Epidemic mitigation in populations modelled as networks

David Sirl

University of Nottingham

david.sirl@nottingham.ac.uk

Bristol, March 2017

Outline

- Motivation.
- SIR epidemics on a network/graph.
- A simple model of a population as a network/graph.
  - The network.
  - Epidemics on it.
  - Their analysis.
- Vaccination.
  - Allocating vaccine to better mitigate epidemics.
- 'Dropping'.
  - Individuals changing their behaviour in response to the presence of infection.

Motivation

- Mathematically tractable epidemic models are valuable tools for understanding, predicting, mitigating, planning, . . . in the context of infectious diseases.
- Classical models include several assumptions of homogeneity, many of which are unrealistic.
- Heterogeneity has been included in many ways, including households, multiple types, multiple severities, ...
- Our focus is on using network structure to reflect population structures like social networks in human/animal populations, network connectivity of computers.

Stochastic SIR epidemic on a network

Given a graph G (undirected), identify nodes with individuals and edges with 'friendships' and define an epidemic model:

- Initially 1 infectious (chosen UAR) and $n-1$ susceptible.
- SIR (suceptible $\to$ infectious $\to$ removed) progression.
- Infectious individuals remain so for random time distributed as I, then become removed.
- Infectious individuals make contacts with each neighbour in G at the points of Poisson Processes of rate $\lambda > 0$; if neighbour is susceptible it becomes infectious.
- Infectious periods and PPs mutually independent.
- Continue until no infectious individuals remain.

Classical model has $G = K_n$. Analysis is as $n \to \infty$.

- Investigate the number of initial susceptibles that are ultimately removed, the *final size*.

Final size behaviour

Histograms of relative final sizes from 10,000 simulations of a network-based SIR epidemic model, $n = 300$.

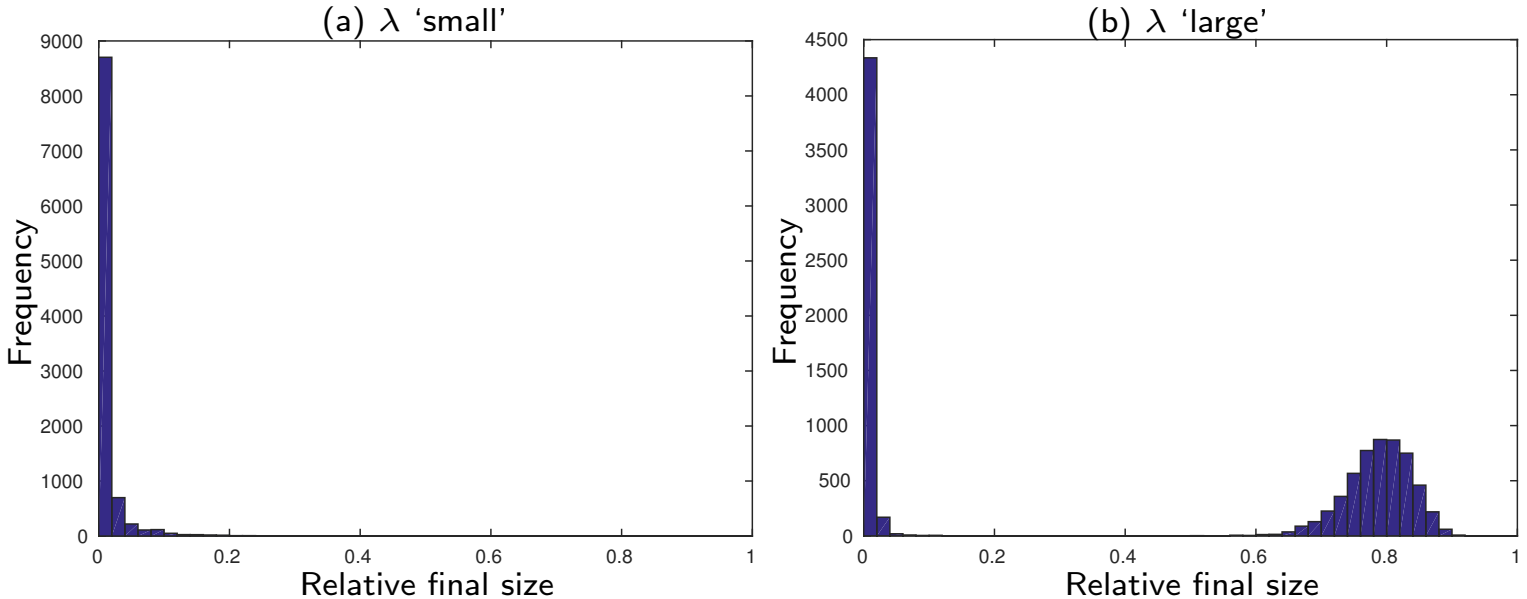

We investigate (i) whether large outbreaks are possible, and if so (ii) how likely they are and (iii) how big they are.

Epidemic model properties

- Main object of interest is the *final size* Z, the number of initial susceptibles that are ultimately removed.
- As $n \to \infty$, we empirically we observe in that either the infection dies out quickly and infects few individuals ($Z = O(1)$) or takes off and infects a significant fraction of the population ($Z = O(n)$).
- By analysing the early stages of the epidemic we find a threshold parameter R_* and the probability of a major outbreak $p_{\mathrm{maj}} = \mathbb{P}(Z = O(n))$; with $p_{\mathrm{maj}} > 0 \iff R_* > 1$.
- We also find the expected relative final size of a major outbreak $z = \mathbb{E}[Z/n \mid \mathrm{Z} = O(n)]$.
- (Can also get CLT for $\sqrt{n}(Z/n - z)$ in the event of a major outbreak.)

BP approximations

- A *forward process* $\mathcal{B}_F$ approximates the spread of infection.
- The criticality of $\mathcal{B}_F$ determines whether a major outbreak is possible.

$$R_* = \text{mean of offspring distn of } \mathcal{B}_F.$$

- Total progeny of $\mathcal{B}_F$ approximates the final size.

$$p_{\text{maj}} \approx \mathbb{P}(\mathcal{B}_F \text{ avoids extinction}).$$

- A *backward process* $\mathcal{B}_B$ approximates the 'spread' of an individual's 'susceptibility set'.
- In the event of a major outbreak, a UAR chosen individual is infected 'iff' its susceptibility set is infinite.
- Total progeny of $\mathcal{B}_B$ approximates the size of an individuals susceptibility set.

$$z \approx \mathbb{P}(\mathcal{B}_B \text{ avoids extinction}).$$

Graphs

- Classical model has $G = K_n$.
- Can also use $G(n, p)$ or random regular graphs[1].
- These represent (more-or-less) homogeneous mixing of homogeneous individuals.
- The degree distribution of these graphs does not reflect what is empirically observed.
  - degree distribution: distribution of the number of neighbours of a randomly chosen vertex.

[1] Neal (2003); Diekmann, de Jong & Metz (1998).

Configuration model

A random graph model with specified degree distribution[2].

- Given n and a degree distribution D,
  - assign each individual $D_i \overset{\mathcal{D}}{=} D$ stubs (half-edges) and
  - pair the stubs UAR.
- This gives a random graph with specified degree distribution, uniformly from all (multi-)graphs on n vertices with that degree distribution.
- There are sufficiently few of these 'imperfections' that they don't affect our analysis.

- No clustering (small loops) or degree correlation (assortativity / disassortativity).

[2] Durrett (2007).

Configuration model

SIR epidemic on a CM random graph[3].

- BP approximations give the following threshold parameter and PGF for z:

$$R_* = p\mu_{\tilde{D}-1} = p(\mu_D + \frac{\sigma_D^2}{\mu_D} - 1),$$
$$f_{B'}(s) = f_{\tilde{D}-1}(1-p+sp), \qquad f_B(s) = f_D(1-p+sp).$$

Here $\mathbb{P}(\tilde{D}=d) = d\mathbb{P}(D=d)/\mu_D$, $p = 1-\phi(\lambda)$, $f_X(s) = \mathbb{E}[s^X]$ and $f^{(i)}(s) = \frac{d^i}{ds^i} f(s)$.

More complex PGFs for p_{maj}, but just as numerically amenable.

[3] Andersson (1997); Newman (2002); Kenah & Robins (2007); Ball, Sirl & Trapman (2009).

Vaccination

- Including the effect of possible interventions is a key use of epidemic models.
- We deal with prophylactic vaccination; vaccination done in advance of any outbreak.
- (Not contact tracing or any other reactive approach.)
- Two key aspects to model:
  - Allocation: who gets vaccinated.
  - Action: the effect on those who are vaccinated.
- We focus on the former.

Vaccination: our baseline models

The CM-SIR epidemic model.

- Configuration model network with degrees $\sim D$.
- Per-pair infection rate λ.
- Infectious period $\sim I$.

- Formulae to compute R_*, p_{maj}, z numerically.

Vaccination modelling

- Vaccine action[4]
  - Perfect: complete protection.
  - All or nothing: complete protection with probability ε, otherwise no effect, independently for each individual.
  - Non-random: rate of incoming PPs multiplied by $a \geq 0$, rate of outgoing PPs multiplied by $b \geq 0$.
- Vaccine allocation
  - Simplest is to vaccinate individuals UAR. Analysis is fairly straightforward.
  - Being more intelligent (i.e. exploiting population structure) allows us to do better.
  - For example, in the standard households model the *equalising strategy* is provably optimal in some circumstances and often (but certainly not always) optimal or near-optimal otherwise[5].

[4] Becker & Starczak (1998).

[5] Ball *et al.* (1997); Keeling & Ross (2015).

Acquaintance vaccination

A challenge is to develop vaccine allocation strategies which target key (well-connected) individuals in the network structure, using only local knowledge about this structure.

One way of doing this is through *acquaintance vaccination*[6].

- Sample individuals UAR with probability p_S.
- Sampled individuals independently name each of their neighbours with probability p_N.
- Named individuals are vaccinated.

[6] Ball & Sirl (2013).

Acquaintance vaccination: BP approximations

Sample individuals UAR w.p. p_S; sampled individuals name each of their neighbours for vaccination independently w.p. p_N.

- Knowing an individual is vaccinated gives some information about its degree.
- We type individuals as $N/V/U$ and S/S^c.
  - Named for vaccination by their infector, Vaccinated but not named by their infector, Unvaccinated.
  - Sampled or unsampled (for possibly naming their neighbours).
- This yields 6-type BPs, from which we derive a threshold parameter R_v and also p_{maj} and z.

Acquaintance vaccination: allocation specifics

- Vaccine coverage is easily shown to be

$$p_V = 1 - f_D(1 - p_S p_N).$$

- This depends on p_S and p_N only through the product $p_S p_N$.
- Performance (measured by R_v, p_{maj} or z) does depend on specific values of p_S and p_N.
- For a perfect vaccine and fixed $p_S p_N$, R_v is increasing in p_N; i.e. it is better to have everyone name a few friends than a few people name all their friends.
- The difference between best and worst is quite small; and for imperfect vaccines the dependence on the precise values of p_S and p_N again seems very small.

Asides before numerical results

- Use the notation $D \sim \text{Hea}(k, \alpha)$ to mean

$$\mathbb{P}(D = d) \propto \begin{cases} k^{-\alpha} & \text{for } d = 0, 1, \ldots, k, \\ d^{-\alpha} & \text{for } d = k+1, k+2, \ldots. \end{cases}$$

- We compare the performance of old and new (best and worst) acquaintance vaccination to vaccinating individuals chosen UAR and to the 'CM-optimal' allocation of vaccinating individuals of highest degree.

Acquaintance vaccination performance 1a

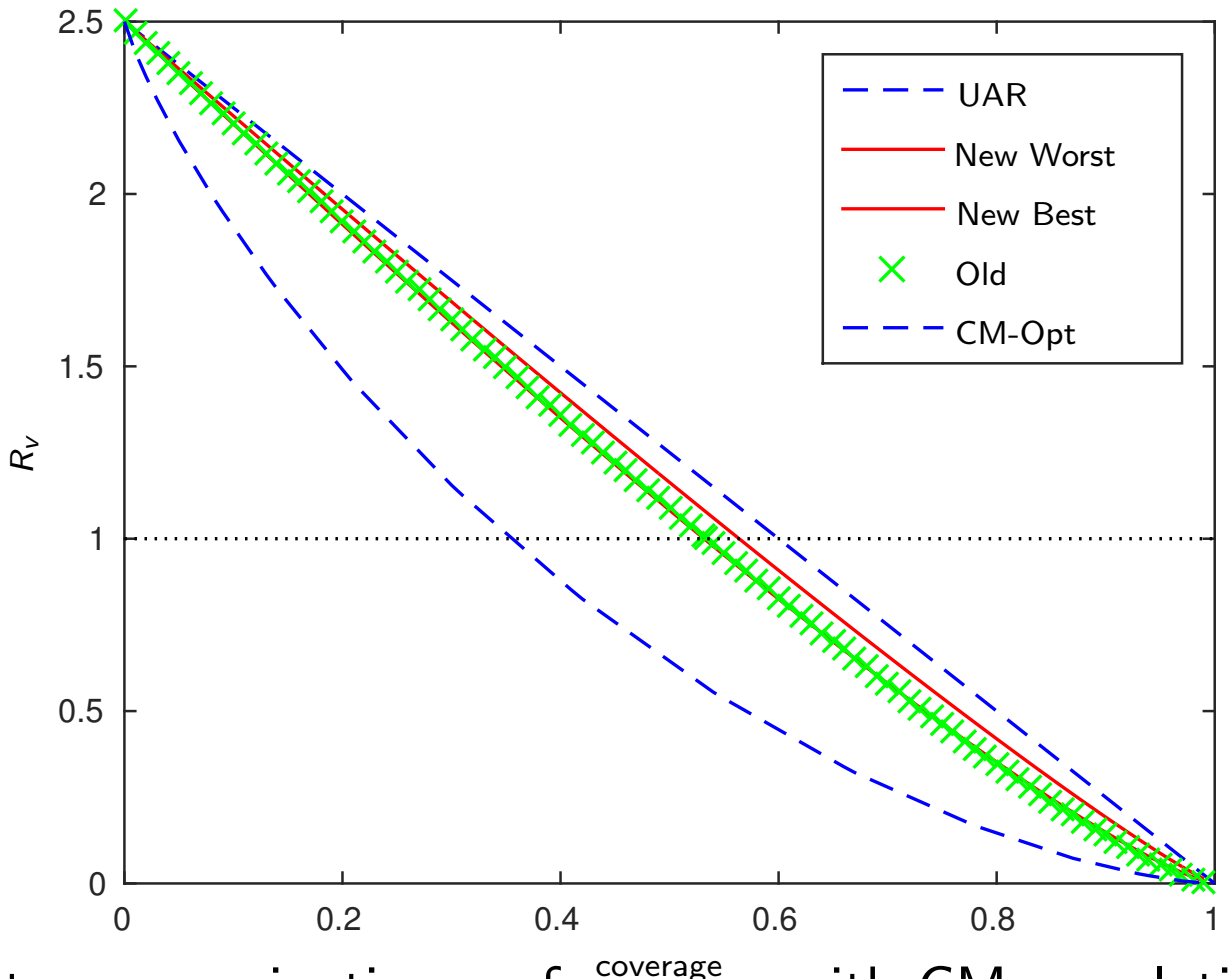

Acquaintance vaccination performance with CM population. Parameters $D \sim \text{Poi}(10)$, $I \sim \exp(3)$ and $\lambda = 1$.

Acquaintance vaccination performance 1b

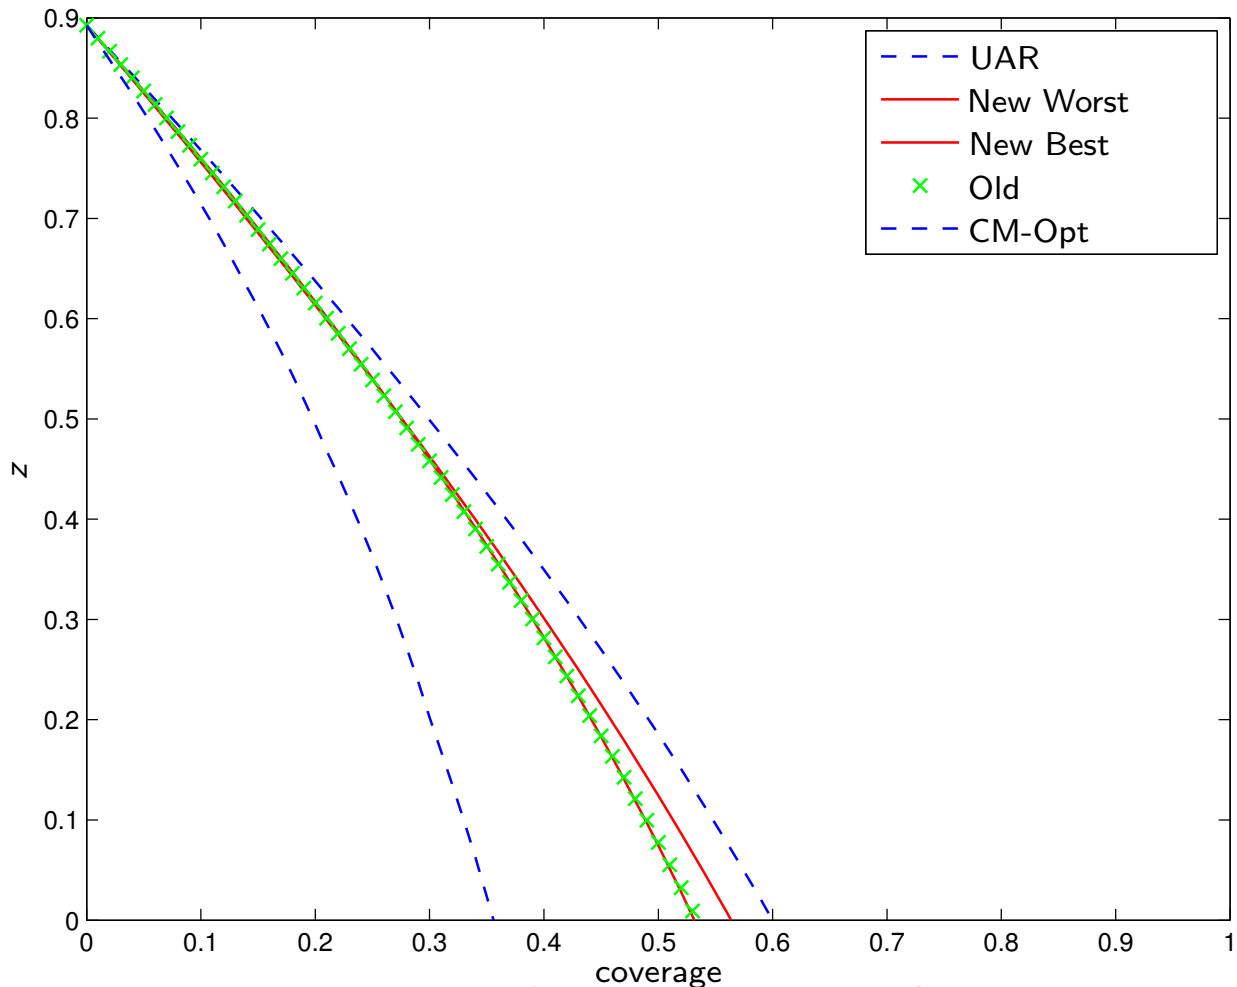

Acquaintance vaccination performance with CM population. Parameters $D \sim \text{Poi}(10)$, $I \sim \exp(3)$ and $\lambda = 1$.

Acquaintance vaccination performance 2a

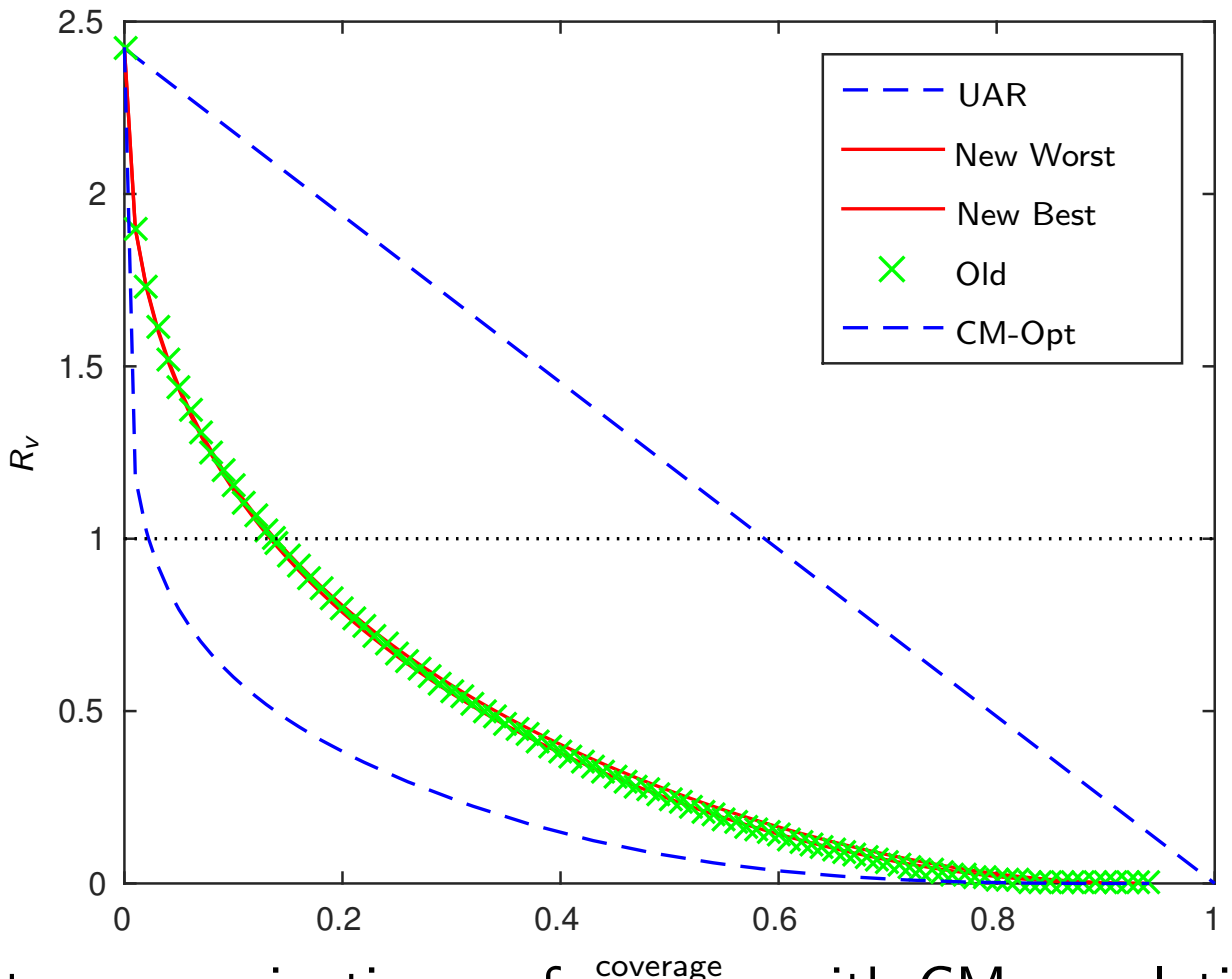

Acquaintance vaccination performance with CM population. Parameters $D \sim \text{Hea}(12, 3.4)$, $I \sim \exp(10)$ and $\lambda = 1$.

Acquaintance vaccination performance 2b

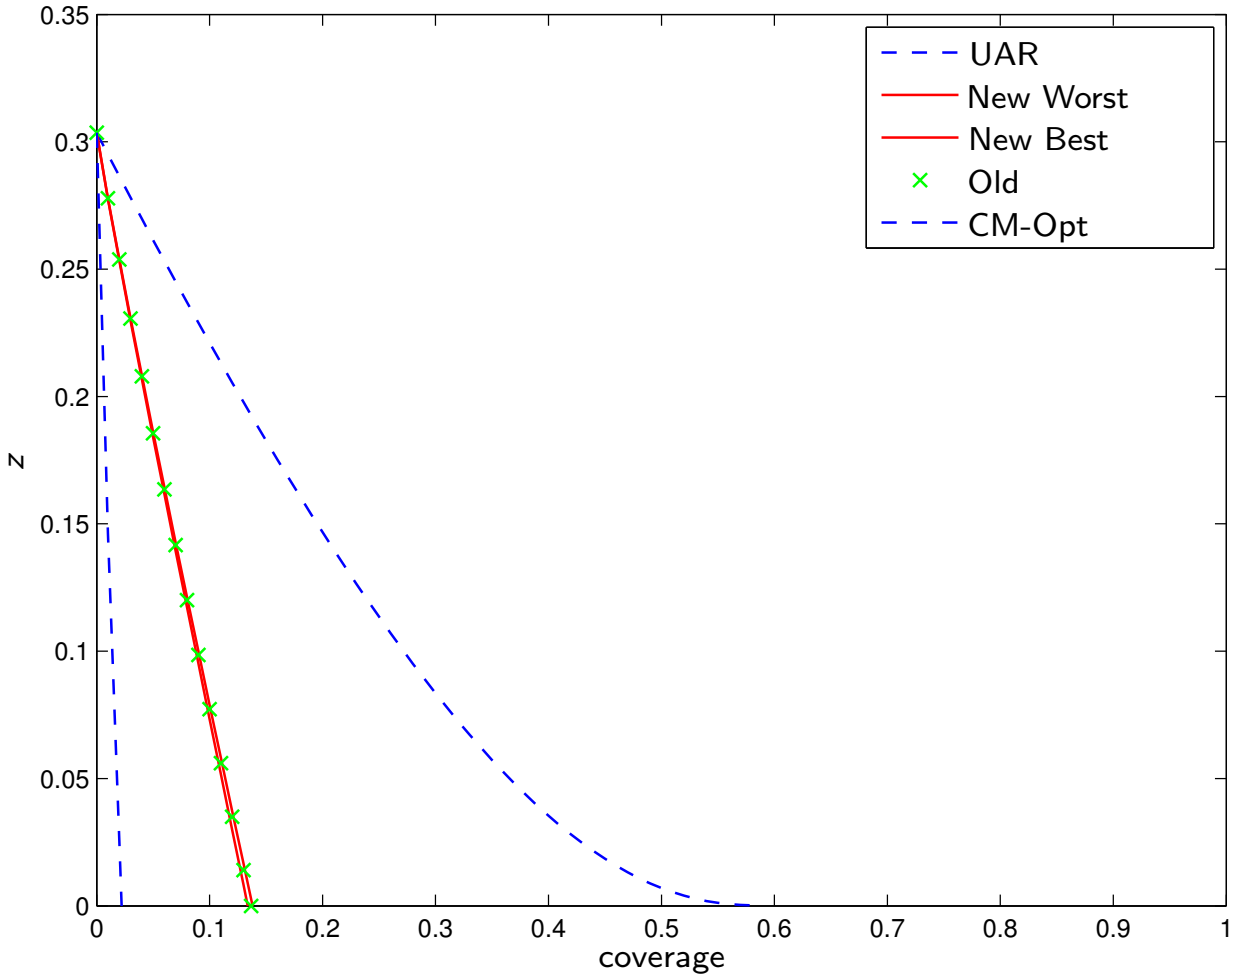

Acquaintance vaccination performance with CM population. Parameters $D \sim \text{Hea}(12, 3.4)$, $I \sim \exp(10)$ and $\lambda = 1$.

k-Acquaintance vaccination in CM population

Sample individuals UAR w.p. p_S; sampled individuals name each of their neighbours w.p. p_N.

- Rather than vaccinating individuals named at least once, vaccinate those named at least k times.
- With higher k this more strongly targets individuals of high degree. (But requires more effort to achieve a given coverage.)
- Branching process of infected individuals ($\to R_v$, p_{maj}, z).
- Now need an 8 type process, typing by whether or not an individual is
  - Named by its infector,
  - Vaccinated,
  - Sampled.
- Numerical results for $k = 2$.

k-Acquaintance vaccination performance 1

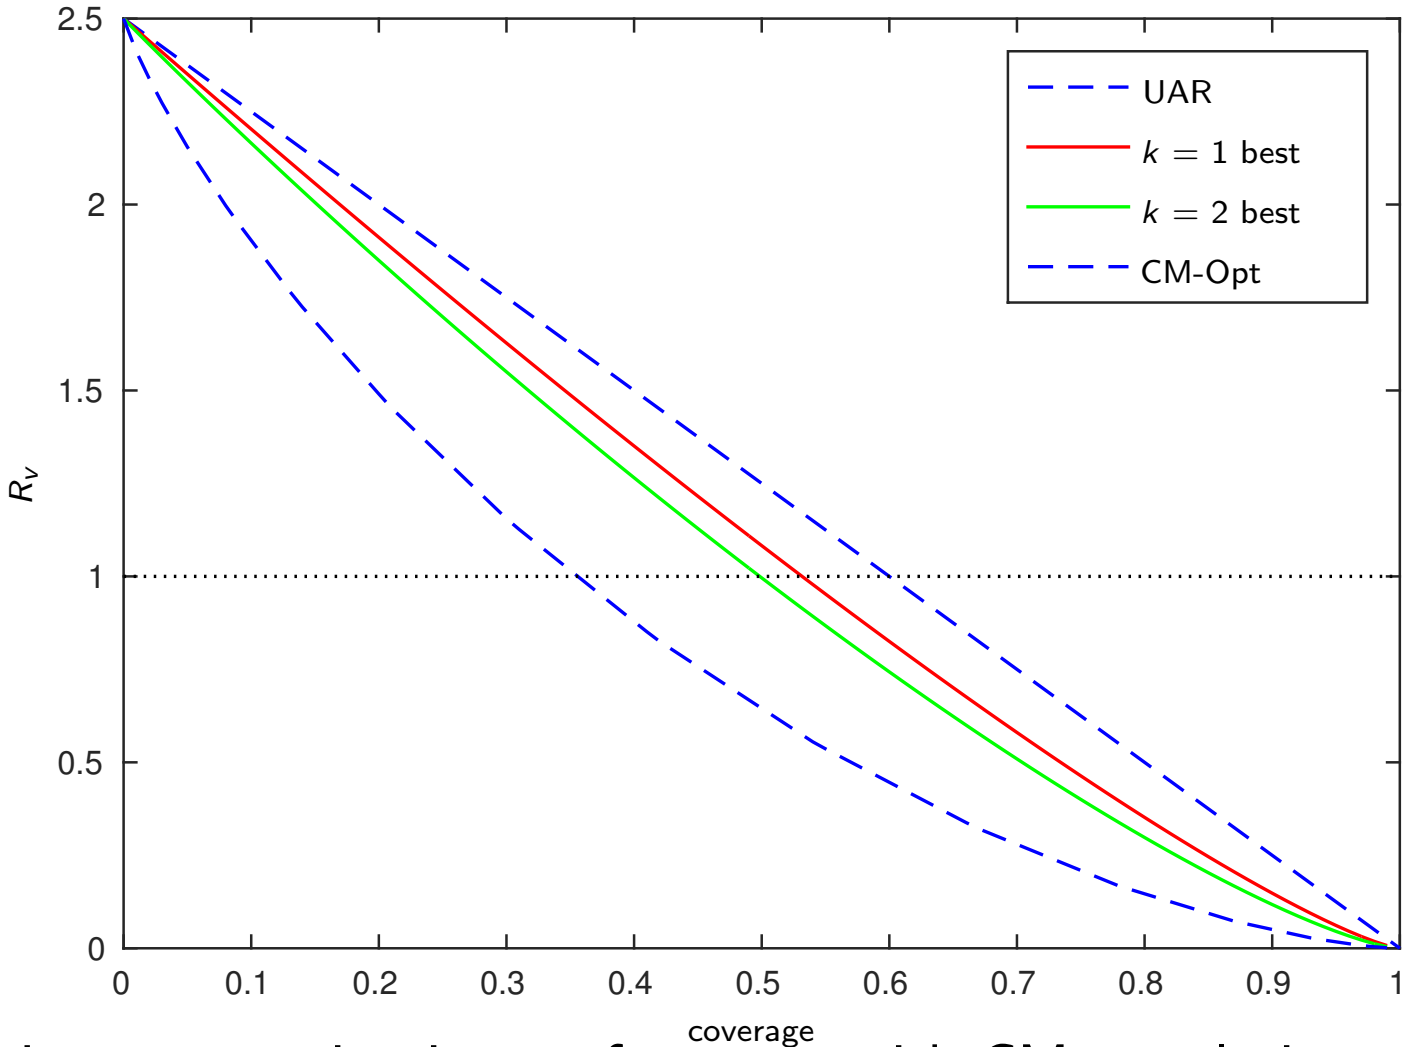

Acquaintance vaccination performance with CM population. Parameters $D \sim \text{Poi}(10)$, $I \sim \exp(3)$ and $\lambda = 1$.

k-Acquaintance vaccination performance 2

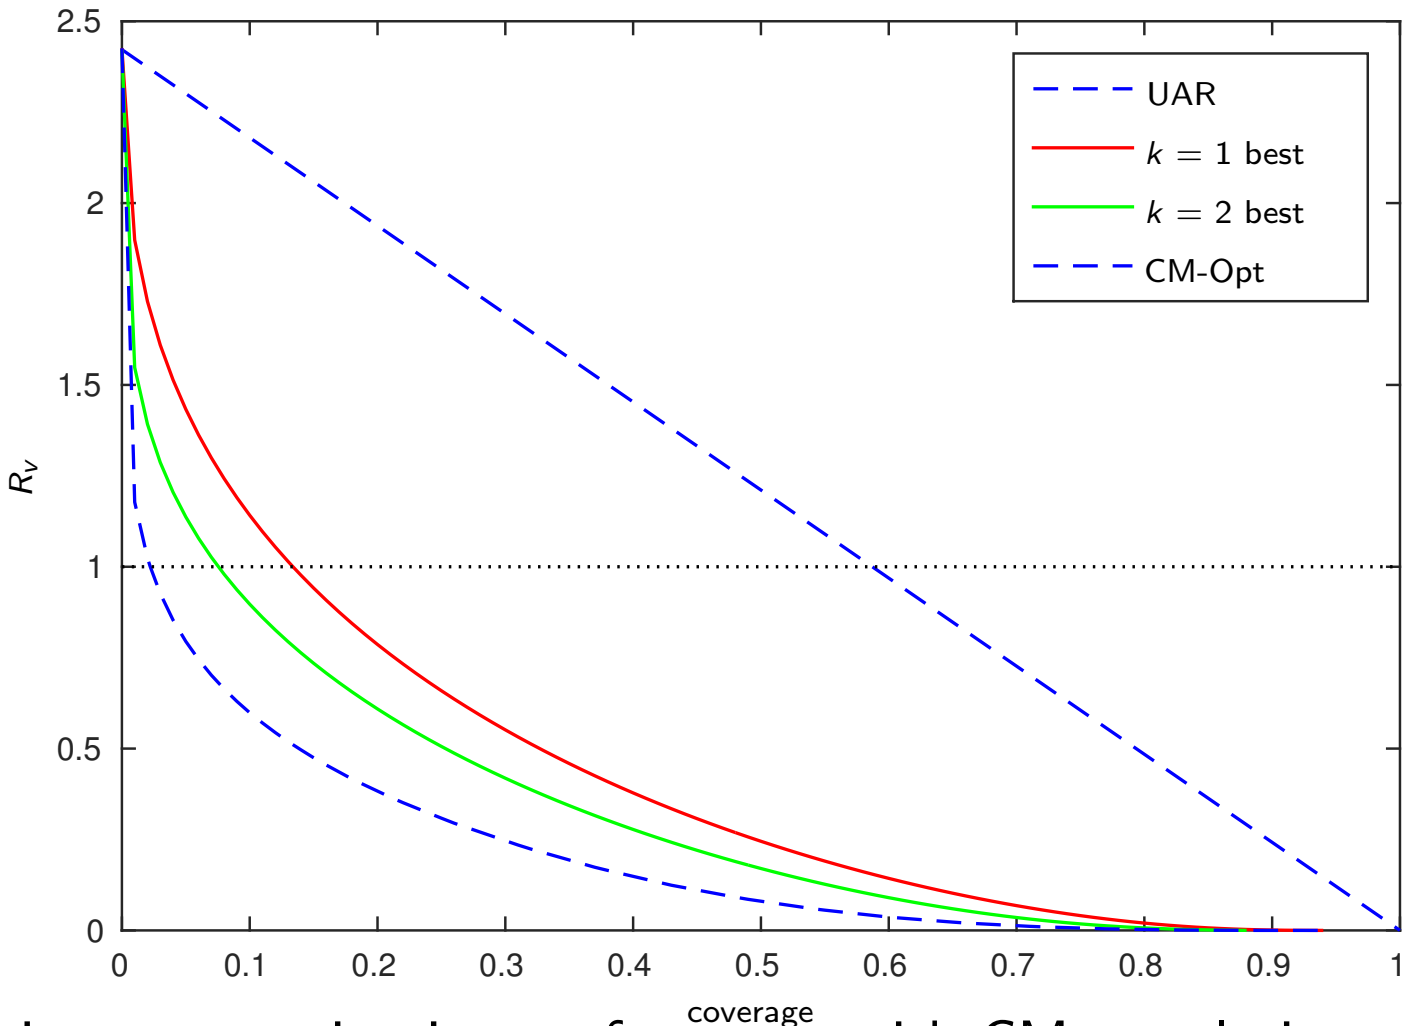

Acquaintance vaccination performance with CM population. Parameters $D \sim \text{Hea}(12, 3.4)$, $I \sim \exp(10)$ and $\lambda = 1$.

Dropping edges in the CM-SIR model

- CM network model, SIR progression as before.
- Infectives
  - infect each neighbour at rate λ,
  - recover at rate γ.
- Also let each neighbour of an infective drop their connection to the infective at rate ω.

Analysis of model with dropping

- Britton *et al.*[7] (i) treat an SEIR model and (ii) allow 'rewiring'.
- Britton *et al.* analyse the early stages (threshold parameters and early exponential growth rate).
  - Branching process approximation.
  - Pair approximation (deterministic large population 'limit' ODEs for the number of singletons, pairs, triples, ... of individuals in the various disease states): system of 7/10 ODEs.
- They find that the threshold parameters disagree, but simulation results are more in agreement with the BP predictions.
- The simpler model is what we[8] have investigated further (as a first step).

[7] Britton *et al.* (2016).

[8] Ball, Britton, Leung & Sirl (in prep.).

Effective degree representation[9]

- Construct the network and the epidemic at the same time.
- Give individuals their (random number of) stubs; an individual's effective degree is the number of remaining free stubs it has.
- Let the epidemic evolve, pairing up stubs only when an infection or *informing* event occurs.
- This leads to a CTMC

$$W(t) = ((X_i(t))_{i=0}^{\infty}, (Y_i(t))_{i=0}^{\infty}, Z_E(t)) \in \mathbb{Z}_+^{\infty} \times \mathbb{Z}_+^{\infty} \times \mathbb{Z}_+.$$

- Here X_i is the number of susceptible individuals of effective degree i, Y_i similarly for infectives and Z_E is the number of unpaired stubs emanating from removed individuals.

[9] Ball and Neal (2008).

Effective degree analysis

- Theory of density dependent population processes[10] gives a functional LLN and CLT, e.g.

$$\lim_{N\to\infty} \frac{W^{(N)}(t)}{N} \to w(t),$$

  where $w(t)$ is the solution of an ODE system and $\to$ can be made precise.
- DD theory applies to finite systems, so must impose a maximum degree (or apply the optimism principle of applied mathematics).
- Infinite system reduces to a single driving ODE; which when $\omega = 0$ is that of Volz/Miller[11].

[10] Ethier & Kurtz (1986, Chapter 11)
[11] Miller, Slim & Volz (2012).

(Expected relative) final size

- Letting $t \to \infty$ in the ODE(s) gives the asymptotic relative final size of the epidemic started with a positive fraction of infected individuals (i.e. a major outbreak).
- This yields implicit equations for z.
- (The same one we saw earlier when $\omega = 0$.)
- Final size in this model with dropping
  $=$ final size without dropping but recovery rate $\gamma + \omega$.

CLT for final size

- Kurtz's theory of DD processes also gives a CLT for the final size of the epidemic:

$$\sqrt{N}\left(\frac{Z^{(N)}}{N} - z\right).$$

  (A little more work is is required to make the DD theory apply.)
- This will suggest a CLT for the size of the giant component in a CM random graph. Previously
  - Derived heuristically for a very special case[12].
  - Asymptotic variance known rigorously[13].

[12] Ball & Neal (2008).
[13] Ball & Neal (2016).

Summary & future

- BP approximations to characterise final size behaviour of SIR epidemic models on random graph population structures.
- Vaccination.
  - Acquaintance vaccination.
  - Including household structure.
  - Further targeting of highly connected individuals.
- Dropping
  - A simple dynamic network.
  - Incorporate 'rewiring'.

- Finding individuals whose vaccination will have the most impact (by some measure).
- Minimising R_v is not always equivalent to minimising z.
- Variations / refinements.
- Performance in other models of network structure.
- A version/variation suitable for implementation?

References I

Andersson (1997). *Math. Biosci.* 140:79–84.

Ball, Mollison & Scalia-Tomba (1997). *Ann. Appl. Probab.* 7:46–89.

Ball & Neal (2008). *Math. Biosci.* 212:69–87.

Ball & Neal (2016). *Ann. Appl. Probab.* To appear.

Ball & Sirl (2013). *J. Appl. Probab.* 50:1147–1168.

Ball, Sirl & Trapman (2009). *Adv. Appl. Probab.* 41:765–796.

Ball, Sirl & Trapman (2010). *Math. Biosci.* 224:53–73. (Erratum *ibid.* 225:81.)

Britton, Juher & Saldaña (2016). *Bull. Math. Biol.* 78:2427–2454.

Diekmann, de Jong & Metz (1998). *J. Appl. Probab.* 35:448–462.

Durrett (2007). Random Graph Dynamics.

References II

Ethier & Kurtz (1986). Markov Processes: Characterization and Convergence.

Keeling & Ross (2015). *Epidemics* 11:7–13.

Kenah & Robins (2007). *Phys. Rev. E* 76:036113.

Miller, Slim & Volz (2012). *J. R. Soc. Interface* 9:890–906.

Neal (2003). *J. Appl. Probab.* 40:779–782.

Newman (2002). *Phys. Rev. E* 66:016128.