chan, H. Nishiura, O. Diekmann, and M.C.J. Bootsma. The manuscript itself builds on earlier work [33], and a brief extract was incorporated in [24].

Chapter 5: Relative opinion dynamics

In the final chapter, we move to the field of opinion formation. We wonder whether opinions (the diffusion object) should be interpreted relative to the population and study analytically what implications this has on long-term behaviour. This viewpoint is motivated by the observation that labels such as ‘political right’ or ‘political left’ only have meaning in relation to other positions present in society. We consider a deterministic model in which agents can both *pull* and *push* each other’s opinions: interactions may lead agents to move closer (persuasion), move away (repulsion), or remain unchanged (e.g. distrust). The heterogeneous and network-based model generalises the DeGroot framework by allowing influence to have both attractive and repulsive components. Our main contribution is — based on matrix/spectral theory — that, when opinion is modelled as a relative object, phenomena such as polarisation and periodic behaviour can arise and analysed analytically without introducing additional key modelling ingredients. By contrast, in non-relative approaches these phenomena typically require extra assumptions (such as bounded opinion ranges) and are often less suitable to analytical analysis.

Source note. This chapter draws on the article ‘A relative approach to opinion formation’ [35].