Games and Germs: A Playful Introduction

Winfried Just
Department of Mathematics
Ohio University

Undergraduate Mathematics Seminar
Ohio University
September 6, 2018

What have you been up to lately?

Bob: What have Ying and you been up to since she defended her dissertation?

WJ: We have been working on problems in behavioral epidemiology.

Bob: What does that mean?

WJ: Epidemiology studies how pathogens, colloquially know as germs, spread in populations of humans, animals, or plants.

Bob: That much I know.

WJ: Behavioral epidemiology studies how people make decisions about adopting control measures, such as vaccination, that can prevent or at least limit the spread of pathogens. It uses mathematical models to understand the likely impact of these decisions.

Bob: Interdisciplinary research with high societal relevance! Perfect fit with the core area *Health and Well-being* of Ohio University's mission! And excellent potential for external funding, I presume?

WJ: You can look at it this way.

Games??

Bob: What mathematical tools do you use for this research?

WJ: Game theory.

Alice: That sounds like a lot of fun!!!

Bob: (Laughs.) Yeah. But now, seriously?

WJ: Seriously.

Bob: How can you play games with epidemics, which can literally be matters of life and death?

WJ: Let's look at games from a mathematical perspective.

Alice: I play a lot of games! I can tell you all about games!!

Bob: Alice!
What did Mom tell you about conversations between grown-ups?

AI about games

WJ: Relax, Bob! Alice can tell us about games, and I will translate it for you into more abstract, or if you will, grown-up language.

Al: I play rock-paper-scissors, chess, poker, soltaire, monopoly, Rubik's cube ...

WJ: And you like playing all of them?

Al: Solitaire—not so much. And Rubik's cube even less.

WJ: And why not?

Al: Because what I really like is playing with other people. Now that I think about it, Solitaire and Rubik's cube perhaps aren't really games at all.

WJ: Would you then say that a "real" game involves interactions between two or more players?

Al: Yeah, this is what I meant.

WJ: We mathematicians can use such games as models for all sorts of real-world situations that involve interactions between people.

Payoffs

WJ: And what else do you like about games?

AI: Winning! I really like winning!! But I don't like losing.

WJ: Are winning or losing the only possible outcomes of a game?

AI: No. In poker or monopoly, you can win or lose by a lot.

WJ: So could we then always think of the outcome of a game as a vector of real numbers that represent each player's payoff?

AI: Maybe. ... In poker with three players, for example, if each player initially puts in 5 chips, and at the end player 1 has 7 chips, player 2 has 8 chips, and player 3 is broke, I would write the outcome as $(2, 3, -5)$.

WJ: This is what we mathematicians call a vector. But how about chess? Here the possible outcomes are "win," "lose," or "draw."

AI: If player 1 wins, we could write the outcome as $(1, 0)$, if player 2 wins, we could write it as $(0, 1)$, if the game ends in a draw, we could write it as $(0.5, 0.5)$. There we have your "vectors."

Maximizing payoffs

WJ: What are the players after?

Al: Each player wants to make her payoff as large as possible.

Bob: Alice! Watch your language! What did Mom and I tell you?

Al: I meant to say, "her or his payoff." Would you mathematicians say: "Each player wants to maximize his or her payoff?"

WJ: Exactly. But your original phrase is more interesting. You said make. How do you "make" your payoff as large as possible?

Al: That really depends on the game, you know.

Bob: We know.

WJ: Can you explain it with an example?

Alice explains the Rubik cube

Al: The Rubik cube has a lot of different configurations. In most of them, the colors are scrambled up, but in one of them, the winning configuration, each side has a single color. We can change configurations by turning one of the nine layers of the cube either left or right, which allows us to make one of 18 moves in each configurations.

WJ: In games like chess or checkers, for example, the configurations would be called "positions."

Al: If I can get from a scrambled configuration to the winning one in a certain number of successive moves, I win, with payoff, say 1; if I don't, I lose, with payoff 0.

Now I maximize my payoff, as you would say, by always making the move that is right for the given configuration or position.

Strategies

WJ: And how do you choose the "right" move for a given configuration?

Al: Oh, there is an algorithm for this. I can show you how it works.

Bob: I'm really proud of you, Al, but maybe some other time ...

Al: But this is really easy! I even taught it to my computer!!

WJ: In game theory, such an algorithm would be called a strategy. We use this word even when the algorithm prescribes moves that are not the best or the right ones for a given configuration.

You said that playing with Rubik's cube is not much fun?

Al: Because it is just like following a recipe. It always works.

WJ: I thought you like winning?

Al: I do!!! But this gets so utterly predictable and boring.

How about solitaire?

WJ: When you play solitaire, do you also follow a strategy that makes your payoff as large as possible?

AI: Yes. I even figured it out myself!!

WJ: So you also follow a recipe, but the game is less boring than Rubik's cube?

AI: Yes, because I don't always win, in the sense of putting all the cards on ordered stacks.
It all depends on how the cards are shuffled.

WJ: We mathematicians would say that the next configuration is drawn randomly from a certain probability distribution.

AI: This sounds like a fancy way of saying that the game involves some chance events.

WJ: You said it succinctly and very well.

Expected payoffs

WJ: Isn't it the case though that in some rounds of solitaire that you lost by random chance you would have won with a different strategy? So that, in a way, you regret your moves?

AI: Yes, this happens sometimes.

WJ: How then would your strategy maximize your payoff?

AI: I didn't say my strategy gives always the highest payoff.

WJ: So what did you mean then?

AI: If I play many, many rounds of solitaire I will have fewer such regrets than with any other strategy.

WJ: So you were talking about the average or expected payoff that you get when you average the payoffs for many rounds.

AI: Exactly! We should have said earlier that the goal of each player is to maximize her or his expected payoff.

On to the really fun games

WJ: You said that solitaire is more fun than Rubik's cube ...

AI: That's because it is less predictable. Sometimes I win, sometimes I lose, even with the best strategy.

Almost like in real games, I mean, in games with other players.

WJ: Only almost?

AI: In real games my expected payoff, does not only depend on my moves, but also on the moves of all other players. The really fun thing about games is that each player tries to outsmart all the others so as to get maximum payoff for her- or himself.

Bob: That doesn't sound like a nice goal to strive for. Wouldn't it be better if players cooperated so as to ensure maximum overall expected payoff for the group?

WJ: We will talk about that later. But first let's go from "moves" to "strategies."

Are the really fun games really that much fun?

WJ: When you play chess, don't you also follow a strategy?

AI: Yes I do!

WJ: And wouldn't you assume that your opponent does the same?

AI: Sure. At least, if my opponent is rational and plays so as to maximize her or his expected payoff.

WJ: Wouldn't then the game boil down to pitting your strategy against the one of your opponent, and the better strategy would win or at least ensure a draw? So that the actual game becomes rather more like Rubik's cube than solitaire?

AI: That would only be true if both my opponent and I teach our strategies to our smartphones and then let the smartphones play.

This would be fun to watch!

But only once. Then it would get totally boring.

WJ: But what would be different for human players?

AI: We people make mistakes. And even if we follow, by and large, a certain strategy, we often make a random choice between two or more moves that appear roughly equally good. So the course of the game is not entirely determined by the strategies.

WJ: You appear to be thinking of strategies of humans as broad conscious outlines that leave some leeway for how the actual sequence of moves would be chosen, perhaps subconsciously?

AI: Yeah, sort of.

WJ: But could we agree to call the broad outline, together with whatever causes the actual choices of moves, "strategies?"

AI: O-OK. Fine with me.

WJ: But then, once the strategies are chosen, we know the resulting payoffs, and the actual sequence of moves is no more interesting than when you play the Rubik cube by recipe.

Choosing strategies

Al: (With a lump in her throat.) But what happened to the fun?

WJ: The whole fun is in the player's choosing their strategies.

Al: How could that be the fun part? Wouldn't this have to happen before the players even start playing and make any moves?

WJ: Yes. ... But can you think of a fun game that is practically over once the players choose their strategies?

Al: Rock-paper-scissors!! This is a lot of fun. You see, when I play it with my friend Ronnie, I choose paper, because he always chooses rock, when I play it with my friend Sequi, I choose rock, because she will choose scissors, and when I play it with my friend Paul, I choose scissors, because he prefers paper.

This way I outsmart them all!

WJ: All of them??

How about Maxi?

AI: Except for Maxi, who is really, really smart.

I think she must be able to read my mind, at least most of the time. But I cannot figure out hers.

And her real name isn't Maxi, but very weird.

WJ: Is it Maximinia?

AI: How did you guess that??? Do you know her?

WJ: Not personally. But in an abstract sort of way, I do.

AI: You mathematicians with your abstractions!

Maxi is a wonderful person and my bestest friend!!!

But how come I never can beat her in any game??

WJ: I will show you how she does it.

Let's go back to chess first.

Alice and Maxi play chess

WJ: Think of Maxi (white) choosing between 4 possible chess strategies SM_1, SM_2, SM_3, SM_4 and you (black) between 4 strategies SA_1, SA_2, SA_3, SA_4. When each of you follows her chosen strategy, the payoff vectors will be as in the table below.

Table: $(1, 0)$—Maxi wins; $(0, 1)$—Alice wins, $(0.5, 0.5)$—draw.

	SA_1	SA_2	SA_3	SA_4
SM_1	$(0, 1)$	$(0, 1)$	$(0, 1)$	$(0.5, 0.5)$
SM_2	$(0, 1)$	$(0.5, 0.5)$	$(1, 0)$	$(0.5, 0.5)$
SM_3	$(1, 0)$	$(1, 0)$	$(0.5, 0.5)$	$(0.5, 0.5)$
SM_4	$(0, 1)$	$(1, 0)$	$(0.5, 0.5)$	$(0.5, 0.5)$

WJ: Which strategy would you pick?

Which strategy would you pick?

Table: $(1,0)$—Maxi wins; $(0,1)$—Alice wins, $(0.5,0.5)$—draw.

	SA_1	SA_2	SA_3	SA_4
SM_1	$(0,1)$	$(0,1)$	$(0,1)$	$(0.5,0.5)$
SM_2	$(0,1)$	$(0.5,0.5)$	$(1,0)$	$(0.5,0.5)$
SM_3	$(1,0)$	$(1,0)$	$(0.5,0.5)$	$(0.5,0.5)$
SM_4	$(0,1)$	$(1,0)$	$(0.5,0.5)$	$(0.5,0.5)$

Al: I would pick strategy SA_1 because it beats three of the four strategies of Maxi and maximizes my expected payoff.

WJ: But wouldn't your payoff depend on what Maxi does?

Al: Oh, yeah! She is really smart and can read my mind. So she would pick SM_3, and then I would lose and regret my choice. I really don't like losing, you know.

Which strategy would you pick, then?

Table: $(1,0)$—Maxi wins; $(0,1)$—Alice wins, $(0.5,0.5)$—draw.

	SA_1	SA_2	SA_3	SA_4
SM_1	$(0,1)$	$(0,1)$	$(0,1)$	$(0.5,0.5)$
SM_2	$(0,1)$	$(0.5,0.5)$	$(1,0)$	$(0.5,0.5)$
SM_3	$(1,0)$	$(1,0)$	$(0.5,0.5)$	$(0.5,0.5)$
SM_4	$(0,1)$	$(1,0)$	$(0.5,0.5)$	$(0.5,0.5)$

Al: OK, I will pick strategy SA_3 because then Maxi cannot beat me with strategy SM_3, and I can still win if she plays strategy SM_1. But that would be really dumb of her.

WJ: Wouldn't then Maxi regret her choice of SM_3 and play SM_2 instead?

Al: Yes, of course! She is sooo smart and can read my mind. And she too likes winning. So I would lose again.

Which strategy should you pick, then?

Table: $(1,0)$—Maxi wins; $(0,1)$—Alice wins, $(0.5,0.5)$—draw.

	SA_1	SA_2	SA_3	SA_4
SM_1	$(0,1)$	$(0,1)$	$(0,1)$	$(0.5,0.5)$
SM_2	$(0,1)$	$(0.5,0.5)$	$(1,0)$	$(0.5,0.5)$
SM_3	$(1,0)$	$(1,0)$	$(0.5,0.5)$	$(0.5,0.5)$
SM_4	$(0,1)$	$(1,0)$	$(0.5,0.5)$	$(0.5,0.5)$

Al: Oh, I see!! I need to pick strategy SA_4! Then I will have no regrets when Maxi picks SA_3, and Maxi will have no regrets either in this case, because she could do no better with any other of her strategies against my SA_4.

WJ: A choice of strategies where no player has any regrets about his or her choice given the choices of all other players is called Nash equilibrium.

Al: Just to make sure: By "no regrets" you mean that no player could achieve a higher expected payoff by switching to another strategy?

WJ: Exactly.

Al: But tell me, can Maxi really read minds?

Can Maxi really read minds?

WJ: Yes, at least minds of very, very smart people like Alice and herself. She assumes that all players are perfectly rational and as smart as she is and will only pick strategies from a Nash equilibrium. Then so does she.

Al: How can she reason all that out?

Table: (1, 0)—Maxi wins; (0, 1)—Alice wins, (0.5, 0.5)—draw.

	SA_1	SA_2	SA_3	SA_4
SM_1	$(0, 1)$	$(0, 1)$	$(0, 1)$	$(0.5, 0.5)$
SM_2	$(0, 1)$	$(0.5, 0.5)$	$(1, 0)$	$(0.5, 0.5)$
SM_3	$(1, 0)$	$(1, 0)$	$(0.5, 0.5)$	$(0.5, 0.5)$
SM_4	$(0, 1)$	$(1, 0)$	$(0.5, 0.5)$	$(0.5, 0.5)$

WJ: Maximinia may simply look at her minimum payoff in each row of the payoff matrix and then choose a strategy for which this minimum in the corresponding row is maximal.

Al: Ah! That's how you knew her! But if there is no Nash equlibrium?

WJ: Every game has at least one Nash equilibrium.

How about rock-paper-scissors?

Al: But that's not true!! Not when I play rock-paper-scissors with Maxi! Look at the payoff matrix:

Table: $(1, 0)$—Maxi wins; $(0, 1)$—Alice wins, $(0.5, 0.5)$—draw.

	RA	PA	SA
RM	$(0.5, 0.5)$	$(0, 1)$	$(1, 0)$
PM	$(1, 0)$	$(0.5, 0.5)$	$(0, 1)$
SM	$(0, 1)$	$(1, 0)$	$(0.5, 0.5)$

Al: You see, Maxi's minimum for each row is zero.

WJ: So Maxi doesn't have a strategy in your table that she couldn't possibly regret,

Al: and neither have I. So there is no Nash equilibrium.

WJ: Not if we consider only the strategies in this table.

Al: But by the definition of the rock-paper-scissors game, as you mathematicians would say, the table lists all possible strategies!!

WJ: You said that you can outsmart Ronnie, Sequi, and Paul, because you know their favorite strategies. But don't they catch on and switch sometimes?

Al: Not Ronnie. He always plays rock. He is sooo ...

Bob: (Frowns.)

Al: ... stubborn. But Sequi switches all the time. She always goes *SPRSPRSPR* .... So I go *RSPRSPRSP* ... and I always win.

Paul is less predictable. He flips a coin, and when it comes up heads, then he plays P. If it comes up tails, he flips another coin, and then plays S when that one comes up heads and R if that one comes up tails.

WJ: We mathematicians would say that Ronnie plays a pure strategy while Sequi and Paul play mixed strategies.
The table on the previous slide lists only the pure strategies.

Al: Sequi's mixed strategy follows a fixed pattern though that I can outsmart.

WJ: But Paul randomizes his choices, and you cannot always win.

Can Alice outsmart Paul?

Al: Not always. But I often challenge him to a match where he plays his coin-flipping strategy for 100 rounds and I play "scissors" all the time. Every time we do this he loses more rounds than I do.

WJ: It seems that the expected payoff for your strategy S is higher when matched against his coin-flipping strategy, which we might call C.

Al: You bet!! The first coin tells him to play P about half of the time, so I win the round with probability 0.5, as you mathematicians would say. And I lose only when both coins come up tails, with probability 0.25, which means about one quarter of the time.

WJ: So the expected payoff E for S against C would be:

$$
\begin{aligned}
E &= (1)Pr(P) + (0.5)Pr(S) + (0)Pr(R) \\
&= (1)(0.5) + (0.5)(0.25) + (0)(0.25) = 0.625.
\end{aligned}
$$

Al: Which is larger than the payoff 0.5 for breaking even. This means that I can outsmart the randomized strategy!!

How would Maxi play?

WJ: You said the randomized strategy. Is there only one?

AI: Now that I think of it, Paul could perhaps use a fair die instead of coins and play R if it comes up 1 or 2, S if it comes up 3 or 4, and P if it comes up 5 or 6. This is another randomized strategy.

WJ: Let's call it D. In this strategy, Paul's choices are drawn from the uniform probability distribution on $\{R, P, S\}$.

AI: If Paul plays D and I play S, them my expected payoff is:
$E = (1)Pr(P) + (0.5)Pr(S) + (0)Pr(R)$
$= (1)(1/3) + (0.5)(1/3) + (0)(1/3) = 0.5$, so we brake even!

WJ: Would Maxi perhaps play D?

AI: Oh no!! Maxi is way too smart for that.
If I play S, she would outsmart me, play R, and win!

WJ: Couldn't you prevent her from outsmarting you by playing D yourself?

AI: Yes!!! She can read my mind, but not the die's mind!!

Nash equilibrium for rock-paper-scissors

WJ: Look at the payoff matrix when we include C and D:

Table: (Maxi's payoff, Alice's payoff).

	RA	PA	SA	CA	DA
RM	$(0.5, 0.5)$	$(0, 1)$	$(1, 0)$	$(0.375, 0.625)$	$(0.5, 0.5)$
PM	$(1, 0)$	$(0.5, 0.5)$	$(0, 1)$	$(0.5, 0.5)$	$(0.5, 0.5)$
SM	$(0, 1)$	$(1, 0)$	$(0.5, 0.5)$	$(0.625, 0.375)$	$(0.5, 0.5)$
CM	$(0.625, 0.375)$	$(0.5, 0.5)$	$(0.375, 0.625)$	$(0.5, 0.5)$	$(0.5, 0.5)$
DM	$(0.5, 0.5)$	$(0.5, 0.5)$	$(0.5, 0.5)$	$(0.5, 0.5)$	$(0.5, 0.5)$

Al: So when we both play D, neither of us has any regrets,

WJ: and (D, D) is a mixed strategy Nash equilibrium.

Al: But wait!! You said earlier "for the strategies in the table."
Aren't there more randomized mixed strategies?

WJ: Infinitely many, in fact. But none of them beats D.

Al: But to be sure, wouldn't you need to check each one of them?

WJ: Mathematicians can reason about infinitely many strategies at once.

Alice wants to know more

Alice: This sounds so cool!!! Can you show me how?

WJ: Not now. In a few years, if you decide to study math ...

Al: But I want to become a game designer!!

WJ: And for that a math degree is very useful.

Bob: (Alarmed) Listen Alice: Mom and I love you very much and want you to have a great career in business or medicine.

WJ: Game theory actually grew out of applications to economics.

Al: And Dr. Just is going to tell us about games in epidemiology, which is really, really similar to medicine!

Costs and benefits of vaccinations

WJ: Now consider an outbreak of an infectious human disease. Not everybody will become infected, but each of those who do will bear a cost C_i of the infection, which is like a negative payoff.

Assume, moreover, that there is a vaccine, and those who get vaccinated prior to the outbreak will not become infected. Each of them will bear a cost of C_v though, which may reflect the price and inconvenience of getting vaccinated, as well as possible side-effects.

Bob: But infectious diseases can be deadly!
So C_v is really negligible compared with C_i.

WJ: Not always. Think about the flu.

Al: Getting the flu is nasty. But I don't like getting jabbed with a needle either. Getting jabbed once or even five times is better than getting the flu though. But getting jabbed 100 times would be worse.

WJ: How about getting jabbed 10 times? Would this be roughly as bad as getting the flu once?

Al: That sounds about equally painful. So let's assume that $C_i = 10C_v$.

Societally optimal vaccination coverage

Bob: But C_i is still much higher than C_v, so everybody should get vaccinated.

WJ: Shouldn't the optimal value v_{opt} of the vaccination coverage v, that is, of the proportion of people in the population who get vaccinated, be such that the total average cost per person is as small as possible?

Bob: Granted. But if we let x denote the probability that an unvaccinated person gets infected, then this average cost is

$E(C) = vC_v + (1-v)xC_i = xC_i + v(C_v - xC_i)$.

As long as $|xC_i| > |C_v|$, which seems a very reasonable assumption, we can always improve $E(C)$ by increasing v.

So vaccinating the entire population must be optimal.

WJ: "As long as" is the key phrase here.

The probability x is a decreasing function $x(v)$ of v.

Bob: But how can x depend on v? If we vaccinate half of the other people but not our Alice, she will remain unprotected and vulnerable to the infection!

Al: But not with the same probability!! Look, Daddy: I can become infected only by catching the flu from some other person, but not from a vaccinated other person. So if many other people get vaccinated, then I will be less likely to catch the flu.

WJ: I couldn't have explained it any better! This function actually reaches 0 for some value $v_{hit} < 1$, which is called the herd immunity threshold.

Al: Do you mean, when a proportion of v_{hit} of people get vaccinated, my chances of getting the infected are zero, regardless whether I got jabbed or not?

Whom should we vaccinate?

WJ: Exactly. Vaccinating a larger proportion than v_{hit} of the population would only increase the cost to the population, without increasing the amount of protection against the disease.

We can see that $v_{opt} \leq v_{hit} < 1$; under some additional assumptions one can even show that $v_{opt} = v_{hit}$.

Al: But how should we select the people who will get vaccinated? Should the government draw up a list?

WJ: That would create a lot of resentment and not everybody would comply.

Al: Right! Nobody likes to get jabbed.

Bob: And who would want to leave their own child unprotected?

The vaccination game

Al: Oh—I get it!! We simply let everybody make their own decision independently, based on rational self-interest.

WJ: This is called the vaccination game. Its mixed-strategy Nash equilibrium will give a vaccination coverage v_{Nash}.

Al: And these individual rational decisions will lead to optimal vaccination coverage for the whole society!

WJ: Unfortunately, no. In this game we have $v_{Nash} < v_{opt} = v_{hit}$, so that more people will suffer from the infection and the overall cost will be higher than for the societal optimum.

Al: Why???
Why could what's best for each of us not be best for all of us???

WJ: If you could assume that individual rational decisions by **everybody else** give herd immunity, would it be rational for **you** to get jabbed?

Al: No, of course not!

WJ: And this would be true for **everybody else.** Thus $v_{Nash} < v_{hit}$.

Is there some hope?

AI: Bummer! Can you mathematicians do anything about it?

WJ: As a society, we can try to find ways to encourage more people to voluntarily get vaccinated.

Mathematicians can help by building more realistic models that take into account how people really make decisions.

AI: So the vaccination coverage v_{Nash} will be reached only if everybody is as smart as Maxi?

WJ: You have put it perfectly!

AI: And are you saying that since most people aren't all that smart, there is some hope?

WJ: I would not say it this way. For example, people often arrive at good decisions by imitating the behavior of successful others.

AI: So, if people make their decisions by imitation, rather than rational calculation, will this lead to more optimal vaccination coverage?

WJ: Excellent question! A lot of mathematical literature basically says imitation would give even worse coverage than rational choice.

So what are we up to?

Al: Bummer again!

WJ: But pretty much all of this literature assumes that imitation can be modeled by so-called Fermi functions. There seems to be no particular reason for modeling imitation only in this way, other than precedence in earlier literature.

Al: Do you mean, some author used them and then all the next authors imitated this assumption?

WJ: Yes, this is what I meant. Some empirical studies indicate we should take into account one more parameter in modeling imitation.

Al: What would this parameter signify and what does it do to the vaccination coverage?

WJ: In our recent work with my former Ph.D. student Dr. Ying Xin and Dr. David Gerberry of Xavier University we conceptualized this parameter as a degree of "open-mindedness." We showed that when this parameter is sufficiently large, the vaccination coverage will become much larger than v_{Nash}, very close to the optimum v_{hit} in fact.

Where can you learn more?

AI: This sounds so exciting!!

Where can I learn more about your work?

WJ: I will give soon a presentation about this work in our Dynamical Systems and Mathematical Biology seminar.

If you are interested, please asked me to add you to the mailing list for this seminar.

The first results of our work also have been recently posted as a preprint:

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`

On the characters of this play

Bob and Alice are purely fictional characters.

Neither of them is based on any one actual person.

The character of WJ . . .

Alice: We know. Let's get out of here and play some games!!