What is behavioral epidemiology and what does math have to do with it?

Winfried (Vinny) Just

Department of Mathematics, Ohio University

Presented at California State University, Channel Islands,

October 18, 2021

Epidemiology, behavioral epidemiology, and medicine

Epidemiology studies how infectious diseases spread through populations of humans, animals, or plants.

The subfield of *behavioral epidemiology* studies the interplay between the spread of diseases and human reactions to observed or predicted spread.

Mathematical modeling in behavioral epidemiology as an area of research started only around 2002 and has seen explosive growth since then. A 2016 survey of the literature on vaccination decisions alone already lists 777 references (Wang, Z. et al., Statistical physics of vaccination. Phys Rep 664:1–113).

Note that epidemiology is very different from *medicine:*

While medicine aims at curing diseases, epidemiology aims at limiting the spread of infections and thereby preventing people from becoming sick in the first place.

In a nutshell: Epidemiology is trying to mess up medicine's business model.

Epidemiology: Some basic terminology

Epidemiology is by nature an interdisciplinary science. On the biological side, we need to know what changes are caused when *pathogens* (viruses, bacteria, fungi, protozoa) invade the organism of a *host* (a human, animal, or plant) and how these pathogens spread from host to host.

Transmission of a pathogen from one host to another may require a *contact* between the two hosts (like in seasonal flu or COVID-19) or an intermediary, called an *insect vector* (like in malaria or in Chagas disease, which is endemic in South and Central America).

Sidebar: I was introduced to epidemiology through joint research on the spread of Chagas disease with Mario Grijalva from OU Heritage College of Osteopatic Medicine and my former Ph.D. student Bismark Oduro.

In contrast to medicine, epidemiology does *not focus on individual cases,* but takes a bird's eye view of how pathogens spread through *populations* of hosts.

What's math got to do with it?

The most important goal of epidemiology is to gain insight into the effectiveness of possible *control measures* or *preventive measures,* such as *vaccination, quarantine,* or *behavior modification* (like social distancing, wearing masks, or lockdowns).

Guiding question of epidemiology: *By how much can we limit the spread of an infection within a population at what cost?*

This question can be studied with mathematical models of how pathogens spread in populations.

- Transmissions of pathogens between individual hosts are stochastic events, so the most natural framework for building such models are stochastic processes.
- But one can often simplify them to useful approximations with ODE or difference equation models.
- The insights from the biological side of epidemiology enter these models as assumptions and parameters.

Staples of mathematical epidemiology: States and compartments

Models of mathematical epidemiology assume that at any given time, each host is in one of several specified *states,* such as:

- *susceptible* (not infected, can become infected)
- *exposed* (infected, but not (yet) infectious)
- *infectious* (able to infect others)
- *removed* (no longer susceptible to becoming infections and can no longer infect others)

Not all these states are considered in all epidemiological models. Models can also have additional states, like "asymptomatic infectious" for studying COVID-19.

A *compartment* of an epidemiological model is the set of hosts in a given state. The model itself is then conceptualized in terms of hosts moving between compartments over time.

Model types: Possible moves between compartments

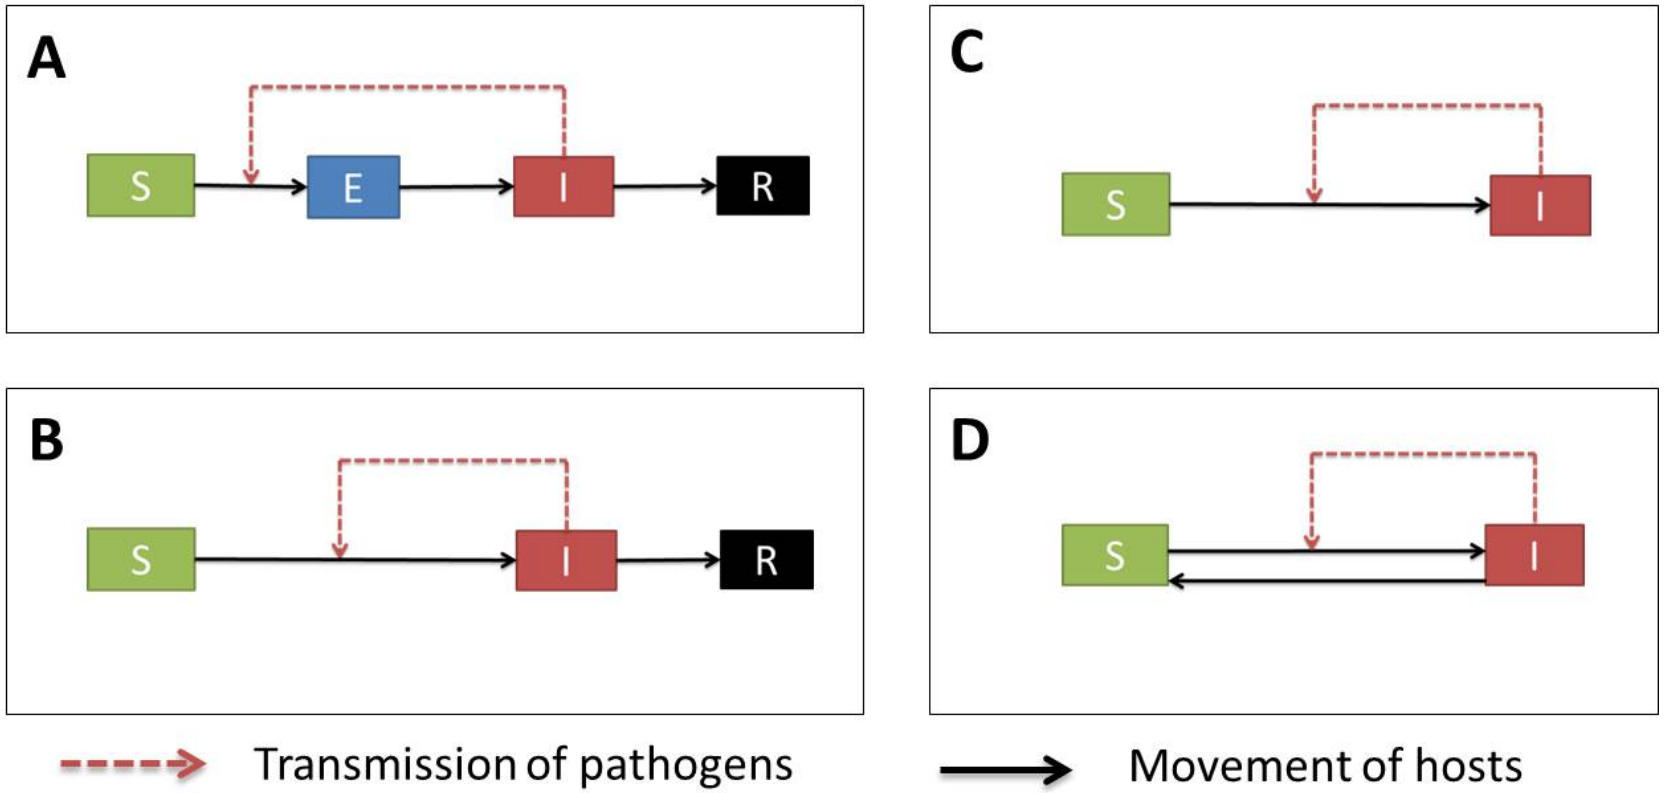

Figure: Four basic model types: SEIR (A), SIR (B), SI (C), SIS (D)

The choice of the model type depends on the characteristics of the infectious disease that we want to study.

Example: An ODE version of the *SIR*-model

The variables represent the proportions of hosts in the S-, I-, R-compartments; $\alpha, \beta > 0$ are model parameters.

$$\frac{ds}{dt} = -\beta si$$
$$\frac{di}{dt} = \beta si - \alpha i = (\beta s - \alpha) i$$
$$\frac{dr}{dt} = \alpha i$$

Let $R_0 := \frac{\beta}{\alpha}$. It is called the *basic reproduction number.*

- When $R_0 > 1$, we will initially observe a near exponential increase of infectious hosts, at least as long as $s(t) \approx 1$.
- When $R_0 \leq 1$, the model predicts decrease of i at all times.

The importance of R_0 in disease modeling

The official definition of R_0 conceptualizes it as the average number of secondary infections that will be caused by introduction of one *index case* (infectious individual) into an otherwise entirely susceptible population.

The subscript 0 in this notation means that R_0 is a special case of the more general notion of the *reproduction number* R or R_t at time $t = 0$.

R_0 is a measure of contagiousness and is the most important parameter in modeling the spread of infectious diseases. For very simple models like the one on the previous and next slides, it essentially determines all the interesting predictions.

This is no longer true though for more realistic and elaborate models.

The value of R_0 for the flu and COVID-19

For the seasonal flu, R_0 changes from year to year and has historically been between 1 and 2, with a mean of 1.3.

For the original strain of COVID-19, it has been estimated that $2.5 < R_0 < 3$.

For newer variants of COVID-19, R_0 can no longer be estimated directly from epidemiological data, since by now we don't have conditions anywhere that would feature an entirely susceptible population that does not practice any preventive measures. But we can compare reproductive numbers between different variants.

It was found that reproductive numbers for the α variant of COVID-19 are about 60% higher than for the original variant.

Similarly, an influential study found that the δ variant is about 50% more "infectious" than the α variant, but this wasn't actually based on a comparison of reproduction numbers. Current estimates of R_0 for the δ variant being between 5 and 9 need to be treated with caution.

The herd immunity threshold

$$\frac{ds}{dt} = -\beta si$$
$$\frac{di}{dt} = \beta si - \alpha i = \beta i\left(s - \frac{\alpha}{\beta}\right) = \beta i\left(s - \frac{1}{R_0}\right)$$
$$\frac{dr}{dt} = \alpha i$$

The sign of $\frac{di}{dt}$ becomes negative when a proportion of at least $HIT := 1 - \frac{1}{R_0}$ of hosts is no longer susceptible.

The quantity HIT is called the *herd immunity threshold.* Based on the estimates of R_0, it was be around 0.65 for the original variant of COVID-19.

Question: Does this mean that if we were to let this original variant of COVID-19 spread freely through the population until around 65% of people will have experienced infection, the remaining 35% would be protected (from the original strain)?

Herd immunity alone won't help much

No.

First of all, it is unclear to what extent models of type SIR or SEIR are adequate for COVID-19. There are documented cases (very few so far) where immunity after recovery from the infection was subsequently lost. If this turns out to be common, we would need to use models of type SEIRS, where there is no herd immunity.

And even in our simple SIR model, new infections only reach a peak, but do not stop, when herd immunity is achieved. For example, when $R_0 = 3$, this model predicts herd immunity when at least $2/3$ of the population have experienced infection, but predicts a *final size* of 0.94, which means that around 94% of the population would eventually experience infection.

This *overshoot phenomenon* tends to be totally ignored in public discussions of controlling COVID-19 through achieving herd immunity.

How do control measures enter the picture?

So far we have only considered the unmitigated spread of an infection and have not talked about control measures.

Essentially, control measures alter some initial conditions or parameters of the underlying disease transmission model, and their effectiveness can to some extent be studied within the framework of traditional mathematical epidemiology. For example:

- Vaccinations will move people into the R-compartment without the experience of actual infections. When people get vaccinated **prior** to a disease outbreak, this changes the initial conditions. If vaccine uptake corresponds to at least the herd immunity threshold, this will then protect the entire population.
- Cancelling certain social activities, like attending large parties or taking vacations in crowded locations, will reduce the number of contacts and thus reduce the reproduction number of airborne infections.

Staples of behavioral epidemiology: Awareness and behavioral reponses

- Similarly, wearing masks and keeping 6 feet apart will reduce the likelihood of transmission during a single contact and will also reduce the reproduction number of airborne infections.

But all this is going to help only if people actually decide to get vaccinated and to adopt these behavioral changes.

Behavioral epidemiology studies how and when people make such decisions.

Generally speaking, they will do so only if they are *aware* of the dangers of an outbreak. A standard assumption in behavioral epidemiology is that adoption of control measures like the ones listed here constitute *behavioral responses* that are induced by awareness.

What if awareness wanes over time?

Awareness of an outbreak may be generated by information about new infections, either through direct observation or media reports. It may also spread from person to person by word of mouth or through social media, not unlike the spreading of an infection.

One would expect that awareness, and the behavioral response induced by it, might wane over time. This would certainly be the case if the behavioral response has driven down new infections.

Question: Could this waning of awareness and the induced behavioral response lead to sustained oscillations of the numbers of new infections if the underlying dynamics of the disease is of type SIS, where hosts do not acquire immunity to reinfection upon recovery?

We studied this question together with Joan Saldaña of the University of Girona and my former Ph.D. student Ying Xin in: Oscillations in epidemic models with spread of awareness. J. Math. Biol. 76, 1027–1057 (2018).

SAIS models and SAUIS models

In the paper, we constructed and explored two ODE models that are similar in spirit to the first one I showed you, but more elaborate. Here I will only outline their basic features and skip the details.

For the first model, which we called a model of type SAIS, we added one compartment of susceptible, but aware individuals (the A-compartment) to a standard SIS-model. Hosts in this compartment adopt a behavioral response that lowers their chances of becoming infected, and also attempt to spread their own awareness to others.

In the second model, which we called a model of type SAUIS, we added a second compartment of aware hosts to an SAIS model, which we called the U-compartment. Hosts in this compartment adopt a behavioral response that lowers their chances of becoming infected, but are unwilling to spread their own awareness to others.

It depends how awareness spreads from person to person

We proved that sustained oscillations are ruled out in models of type SAIS and all trajetories will always approach equilibria.

On the other hand, we rigorously demonstrated that models of type SAUIS permit sustained oscillations for certain parameter settings. Technically speaking, SAUIS models permit *Hopf bifurcations* in their dynamics, while SAIS models do not.

According to our interpretation of these results, the cause of these oscillations is not the waning awareness all by itself, but the combination of this phenomenon with the differential pattern of sharing awareness.

An interesting open question is what happens if the underlying disease dynamics in these models is of type SIRS (or SEIRS) instead, as it may well be for COVID-19. My Ph.D. student Ying had done some simulations for such models that indicated even more exotic bifurcations than mere Hopf bifurcations. But then she graduated, and we did not pursue this topic further.

To get vaccinated or not to get vaccinated?

That became the question for each of us.

Here is how a completely rational person would approach the problem (for a hypothetical perfect vaccine; but the principle is similar for a 95% effective one):

- Getting vaccinated has a small cost $c_v > 0$ (time spent, mild side effects, etc.).
- Not getting vaccinated has a much larger cost $c_i \gg c_v$ **if** I do get infected, but has no cost at all **if** I luck out and don't catch the disease anyway.
- If x is the probability of an unvaccinated person getting infected, I should get vaccinated if, and only if, $xc_i \geq c_v$.

Question: What's missing from the above argument?

The probability x is not fixed, but depends on the *vaccination coverage,* that is, the overall proportion V of people in the population who decide to get vaccinated.

Vaccination games

So, a population of perfectly rational people might decide about vaccinations as follows: Let's each make independently our decisions to get vaccinated with probability V^*. This will give a population-wide vaccination coverage of (roughly) V^*.

If V^* is chosen in such a way that

$$c_v = x(V^*)c_i,$$

then nobody will regret their decision, and V^* represents the *Nash equilibrium* of a model called a *vaccination game.*

Alas, the Nash equilibrium is not optimal in the sense of minimizing the expected cost to society as a whole.

Under mild additional assumptions, a vaccination coverage $V = HIT$ would be societally optimal, but $V^* < HIT$.

To see the latter, it suffices to know that for large populations $x(HIT) \approx 0$, so that $c_v > x(HIT)c_i$.

But people are not rational. So there is some hope.

Real people use heuristics such as imitating successful others for making their decisions. Vaccination decisions through imitation have been extensively studied. The seminal paper Fu F, Rosenbloom DI, Wang L, Nowak MA (2011) Imitation dynamics of vaccination behaviour on social networks. Proc R Soc B 278(1702):42–49 roughly uses the following setup for studying vaccinations against the seasonal flu:

- Hosts independently make their vaccination decisions before each flu season.
- Each host compares their own cost from last season with the cost of one randomly chosen other member of the population.
- Each host then switches their strategy (vaccinate/don't vaccinate) with a probability that is given by a *Fermi function* of the difference in costs.
- The resulting vaccination coverages and infection probabilities are then calculated in a difference equation model based on expected values of V and on an SIR model for the flu.

So, does imitation work better than rational self-interest?

No. At least not under generic assumptions.

That's what the vast literature says. But pretty much all of this literature uses the same Fermi function for the switching probability as the seminal paper:

$$p_{switch} = \frac{1}{1 + e^{-\beta(C(\text{your strategy}) - C(\text{other}))}}$$

Here $\beta > 0$ is a parameter that can be interpreted as a degree of certainty about the actual cost difference.

Well, if it comes from physics and if Professor Nowak from Harvard introduced it into this context, it must be good stuff, right?

Except that psychological experiments on real decision-making didn't find a good match and needed additional parameters to find a decent fit; see Traulsen A, Semmann D, Sommerfeld RD, Krambeck HJ, Milinski M (2010) Human strategy updating in evolutionary games. Proc Natl Acad Sci USA 107(7):2962–2966.

Generalized Fermi functions

Together with David Gerberry of Xavier University and Ying Xin, in Open-minded imitation can achieve near-optimal vaccination coverage. J. Math. Biol. 79, 1491–1514 (2019), we generalized Fermi functions as follows:

$$p_{switch} = \frac{1}{\alpha + e^{-\beta(C(\text{your strategy})-C(\text{other}))}} = \frac{\alpha^{-1}}{1+\alpha^{-1}e^{-\beta(C(\textit{your strategy})-C(\textit{other}))}}.$$

The parameter $\alpha > 1$ can be interpreted as a degree of open-mindedness, a willingness to try out a new strategy if it doesn't seem to be too much worse than what you have been doing previously.

For the same generic case as in Fu et al., we found both analytically and and by simulation studies that by choosing α and β large enough, we can get equilibrium vaccination coverages above V^* and arbitrarily close to the societal optimum HIT.

Lots of open questions remain

While the results of this paper are encouraging, they only cover the simplest case. It remains open to what extent our results generalize:

- To imperfect vaccines.
- To models that incorporate the structure of contacts in the population.
- To imitation of more than one other.
- To diseases that resemble COVID-19 more than the seasonal flu.

The list goes on, and we have some very interesting partial results for some of these problems.

But I currently don't have committed collaborators for this topic.