Should I get a flu shot?
How well did this go last year?

Winfried Just, David Gerberry, Ying Xin
Departments of Mathematics
Ohio University and Xavier University

ECMTB 2018, Lisboa, Portugal

July 25, 2018

Meet the presenters

WJ: Hi! I'm Vinny from Ohio University.

DG: Hi! I'm David, a mathematician from Xavier University.

WJ: Oh! I'm no good at math.

DG: So you are in the wrong place then!

WJ: But so glad we met. I have a problem that you as a mathematician might be able to help me with.

DG: Would be happy to.

WJ: You see, I am trying to decide whether or not to get a flu shot this year, I'm trying to make a rational decision.
How would you mathematicians approach this problem?

Costs

DG: Well, there is a cost of vaccination.

WJ: You mean, like, waiting in line, getting poked with a needle ... and there might be nasty side effects. But perhaps I will not have to wait in line or suffer any side effects?

DG: Let's consider the average cost, and let c_v denote it.

WJ: Catching the flu is a lot nastier though than getting a flu shot.

DG: You are saying that the (average) cost of infection c_i is a lot larger than c_v,

that is, $c_i \gg c_v > 0$.

Probability of infection and vaccine efficacy

DG: Do you always catch the flu when you don't get vaccinated?

WJ: No, last year I remained unvaccinated and did not catch the flu.

DG: Let x denote the probability that an unvaccinated person will catch the flu. This probability will depend on the vaccination coverage V, so that $0 \leq x(V) < 1$.

DG: And is it true that a person who does get vaccinated never catches the flu?

WJ: No!!! Two years ago I did get vaccinated. And then I caught the flu nevertheless. That was really bad.

DG: I am sorry to hear this. So we might need to consider another parameter rE, called the efficacy of the vaccine. For an ideal vaccine, we would have $rE = 1$. In general, let us assume here that the probability of a vaccinated person catching the flu is $(1 - rE)x$.

Summary of our notation

- There is a cost associated with getting a flu shot.

  Let c_v be the average cost of the vaccination.

- There is another cost of getting infected with the flu.

  Let c_i be the average cost of infection.

- We have: $c_i \gg c_v > 0$.

- Let $0 \leq x < 1$ denote the probability that an unvaccinated person gets infected.

- Let $0 < rE \leq 1$ be the vaccine efficacy.

  For illustrative purposes we interpret rE in these slides in such a way so that the probability for a vaccinated person to catch the flu is always $(1 - rE)x$.

Expected costs for vaccinators and nonvaccinators

DG: Now you can calculate your expected costs when you vaccinate and when you don't vaccinate.

WJ: If I don't vaccinate, my expected cost will be

$C_u = C_u(V) = c_i x(V)$,

and if I vaccinate my expected cost will be

$C_v = C_v(V) = c_v + c_i(1 - rE)x(V)$.

So since $c_i \gg c_v$, the cost for not vaccinating will be higher and everybody should vaccinate!

Herd immunity

DG: Not necessarily. When rE is not too small, then there exists a vaccination coverage $V_{hit} < 1$, called the "herd immunity threshold," such that for all $V \geq V_{hit}$ we have $x(V) = 0$.

WJ: Great! So then it would suffice to vaccinate a proportion of $V_{hit} < 1$ of the population to provide perfect protection for all. We could then save the cost of vaccinating a proportion of $1 - V_{hit}$ of the population.

DG: Exactly! So it would not be necessary or optimal for **everybody** to get vaccinated.

But who should get vaccinated?

WJ: But who should and who shouldn't get vaccinated? And should **I** or shouldn't **I**? This is exactly my dilemma.

DG: If the government were to draw up a list ...

WJ: You must be kidding ... How could you trust **them** with a problem of minimizing costs??
I don't want no government making health care decisions for me.

DG: Who should decide then?

WJ: We, the people. Like you and me. By making rational decisions as individuals, we will arrive at the vaccination coverage V_{hit} that's best for everybody.

Individual vaccination decisions

DG: How would this work?

WJ: From what you said earlier, the probability $x(V)$ that an unvaccinated host catches the flu will depend on the vaccination coverage V. Then the costs also must depend on V.

When the vaccination coverage is too low, we will have $C_u(V) > C_v(V)$, so that rational people will choose to vaccinate, which will increase V.

When the vaccination coverage is too high, we will have $C_v(V) > C_u(V)$, so that rational people will choose not to vaccinate, which will decrease V.

In this way, individual choices by rational people will drive the vaccination coverage to some equilibrium where $C_v(V) = C_u(V)$.

Nash equilibria in the vaccination game

DG: What you described here is the outline of a mathematical model called vaccination game.

We can consider everybody making a vaccination decision as an individual player who tries to maximize their expected payoff by minimizing costs.

One can choose to vaccinate or not to vaccinate, these are the pure strategies.

When the payoffs $-C_v(V)$ and $-C_u(V)$ of the pure strategies are equal, then the population has reached a Nash equilibrium.
Let V_{Nash} denote the vaccination coverage at the Nash equilibrium.

.

Nash equilibria in the vaccination game

In mathematical models called vaccination games we consider vaccination decisions made by individual players who try to maximize their expected payoff by minimizing costs.

One can choose to vaccinate or not to vaccinate, these are the pure strategies.

When the payoffs $-C_v(V)$ and $-C_u(V)$ of the pure strategies are equal, then the population has reached a Nash equilibrium with vaccination coverage V_{Nash}.

At a Nash equilibrium no player has any regrets about their strategies given the choices of all other players and no incentive to switch to another strategy.

Nash equilibria *vs.* societal optimum

WJ: So here is my point. Consider a population with some people always vaccinating and some people never vaccinating. But if, for example, $C_u(V)$ for such a population is larger than $C_v(V)$, then some of the non- vaccinators would switch to vaccinating based on rational self-interest. This process will then lead to the Nash equilibrium, where $C_v(V_{Nash}) = C_u(V_{Nash})$, and nobody has any regrets whatsoever.

DG: Yes, this is what models of the vaccination game predict. And Nash equilibria are usually defined as situations where no player has any regrets about their strategies given the choices of all other players and no incentive to switch to another strategy.

WJ: Beautiful! So perfectly rational people will arrive at a no-regrets-whatsoever situation with optimal vaccination coverage as in your herd immunity threshold by just making rational decisions on how to randomize their individual vaccination choices. No government meddling required!

Is $V_{Nash} = V_{hit}$?

DG: You are assuming here that the vaccination coverage at Nash equilibrium is optimal and is equal to the herd immunity threshold. But this is not true.

WJ: Now give me a break: Doesn't "optimal" mean the same thing as "no regrets whatsoever?"

DG: We are talking about different types of regrets. At Nash equilibrium, nobody has any regrets about their **individual** decision. At V_{hit}, we have no regrets about the cost to society as a whole.

WJ: So how could what's best for each of us individually not be best for all us?

DG: It's possible.
What would be the rational choice for you, and thus for everybody, when $V = V_{hit}$ so that $x(V) = 0$?

WJ: Then $C_u(V) = c_i x(V) = 0$ for an unvaccinated person, and
$C_v(V) = c_v + c_i(1 - rE)x(V) = c_v$ for a vaccinated person.
Thus $C_v(V_{hit}) > C_u(V_{hit}) = 0$. So $V_{Nash} < V_{hit}$.

What's to be done about it?

DG: In other words, individually optimal decisions lead to a suboptimal outcome for the whole society.

WJ: Bummer! Anything you mathematicians can do about it?

DG: That would take an effort of the whole society. As mathematicians, we can only carefully study whether our models are realistic and make accurate predictions.

First, we need to carefully check our assumptions. The prediction of a Nash equilibrium is based on the idea that everybody makes perfectly rational decisions.

WJ: Are you saying that since most people aren't all that smart, there is some hope?

DG: I would not say it this way. But as a society we could help people in making more beneficial decisions.

How do real people make decisions?

WJ: OK, but what I meant was this: People like me wouldn't even know how to make the calculation for the Nash equilibrium.

DG: So how do you usually arrive at your vaccination decisions?

WJ: You didn't notice?

DG: ?? Notice what?

Imitation

WJ: I might ask an expert, like you.

DG: Well, thank you, but ...

WJ: And if I hadn't met you by accident, I would ask my friend George how things went for him last year. If what he did worked reasonably well, I might do the same this year.

DG: So you might then imitate George's strategy.

WJ: You can call it this way. It think if people were to imitate good decisions of other people, that would lead to better outcomes for the society as a whole.

DG: This conjecture has been widely studied.

Imitation of randomly chosen people

WJ: So what have these studies found?

DG: The literature reports that when $c_i > 2c_v$ the population will always arrive at a vaccination coverage that is even lower than V_{Nash}, with an even higher cost to the overall population.

WJ: Bummer again!

DG: But this may be an artifact of how the process of imitation is conceptualized.

WJ: What do you mean?

DG: I will tell you, but first explain to me how, exactly, you imitate your friend's George's strategy.

Imitating George

WJ: Well, most of the time, I would just do what I did last year. But once in a while, I would ask George what his cost was last year. If it was lower or at least in the same ballpark, I would most likely switch to his strategy. But if his cost was a lot higher than mine, I will most likely stick to my own previous strategy.

DG: You said "most likely." So: Not always?

WJ: Yeah. He or I might have just lucked out with not getting vaccinated.

Fermi functions

In the literature the probability p_{switch} of switching is usually modeled by a so-called Fermi function:

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

with $\beta > 0$. When β gets larger, this becomes closer to a best-response function.

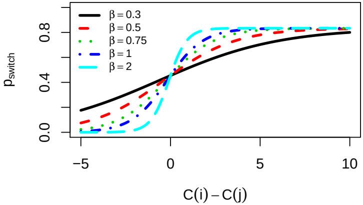

Are Fermi functions realistic?

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

WJ: But wait! When my strategy has a higher cost, then this model predicts that I would switch to the other with probability > 0.5. This isn't what I do. Most of the time I just stick with my strategy for the previous year.

DG: My collaborator Winfried Just of your institution, Ohio University, noticed the same thing.

WJ: Really? I know this guy! Sounds like trouble brewing. So what happened next?

Generalized Fermi functions

Some empirical research in the psychological literature supports more flexible functional forms of the switching probabilities. We generalized the Fermi function by introducing a parameter $\alpha \geq 1$ so that:

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

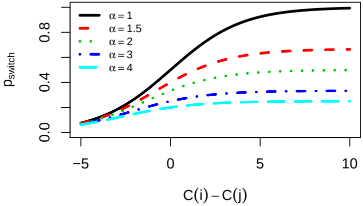

Generalized Fermi functions, continued

WJ: Now that would be closer to my way of imitating!

DG: You can think of α^{-1} as the probability of making your decision in a given year by possibly imitation somebody else.
If α is large, then p_{switch} would always be close to 0. Recall that:

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

The following figure shows how the second fraction, which represents the conditional switching probability, depends on α.

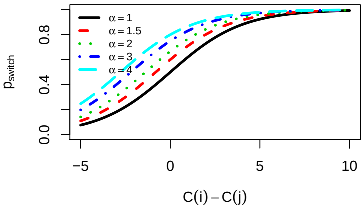

WJ: Right! If I do consider imitating somebody else, then my switching probability is close to 1, unless that other person did really poorly.

DG: So you would be open-minded about trying out the other's strategy.

WJ: You can call it this way. But would that make any difference in vaccination games?

DG: Yes. In our simulations, we found that when $\alpha = 1$, as in the previously published papers, then the population converges to a vaccination coverage $V^* < V_{Nash}$, with $C(V^*) > C(V_{Nash})$.

However, for sufficiently large values of α we found many parameter settings where $V^* > V_{Nash}$, with $C(V^*) < C(V_{Nash})$.

Results of our simulations

- The equilibrium V^* for V increases with α to $V_{Nash} < V^* < V_{hit}$.
- The average costs for the population decrease accordingly.

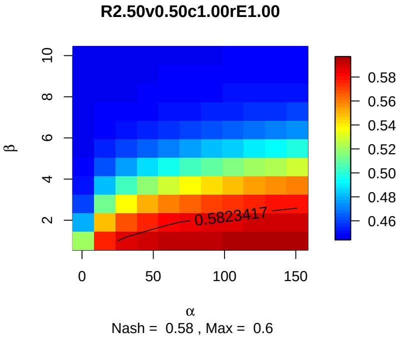

Figure: Dependence of equilibrium V^* on α and β for $R_0 = 2.5$.

Do imitate on occasion and be open-minded!

WJ: Wow! Your findings show that imitation is beneficial for the population, even imitation of randomly chosen strangers, but it should be done only rarely. Most of the time people should rely on their own wits!

DG: This is a nice way of putting it, but in science we would be more cautious about making such sweeping pronouncements.

WJ: More cautious, in what sense?

DG: We found that it is not so much the overall frequency of imitation, but the open-minded way in which you make the decision that gives these high vaccination coverages.

Moreover, in mathematics, we like to obtain confirmation of simulation results by proving rigorous theorems, like this one:

A theorem

Theorem

Assume that the vaccine efficacy $rE = 1$ *and fix any* $V^- < V_{hit}$. *Choose any* $\beta(V^-) > 0$ *large enough such that*

$$1 - e^{-2\beta(V^-)(c_i - c_v)} - \frac{2(1-x^-)}{x^-} e^{-\beta(V^-)(c_i - 2c_v)} > 0.$$

Let

$$\alpha(\beta) > max\{1,\ e^{\beta(c_i - c_v)} + e^{-\beta(c_i - c_v)} - 2e^{\beta c_v} - 2e^{-\beta c_v}\}$$

Then for any $\beta > \beta(V^-)$ *and* $\alpha > \alpha(\beta^*)$ *and initial vaccination coverage* $V_0 \in (0,1)$ *the system will approach an equilibrium* V^* *that satisfies the inequality*

$$V^- < V^* < V_{hit}.$$

Future plans

We have already some preliminary results for the case when $rE < 1$:

- Again, for suitable choices of $\alpha > 1$ and β we obtain $V_{Nash} < V^* < V_{hit}$.
- We have both a theorem about this and simulation results.
- However, the dependence of V^* on α may no longer be monotone. Also, we believe that our theorem tells only part of a more complex picture and are still working on extending it to other regions of the parameter space.

We are also working on more general models that allow for more than two strategies and adaptive procedures other than imitation.

But now, seriously ...

WJ: Seems you have discovered some really interesting things.

Thoroughly enjoyed our conversation, but gotta go now.

Would like to stay in touch and hear more about your future research on this.

But now tell me, seriously:

Should I get that flu shot or not?