Vaccination Strategies

Alberto Maria Segre
Department of Computer Science
The University of Iowa
Iowa City, IA 52246
alberto-segre@uiowa.edu
http://compepi.cs.uiowa.edu/

Vaccination

We discussed vaccination early in the term as one of the available interventions for an infectious disease. Vaccination enabled the global elimination of smallpox in 1979, and lead current efforts to eradicate polio. Once eradicated, vaccination is no longer necessary (or maybe not).

Every intervention has a cost and a benefit. For vaccination, costs are the obvious economic ones, including the cost of manufacturing and administering the vaccine as well as the — generally quite small — risk to individuals who suffer some reaction to the vaccine itself.

The benefits of vaccination are twofold: protection of the immunized individual, and protection of the community by reducing the transmissibility of the disease throughout the herd.

Compared to social distancing, the costs are lower (compare to the economic costs to society we are now or will soon experience) and the benefits are higher, because vaccination can be performed *a priori* in a relatively low-impact fashion.

Vaccination Modeling

Recall the SIR model from Dimitrov and Myers:

$$\frac{dS}{dt} = -\beta \frac{I}{N} S \qquad \frac{dI}{dt} = \beta \frac{I}{N} S - \gamma I \qquad \frac{dR}{dt} = \gamma I$$

> The model then evolves in discrete time steps, with all individuals simultaneously acting as follows in each time step:
>
> (1) Each susceptible individual draws a uniformly random person from the population. If the person drawn is infected, then the susceptible individual changes his state to infected with probability β.
> (2) Each infected individual changes his state to resistant with probability γ.
> (3) Each resistant individual remains resistant.

The easiest way to incorporate vaccinated individuals into this model is to initially label them as R rather than S. Because $\lambda = \beta I / N$ is the "force" that propels infection, vaccination reduces the impact of this force by reducing S. Eventually, the probability of someone going from S to I is vanishingly small, thus achieving "herd immunity."

Non Uniform Mixing and Vaccination Budgets

The model we just saw assumes that (*i*) the vaccine is 100% effective (so that vaccinated individuals are completely out of play), (*ii*) there is plenty of vaccine to go around, and (*iii*) all individuals are interchangeable.

The first assumption is an issue for certain types of infections, where vaccinated individuals might still be vectors or carriers: we'll address this later today.

The second assumption can be relaxed, but isn't really relevant if the third assumption holds (vaccinating some is better than none).

But the third assumption is really suspect because we know mixing is not truly random, and so it follows that vaccinating specific individuals is more advantageous.

If individuals are not interchangeable, the next obvious question is who should be vaccinated if there is not sufficient vaccine to go around?

Medlock & Galvani

Medlock and Galvani address a *resource allocation question*: in situations where vaccine doses are limited, who should be vaccinated to maximize the overall benefit to the population?

The question is only meaningful if the population is meaningfully differentiated. Otherwise, vaccinating any random individual will always have exactly the same impact (reducing S by 1).

Medlock & Galvani choose to compartmentalize the population by age. Other choices might be location, gender, types of pre-existing conditions, and so on. But age here makes sense from the disease standpoint:

- The old and infirm die at differentiated rates from the flu.
- The very young don't have immunity, and are at higher risk.
- Vaccines are often not as effective in older people.

An Explosion of Compartments, An Explosion of Parameters

When you compartmentalize the population, the number of parameters you need to estimate increases.

You will need to know age-appropriate values of β and γ, the age distribution of the entire population and within each age group. You'll also need to know R_0 and mortality rates by age group (here, fit to 1918 and 1957 pandemics, both bad flu years) along with other disease parameters.

We also know that social behaviors differ by age group (*e.g.*, school, parent/child, work, etc.), so we'll need to estimate within and across group contact patterns, which themselves are clearly not uniform.

If pushed to its natural limit, each compartment would contain a unique individual, and you would need to estimate these parameters for every single individual in your simulation (hence: agent based modeling).

Vaccination and Optimization

First, we'll cover the main points of the paper. Once we are clear on those, we'll look at how Medlock & Galvani handled the modeling issues.

Medlock & Galvani set out to answer the question *what is the optimal vaccination strategy* that provides the *best outcome* for the *lowest cost*.

There are lots of variants of this kind of *optimization problem*, which commonly occurs in resource allocation contexts. The involve finding the optimal way to spend your money (resources) to best achieve the outcome you want.

First, you need to define the problem more precisely:

- what does "best" really mean, and how do you measure it?
- is the budget potentially unlimited, and you are just being cheap?
- or is the budget otherwise limited by external factors?

Best Outcome

Medlock & Galvani describe five different (some more reasonable than others) outcome measures, and solve the problem for each of them:

- minimizing infections;
- minimizing deaths;
- minimizing YLL, or "years of life lost" (weighted by age);
- contingent valuation (from surveys); and
- actuarial cost (including of vaccine, future earnings, lost work, etc.).

They then derive how many doses of vaccine would be required to extinguish the pandemic ($R < 1$) using the 2007 US population (300M) and reasonable disease parameters.

The also solve the problem for three different fixed vaccination budgets (20M, 40M and 60M vaccine doses).

In both cases, they solved for "what is the allocation of the budget by age group to achieve the best outcome?" five times, once for each outcome measure.

Optimal Vaccination Doses by Initial R_0

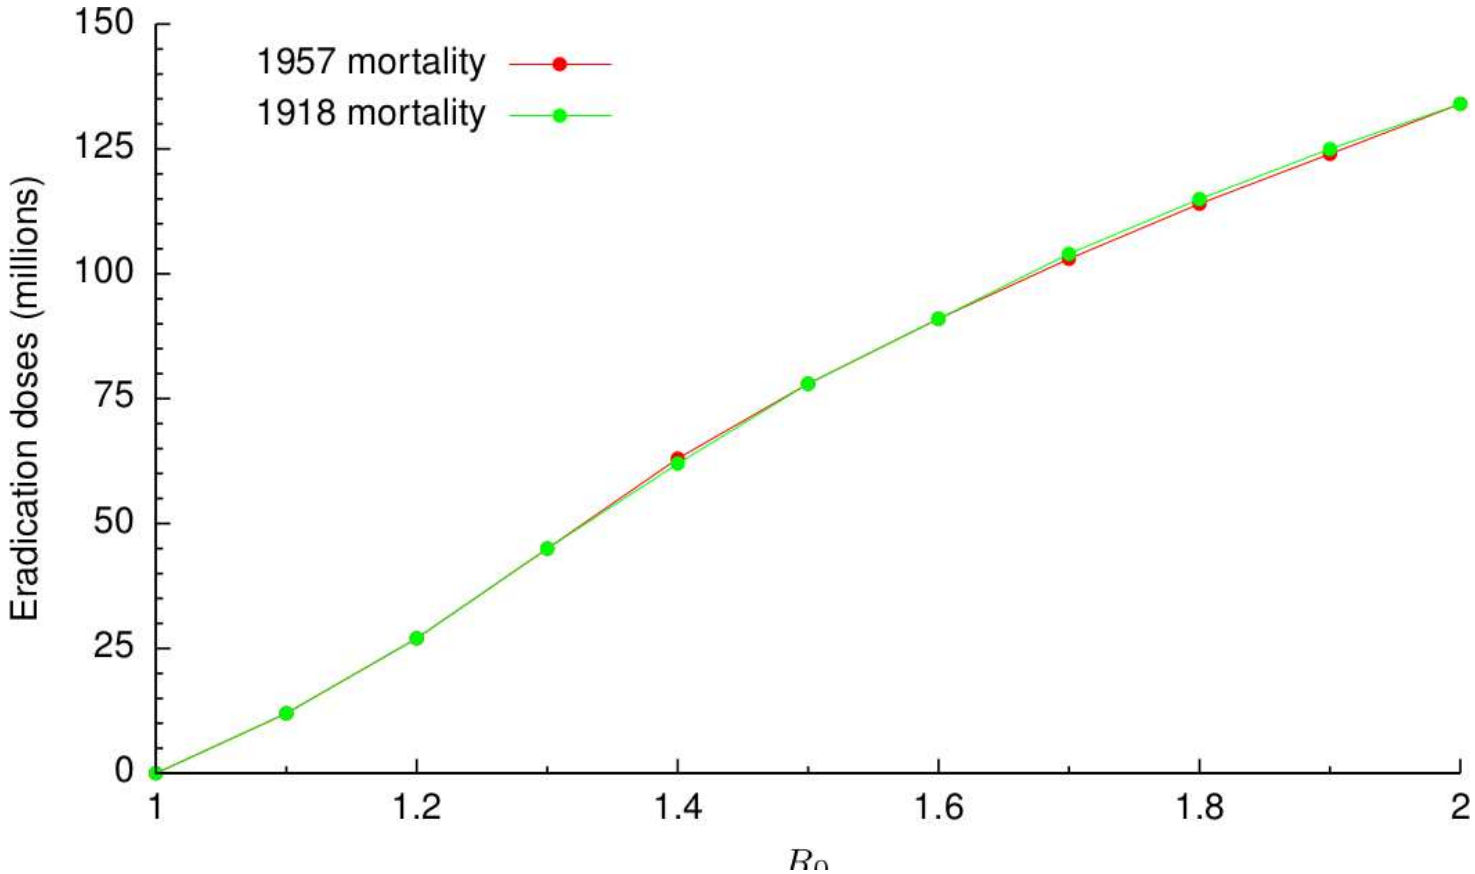

Number of doses required under optimal allocation to achieve eradication ($R < 1$) by initial R_0 (reproductive number in a completely susceptible population).

Optimal Distribution for Eradication ($R_0 = 1.4$)

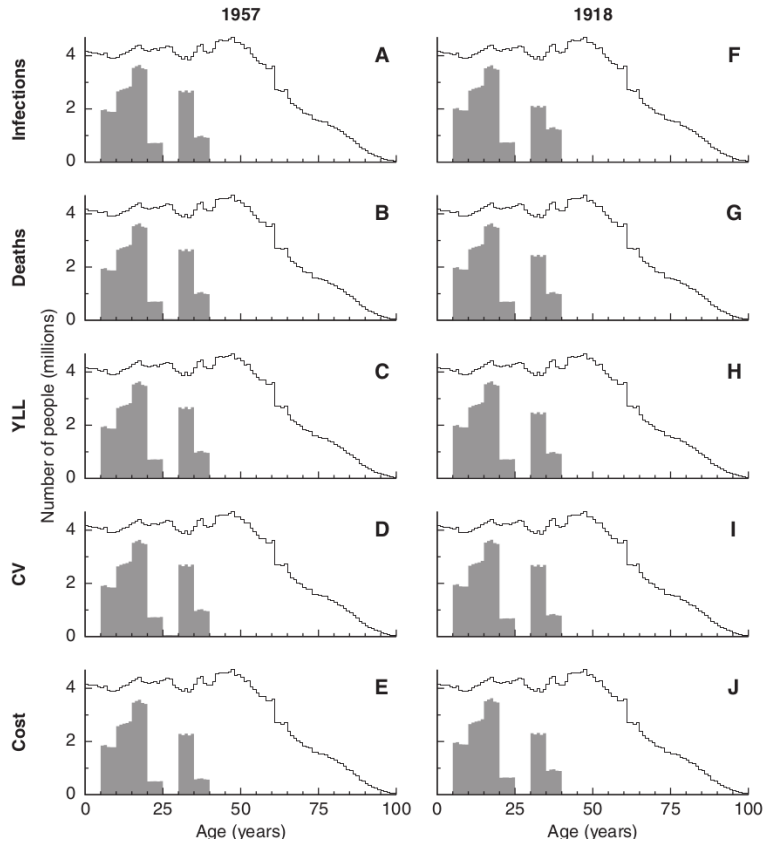

Eradication ($R < 1$) for all five outcomes fit against 2007 population (L: 1957, R:1918 mortality rates): focus on 5-19 and 30-39 reflects contact patterns. Here, 62-63M doses; compare to ˜85M doses dispensed in 2009.

Limited Vaccination Budgets

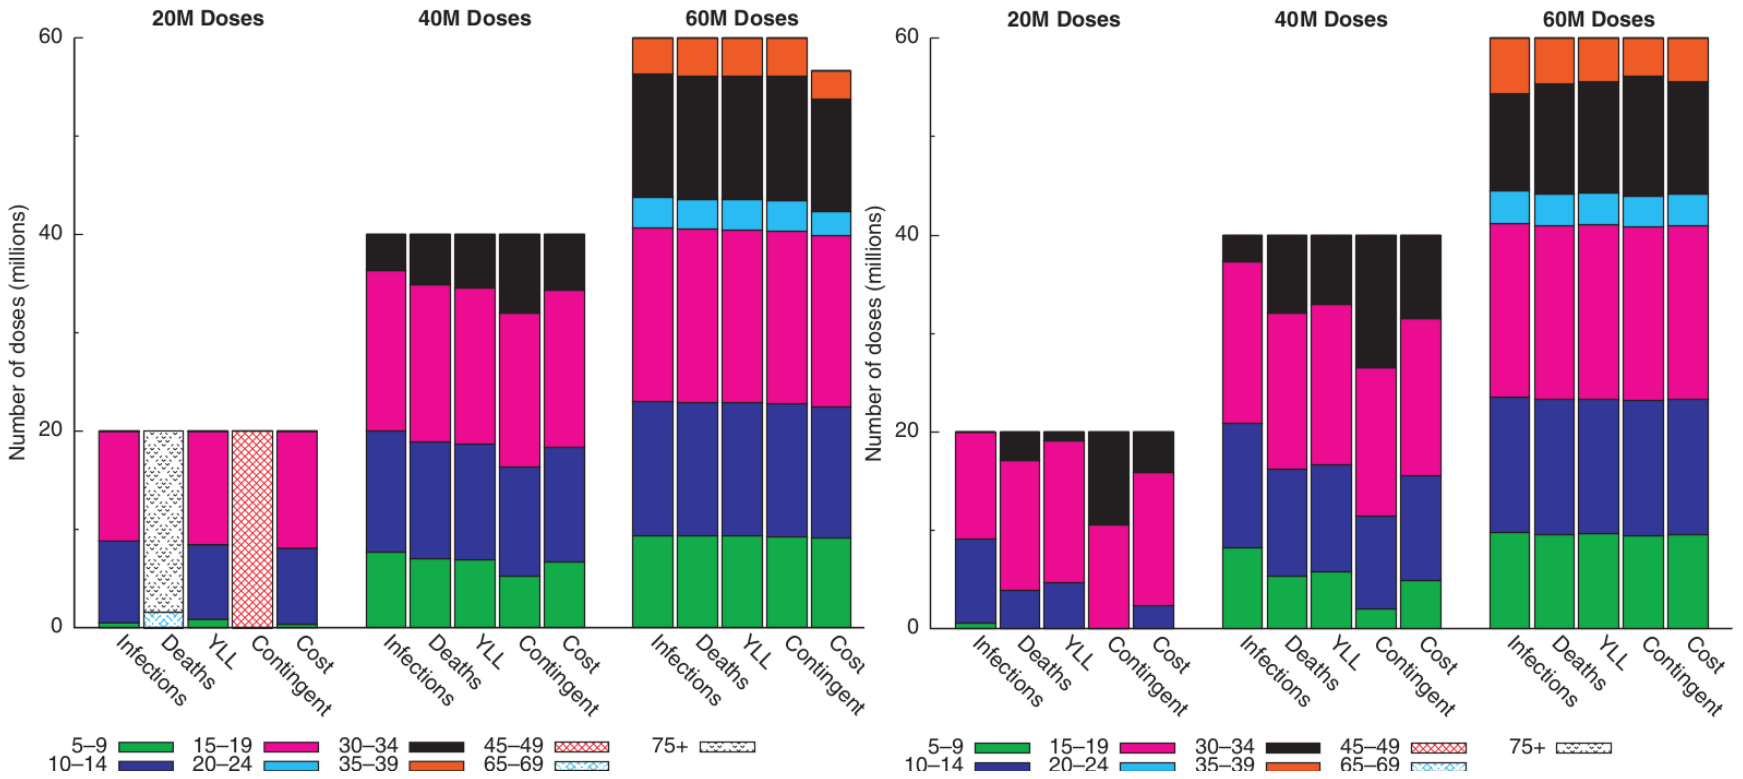

With fewer doses available, we can't eradicate influenza. If above 37M doses, ages 5-19 and 30-39 are prioritized for all outcome measures and both mortality models (L: 1957, R:1918). Left: below 37M outcomes based on deaths prioritize the elderly (left), and below 36M outcomes based on valuation prioritize ages 45-49 (combination of economic value and infection severity).

Limited Vaccination Budgets

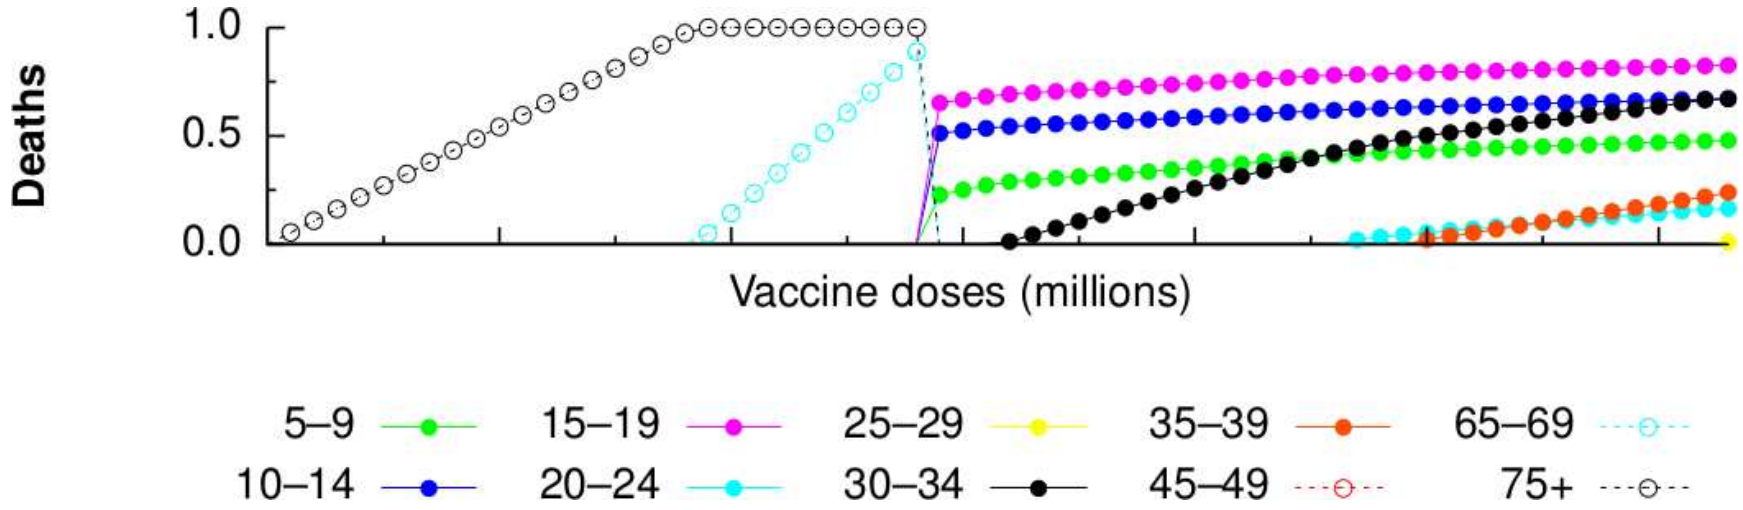

In the previous slide, we saw the extra value of vaccinating the elderly when budgets were small and deaths were the primary outcome. As more doses become available, it is more effective to vaccinate those who are "around" the elderly.

These all make sense but are quite fiddly depending on doses, R_0 and outcome being optimized.

How Modeling Informs Policy

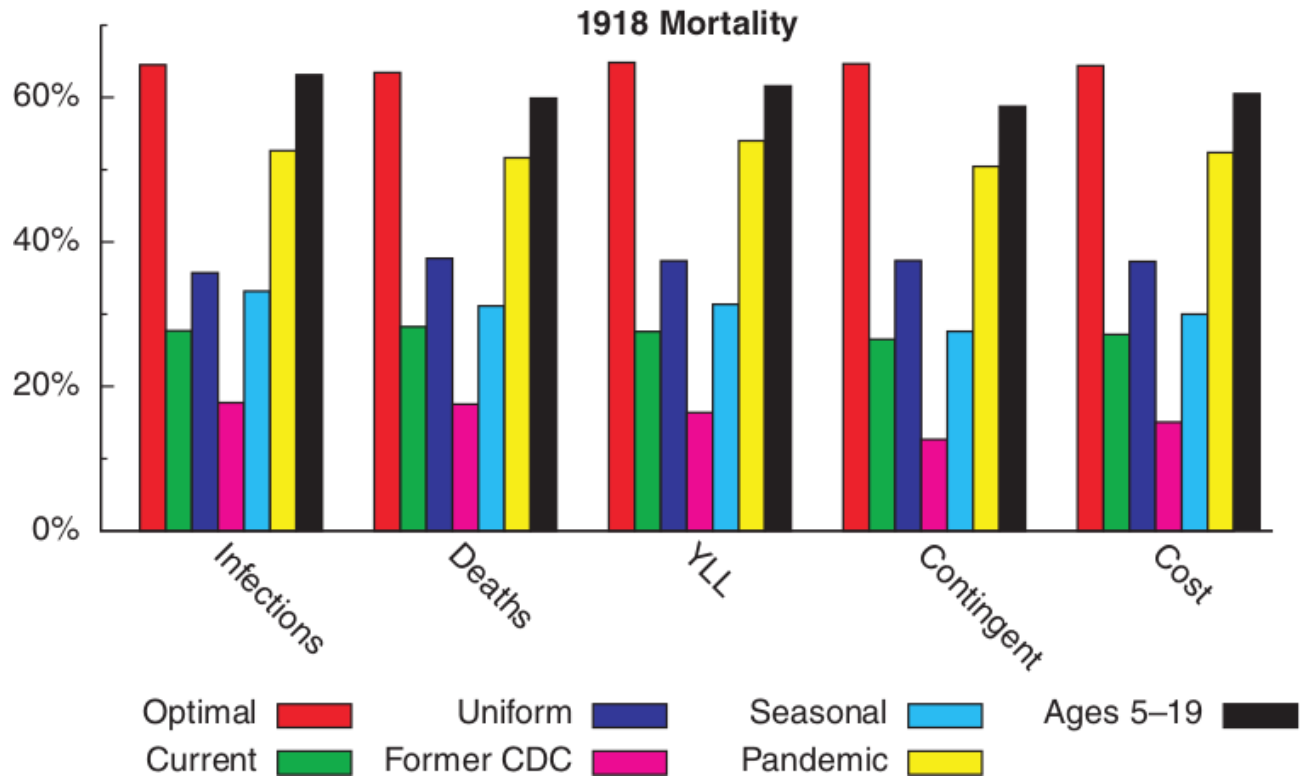

Outcome reduction (vs no vaccine) by policy (1918 mortality). Shows how models can inform policy: "current" is actual distribution; "uniform" vaccinates all age groups equally; "former" CDC is old policy; "seasonal" is new policy for seasonal influenza; "pandemic" is for H1N1-like flu (novel swine-origin); and "ages 5-19" allocates all vaccination to that age group.

How Did They Do That?

There are lots of results packed in this paper and the supplement. They consider, for example, what changes when the vaccine is not 100% effective, as well as what happens when vaccine effectiveness changes by age group.

The paper itself is almost the TL;DR version of their results!

In the supplement, Medlock & Galvani do the real work of describing their models and what they observe, even if some of their results seem whacky (does anyone really believe the contingent valuation formula?).

Its worth examining the model in some detail to see if we can really believe, *e.g.*, the policy implications given above.

I'll try to give the gist of it, but will have to skip some details (like difference between vaccine efficacy against infection, ε and against death, δ).

Medlock & Galvani Model

Starting from an SEIR model, Medlock & Galvani split the population into two "tracks," with vaccination status determined prior to simulation:

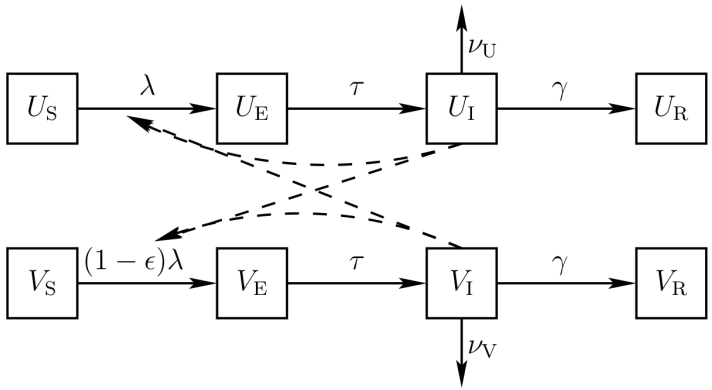

Where λ is the infection rate, τ is the progression rate, and γ is the recovery rate (β replaced by λ and τ for SEIR).

ε is the effectiveness of the vaccine, which can fail (hence $V_S \rightarrow V_E$, but at a lesser rate than $U_S \rightarrow U_E$).

V_R and U_R are semantically identical; ν_U and ν_V are death rates.

Medlock & Galvani Model

As with our SIR model, we have a system of differential equations, but with this two-track SEIR model, there are eight and not three equations:

$$\frac{dU_S}{dt} = -\lambda U_S \qquad \frac{dU_E}{dt} = \lambda U_S - \tau U_E$$

$$\frac{dU_I}{dt} = \tau U_E - (\gamma_U + \nu_U) U_I \qquad \frac{dU_R}{dt} = \gamma_U U_I$$

$$\frac{dV_S}{dt} = -(1-\varepsilon)\lambda V_S \qquad \frac{dV_E}{dt} = (1-\varepsilon)\lambda V_S - \tau V_E$$

$$\frac{dV_I}{dt} = \tau V_E - (\gamma_V + \nu_V) V_I \qquad \frac{dV_R}{dt} = \gamma_V V_I$$

Now imagine cloning this model into 17 layers, one per age group (1-4, 5-9, 10-14, ... 70-74, 75+), and fixing the interactions across layer to reflect the contact parameters across ages.

Medlock & Galvani Model

We'll need to come up with (at least) $17 \times 4 \times 2 = 136$ different parameters: λ, τ, γ and ν for each tracks in each of the 17 layers.

Some of these parameters are independent within the track, while others depend on the other track (or even the other layers).

For example, remember those dotted lines?

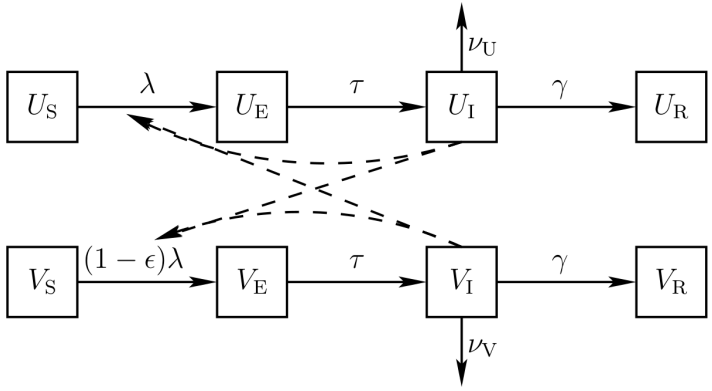

These show influence in the forcing functions, where, $U_S \rightarrow U_E$ has terms from both U_I and V_I; the λ must reflect these values.

Medlock & Galvani Model

Medlock & Galvani calculate age-specific λ as follows:

$$\lambda_a = \sum_{\alpha=1}^{17} \frac{\phi_{\alpha a}(\beta_U U_{I\alpha} + \beta_V V_{I\alpha})}{N} = \frac{1}{N} \sum_{\alpha=1}^{17} \phi_{\alpha a}(\beta_U U_{I\alpha} + \beta_V V_{I\alpha})$$

Where $\phi_{\alpha a}$ is the contact rate for infectious individuals from age group α with susceptible individuals of age group a, weighted by relative population size N_α (note $\phi_{\alpha a} = \phi_{a\alpha}$ by design).

Note β_V and β_U are probabilities of transmission given a contact by a vaccinated infectious person and an unvaccinated infectious person, respectively.

Medlock & Galvani Model

The ϕ contact parameters were obtained from a contact diary-based study in eight European countries. Averaging over all eight countries and using means to make the parameters symmetric, we get (log scale):

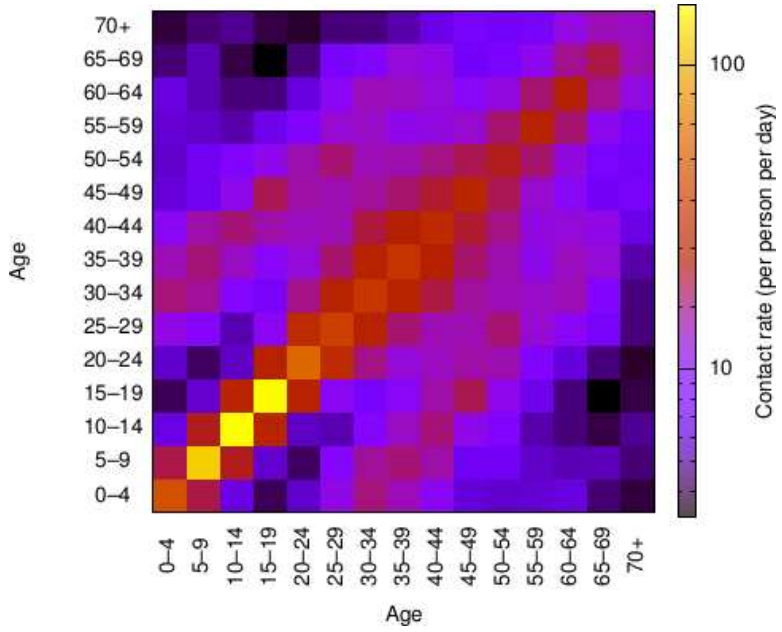

Notice strong within-age (diagonal) and child-parent (sub/super diagonals).

Solving the Model

You should have a pretty good sense of what it takes to fit such a complex model; there's more detail of how Medlock & Galvani calibrate their parameters, such as ν_V,ν_U, ε_a and so on, sometimes against known pandemics (1918 and 1957) in the supplement.

Once the model is complete, they solve for the age-specific distribution of vaccine doses, p_a, the proportion of age group a that is in the vaccinated track.

This is a nonlinear constrained optimization problem, optimizing a pre-selected objective (outcome) function while honoring two constraints: $0 \leq p_a \leq 1$ and $\sum_{a=1}^{17} p_a N_a \leq W$ where W is the vaccination budget.

Unfortunately, the outcomes have no easily computed first derivative, so the usual gradient descent inspired mechanisms are not feasible here.

COBYLA and Sensitivity Analysis

The COBYLA algorithm iteratively solves nonlinear constrained optimization problems.

Each iteration, COBYLA solves linear approximations of the objective and the constraints in the area of the current solution, and then moves within this constrained choice set.

It's just one of many such algorithms; choosing the "best" one often reduces to what you know about the problem, and how complex the objective function and constraints are.

Final note: Medlock & Galvani go the extra mile in the supplement to see what happens as their chosen parameter values vary. Much appreciated, and extremely appropriate.