Open-minded imitation in vaccination games

Winfried Just, Department of Mathematics, Ohio University

based on joint work with

David Gerberry, Department of Mathematics, Xavier University
and
Ying Xin, Department of Mathematics, Montana State University

Dynamical Systems and Mathematical Biology Seminar,
Ohio University, Athens, Ohio
September 25, 2018

Timeline of this presentation

The Past: Transcript of a recent conversation.

The Present: A very short regular research talk with real theorems.

The Future: A long list of open problems.

In other words, things I really would like to talk about next time, if only I could find new collaborators.

Two people meet by accident ...

WJ: Hi! I'm Vinny!

YX: Hi! I'm Ying Xin, a Ph.D. student in mathematics.

WJ: Oh no! I'm terrible at math!

YX: I find that hard to believe.

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

YX: Would be happy to.

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

Costs

YX: 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?

YX: 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.

YX: 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

YX: 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.

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

YX: 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.

YX: 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 $(1 - rE)x$.

Expected costs for vaccinators and nonvaccinators

YX: 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

YX: 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.

YX: 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.

YX: 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.

YX: 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

YX: 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 vaccination games

YX: 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 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.

YX: Yes, this is what models of the vaccination game predict.

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!

YX: 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.

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

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

YX: 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?

YX: 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?

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

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

YX: 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?

YX: 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.

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

WJ: You didn't notice?

YX: ?? Notice what?

WJ: I might ask an expert, like you.

YX: 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.

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

Imitation of good decisions

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.

YX: This conjecture has been widely studied.

WJ: So what have these studies found?

YX: 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!

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

WJ: What do you mean?

Imitating George

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

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.

YX: 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.

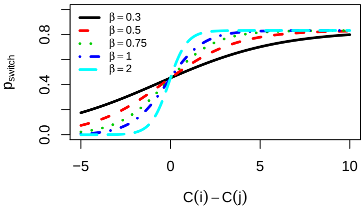

Are Fermi functions realistic?

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

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.

YX: My Ph.D. advisor Prof. Just noticed the same thing.

WJ: I know this guy! He is your advisor? I could tell you stories ...

YX: Maybe not now ...

WJ: OK. Some other time. 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:

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

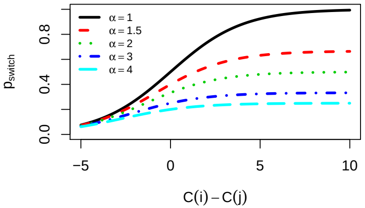

Generalized Fermi functions, continued

WJ: Would that be closer to my way of imitating?

YX: You can think of α^{-1} as the probability of considering to base your decision in a given year on imitation. If α is large, then p_{switch} would always be close to 0.

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

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

WJ: You can call it this way.

The parameter α as a degree of open-mindedness

YX: 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.

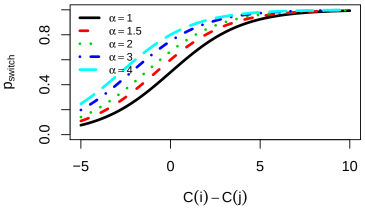

Vaccination games with generalized Fermi functions

WJ: But would that α make any difference in vaccination games?

YX: 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.

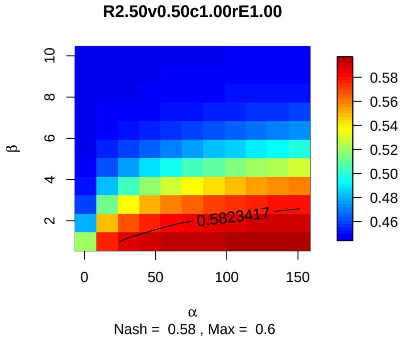

Figure: Dependence of equilibrium V^* on α and β for $R_0 = 2.5$.

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!

YX: 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?

YX: We found that it is not so much the overall frequency of imitation, but the open-minded way of decision-making that gives these high vaccination coverages.

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

Time to say good-bye

WJ: Theorems? Er ...

Thoroughly enjoyed our conversation, but gotta go now.

YX: Same here. Thoroughly enjoyed our conversation, but need to go now on a long trip to Montana.

WJ: Wow! To do mountain climbing?

YX: No, work on a postdoc project.

WJ: Well, good luck with whatever you are going to do over there!

But now tell me: Should I get that flu shot or not?

XY: You should keep an open mind about it.

Research talk: Variables and parameters of our model

- c_v: average cost of the vaccination
- c_i: average cost of infection.
- We assume throughout that $c_i > 2c_v > 0$.
- $0 \leq y \leq x < 1$: probabilities that a vaccinated and an unvaccinated person, respectively, gets infected. These probabilities depend on vaccination coverage and vaccine efficacy.
- Let $0 \leq rE \leq 1$: vaccine efficacy, assumed fixed. When $rE = 0$, then $y = x$, when $rE = 1$, then $y = 0$.
- R_0: basic reproductive ratio of the underlying disease transmission model, assumed fixed.
- n: number of flu season.
- $0 \leq V_n \leq 1$: vaccination coverage in season n.
- We model the change of V_n from season to season with a difference equation model.

Vaccination decisions

At the beginning of each season, each player is assumed to randomly pick one other individual and compare his or her cost in the previous season with the cost of that individual. Then the player will switch to the other's strategy with probability

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

where $\alpha \geq 1$ corresponds to a degree of open-mindedness, and β measures accuracy of perception.

This determines V_n.

The probabilities $x = x(V_n)$ and $y = y(V_n)$ for the subsequent flu outbreak are then calculated from a standard SIR model.

Some minor results

- V_{hit} is the societal optimum.
- $V = 0$ and $V = 1$ are always equilibria.
- The equilibrium $V = 1$ is always unstable.
- The equilibrium $V = 0$ is unstable iff an interior equilibrium $V^* \in (0, 1)$ exists. This can be the case even when $V_{Nash} = 0$.
- When V^* exists, it is always in $(0, V_{hit})$ and is always unique.
- V^* may be stable or unstable, but for sufficiently large α it will always be locally asymptotically stable.

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*

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

A couple of references

The results presented on the previous two slides are preoved in:

Y. Xin, D. Gerberry, and W. Just (2018); Open-minded imitation can achieve near-optimal vaccination coverage. *arXiv:1808.08789*
`https://arxiv.org/abs/1808.08789`

Our model is based on the model of:

F. Fu, D. I. Rosenbloom, L. Wang and M. A. Nowak (2011); Imitation dynamics of vaccination behaviour on social networks. *Proc. R. Soc. B* **278** 42–49 doi:10.1098/rspb.2010.1107

The near future: Generalizations for $rE < 1$

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.

The near future: Models with more strategies

Why restrict ourselves to 2 strategies "vaccinate" and "don't vaccinate"?

We have a more general model where each strategy specifies probabilities of vaccinating based on the history in the previous season, that is, whether the player vaccinated/did not vaccinate and experienced infection or did not.

We have already several results for this type of model.

The near future: What if $c_i < c_v$?

Then nobody should vaccinate, obviously.

But suppose some misguided vaccination policy has been adopted nevertheless, and some proportion of the population is already vaccinated, but not a sufficiently large proportion to achieve herd immunity. We observed that in such situations the most cost-efficient course of action would be to vaccinate even more people to full achieve herd immunity. To throw good money after bad, so to speak.

This fairly simple and rather paradoxical observation apparently has not so far been made in the literature and we plan on investigating precise conditions on when such a situation would occur.

The future: What if R_0 and rE vary from season to season?

Recall our modeling assumptions:

- Let $0 \leq rE \leq 1$: vaccine efficacy, assumed fixed. When $rE = 0$, then $y = x$, when $rE = 1$, then $y = 0$.
- R_0: basic reproductive ratio of the underlying disease transmission model, assumed fixed.

Seasonal flu is caused by different strains of the virus in each season. So in reality R_0 and rE vary over time.

Curiously, while the literature on vaccination games for seasonal infections is extensive, only a couple of related papers consider variability in these parameters. We plan on incorporating such variability into our model and study the resulting predictions. One of our undergraduates, Morgan Balcerek, is getting on board with this project.

Why imitate only one other?

Recall that at the beginning of each season, each player is assumed to randomly pick one other individual and compare his or her cost in the previous season with the cost of that individual. Then the player will switch to the other's strategy with probability

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

Why only one? There are some similar models in the literature where comparisons with the average costs for a larger sample of other hosts are being made. It would be interesting to extend our model in this direction.

Why look at only the last season?

Recall that at the beginning of each season, each player is assumed to randomly pick one other individual and compare his or her cost in the previous season with the cost of that individual. Then the player will switch to the other's strategy with probability

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

Why only the last season? What if we instead compare weighted averages of the costs for several preceding seasons? There are models in the literature that assume such longer-term memory, albeit not ones that are very similar to ours. It would be interesting to extend our model in this direction.

Why imitate a random other?

Recall that at the beginning of each season, each player is assumed to randomly pick one other individual and compare his or her cost in the previous season with the cost of that individual. Then the player will switch to the other's strategy with probability

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

Wouldn't it be more realistic to assume that one imitates a close friend, like George, rather than a randomly chosen stranger? Then we would need to consider both imitation and disease transmission on (not necessarily the same) contact networks. There is a large literature on this type of models. It would be interesting to extend our model in this direction.

Why imitate in the first place?

In almost all of the literature on the subject, it is assumed that strategies get modified either by rational calculation of the expected Nash equilibrium or by imitation, or some mixture of these two adaptive procedures. Is there another way that might work better?

We have designed such an option and called it "unhappiness minimization." Preliminary simulations indicate that it seems to work better than rational choice and imitation, even open-minded imitation.

Can one prove that it works better?

Or determine conditions on the model parameters when it works better?

And if so, why does it work better?

Any other ideas?

Many more potential directions for follow-up work can be found in

Y. Xin, D. Gerberry, and W. Just (2018); Open-minded imitation can achieve near-optimal vaccination coverage. *arXiv:1808.08789* `https://arxiv.org/abs/1808.08789`

What other directions would you suggest?