A decision-theory based model of HPV vaccination and infection dynamics

Alison Elgass, Mentored by Marisa Eisenberg

1. Introduction

HPV, or *human papillomavirus*, is the most prevalent sexually transmitted infection (STI) in the United States, with an estimated 79 million Americans infected [1]. There are over 150 different types of HPVs, some which affect the skin and cause common warts, and others which affect the genital area and can be transmitted only through sexual contact. Within the genital HPVs there is a further breakdown of types into low-risk and high-risk. Most HPV infections are harmless and clear on their own; in fact, most sexually active people will become infected at some point and not even know it. But in some cases HPV can manifest itself and cause symptoms such as genital warts, which are caused by low-risk types, or it can lay dormant and later develop into cancer, which is primarily caused by high-risk types [2].

Cervical cancer is the most common HPV-related cancer in women. Each year in the U.S., about 11,700 new cervical cancers are diagnosed, and it is estimated that 90% of these cancers are caused by an HPV infection. Several other genital cancers are often caused by HPV, including anal, rectal, penile, vaginal, and vulval. In addition, 70% of cancers of the oropharynx (back of the throat) are caused by HPVs. Oropharyngeal cancer is the most prevalent HPV-related cancer overall, and these cancers actually appear predominantly in men [3]. In total, about 40,000 cancers are diagnosed every year in sites where HPV is found, and HPV causes about 31,500 of these [4].

A vaccine for HPV was approved in 2006, with 2 or 3 doses required for it to be fully effective [5]. The vaccine protects against the most common high-risk HPV types, which together account for about 90% of HPV-related cervical cancers, as well as the majority of other HPV-related cancers and genital warts [6]. Initially the vaccine was only recommended for girls, beginning at ages 11-12. Then in 2009 the FDA began recommending that boys also receive the vaccine at the same age [5].

Unfortunately, the HPV vaccine has seen lower uptake rates than other child and adolescent vaccines in the United States. As of 2016, the CDC reports that 60% of teens ages 13-17 had received at least 1 does of the HPV vaccine, and only 43% were fully vaccinated. The numbers are slightly higher among females (65% received $\geq$1 dose and 50% were fully vaccinated) compared to males (56% received $\geq$1 dose and 38% fully). In contrast, 88% of teens have received a tdap (tetanus, diphtheria, and acellular pertussis) vaccine, and 82% have received a meningococcal vaccine, both of which are also recommended at ages 11-12. Childhood vaccines have even higher uptake rates, with over 90% of teens vaccinated against chickenpox, hepatitis B, and MMR (measles, mumps, and rubella) [7].

Experts have speculated about the reasons for this discrepancy in vaccination rates. Proposed explanations include the relative newness of the vaccine, the association with sexually-transmitted disease, and the general anti-vaccine movement, among others [8]. We believe that many factors play a role in the vaccination decision, and that these factors are largely influenced by parents' perceptions of the disease and the vaccine itself. Thus instead of studying one possible reason, in this paper we aim to construct a more holistic model of the

vaccination decision, using the relative costs of vaccination and infection. Our two main goals are: 1) To use decision theory and game theory to predict parents' decisions of whether or not to vaccinate their child against HPV. We draw on the work of Bauch & Earn in their paper on the application of game theory to childhood vaccination [9]. We use their method to calculate the Nash equilibria of our vaccination "game," in which parents decide whether to vaccinate or not based on the relative expected payoffs of vaccination versus infection, and the prevalence of the disease in the population. We then group parents into three categories based on their preconceived opinions on HPV vaccination (preconceived here meaning before they take their child to the doctor) and, using results from an extensive search of previous literature, choose appropriate parameters for each of the three groups. 2) To couple this decisionary model to a simple SIS model of HPV infection with vaccination and to study the dynamics the combined model produces. Specifically, we wish to observe the back-and-forth interaction between the two models, with the decision model affecting the vaccination rate and thus the disease progression, and, conversely, the disease model affecting parents' decisions via the overall disease prevalence.

2. The vaccination decision model

Our first goal is to create a model of the vaccination decision making process, so that we can produce realistic estimates of the percentage of parents who will choose to vaccinate, dependent upon the parameters we use. We utilize decision theory in our model because it is ideal for modeling situations in which an individual must choose between different options or "strategies" based on their relative payoffs (where payoff here represents the overall benefits and/or harms that result from a particular option). Logically, an individual will want to maximize his expected payoff or, equivalently, minimize his expected loss.

In the context of this paper, the individual or "player" is the parent, and the two strategies are vaccinate (V) or do not vaccinate (NV). That is, when a parent takes his or her child to the doctor,[1] they choose either to have them vaccinated or not. Since this decision takes place in the doctor's office, we must also consider the effect of the doctor on the decision. The doctor also has two options: they can either recommend the vaccine or not.[2]

We chose to construct this initial model using a decision theory (1 player) rather than game theory ($\geq$ 2 players) framework because, although the parent's decision is partially influenced by their doctor, we are only concerned with the parent's chosen outcome and not the doctor's. Additionally, we assume that the doctor does not change his strategy based on the parent's initial preferences or chosen strategy; in other words, any given doctor is either going to recommend vaccination or not, regardless of the specific parent.

[1] We assume in our model that all vaccination decisions are made by the parent of the vaccine-recipient, since although vaccination is approved for boys and girls up to ages 26, most vaccinations occur in children under 18

[2] We define the "no recommendation" category to include doctors who may mention the vaccine but not suggest or advocate for it, doctors who fail to offer the vaccine (either because their office does not stock it or because of their personal beliefs), and doctors who explicitly advise against the vaccine

2.1 Constructing the decision theory matrix

Using this approach, we can create a matrix of the expected payoffs for each of the two strategies, vaccinate or don't vaccinate, and for each whether the parent receives a recommendation from his or her doctor. Thus we obtain the following two-by-two matrix with four possible outcomes:

	Doctor Recommends	Doctor Does Not Recommend
Vaccin-ate	1 $-r_v + \alpha$	2 $-r_v$
Don't	3 $-\pi_p(r_i + \pi_c r_c)$	4 $-\pi_p(r_i + \pi_c r_c)$

Figure 1: Decision theory-based matrix of expected costs

Parameters
r_v = perceived cost/risk of vaccination
r_i = perceived cost/risk of HPV infection
r_c = perceived cost/risk of HPV-related cancer
π_p = probability of contracting HPV, given the level of vaccine coverage in the population is p
π_c = probability of HPV infection developing into cancer
α = "boost" for agreeing with doctor

1. The parent receives a doctor's recommendation and chooses to vaccinate

The expected payoff here is the perceived risk of vaccination, which could include any expected side effects or complications of receiving the vaccine, as well as any other negative consequences, for example if the parent believes that receiving the vaccine will give his or her child a false sense of security and cause him or her to be more sexually risky.[3] We then add a constant alpha term as a "bonus" for the parent agreeing with the doctor. This could be interpreted as the peace of mind that comes with following the doctor's orders, increasing the overall expected payoff (or, more specifically, offsetting some of the negative cost of the perceived risk of vaccinating).

As is the case for all four possible outcomes in the matrix, the expected payoff is negative because it is a cost, not a benefit. Additionally, it's important to note that the risks are as *perceived* by the parents. Therefore different parents will necessitate different perceived risk

[3] This belief was discredited in the paper by Mayhew et al [11], which found that girls' risk perceptions immediately after vaccination were not associated with subsequent sexual risk-taking

parameters, and for some these may not be “accurate” from a public health perspective (for example, some parents might perceive the risk of vaccination to be very high, when in reality the HPV vaccine has been proven safe and its side effects relatively mild [10]). We discuss the issue of choosing appropriate parameters in section 2.4.

2. The parent does not receive a doctor’s recommendation and chooses to vaccinate anyway

The expected payoff here is the same as it is for outcome 1, but without the added alpha term, since in this case the doctor did not recommend the vaccine.

3. The parent receives a doctor’s recommendation but chooses not to vaccinate

The expected payoff here is the probability that the unvaccinated child will contract HPV at some point in the future times the perceived risk of HPV infection, plus the probability that an HPV infection will develop into cancer times the perceived risk of HPV-related cancer.

In this initial framing, we treat the probability of infection, π_p, as a fixed (constant) parameter; however, more realistically π_p changes depending on the transmission dynamics and the vaccination patterns across the population, which we will examine later in the paper.

We also note that we considered adding a negative constant term here as a “penalty” for going against the doctor’s recommendation. However for simplicity sake, we forgo this and note that we can simply adjust alpha to be higher so that it reflects the benefit of choosing to vaccinate given that a doctor recommends it.

4. The parent does not receive a doctor’s recommendation and chooses not to vaccinate

The expected payoff here is the same as it is for outcome 3, since the perceived risks associated with non-vaccination are the same regardless of whether or not the doctor recommends the vaccine.

2.2 Calculating the Nash equilibria

In order to proceed with the analysis of our decision theory matrix, we think of it in terms of mixed strategies. Thus for any given parent, instead of their two strategies being simply vaccinate or don’t vaccinate, the former strategy becomes their probability V of vaccinating, and the latter their probability 1-V of not vaccinating. We also expand the decision theory model to a game theoretic framework, by noting that the probability of infection, π_p, in actuality depends on the vaccination choices of all the other parents in the population as well—as more parents choose to vaccinate, the risk of infection decreases. Thus, the decision theory model can be interpreted as a game theoretic one, where parents in the population choose their strategies depending on what the other parents in the population choose. With this framing, we can calculate the expected payoff for the strategy of vaccinating with probability V, given that the overall level of vaccine coverage in the population is p (so that now π_p is interpreted as a function of p), to be

$$E(V, p) = [V(-r_v + \alpha) + (1-V)(\pi_p(r_i + \pi_c r_c))]*R + [V(-r_v) + (1-V)(\pi_p(r_i + \pi_c r_c))]*(1-R)$$

where R is the fraction of doctors who recommend (or probability of the doctor recommending) the vaccine. Cancelling terms, we obtain

$$E(V, p) = V\alpha R + V(-r_v) + (1-V)(\pi_p(r_i + \pi_c r_c))$$

$$= V(-r_v + \alpha R) + (1-V)(\pi_p(r_i + \pi_c r_c))$$

We can simplify further by creating a term r which represents the relative risk of vaccination (minus the alpha "booster" term) versus infection.

$$r = (r_v - \alpha R) / (r_i + \pi_c r_c)$$

We then scale the equation by dividing out the constant $(r_i + \pi_c r_c)$[4], giving us the following

$$E(V, p) = -rV - \pi_p(1-V) \qquad **$$

Next we pivot away from our decision theory mindset, and we view the system instead as a game where other players' decisions to vaccinate or not affect each other via the overall disease prevalence. This will allow us to calculate a Nash equilibrium for our system. We follow the method of Bauch & Earn in their paper "Vaccination & the Theory of Games" [14].

We assume that some portion ε of population vaccinate with probability V, and the rest (1-ε) vaccinate with probability Q. So the overall vaccine coverage level is

$$p = \varepsilon V + (1-\varepsilon)Q$$

The expected payoff for each strategy follows directly from equation **

$$E_V = -rV - \pi_{\varepsilon V + (1-\varepsilon)Q}(1-V)$$

$$E_Q = -rQ - \pi_{\varepsilon V + (1-\varepsilon)Q}(1-Q)$$

The expected payoff gain of switching strategies from Q to V is just the difference between their individual expected payoffs

$$\Delta E = E_V - E_Q = (\pi_{\varepsilon V+(1-\varepsilon)Q} - r) (V - Q)$$

In order for strategy V to be a Nash equilibrium, ΔE must be positive for all values of the overall vaccination level p. (Note that as p increases, the probability π_p of contracting the disease must

[4] We assume r_i is constant among all parents based on the study by Kahn et al, which found no significant difference in "perceived severity of HPV infection" (a separate category from HPV-related disease) between parents who intended to vaccinate and parents who didn't or were unsure [12]. Similarly, we assume r_c is constant based on the study by Reiter et al, which found no difference among parents in their "perceived severity of cervical cancer" [13]. π_c is constant because the probability of developing cancer (assuming no vaccination) does not vary based on parental beliefs.

decrease continuously, making π_0 the maximum value of π). Intuitively, this means that if V is an equilibrium, then it is always beneficial to switch to strategy V, and then to stay with this strategy. This presents us with 2 cases:

1. $r \geq \pi_0$

Since π_p decreases strictly as p increases, if $r \geq \pi_0$ then $r \geq \pi_p \ \forall p$, so the first term in our ΔE equation will always be negative. Thus for ΔE to be always positive, V-Q must also always be negative, so V must be 0.

This should make logical sense, since $r \geq \pi_0$ means that the relative risk of vaccination outweighs the probability of contracting the disease even at its highest value (when nobody is vaccinated, p=0). If this is the case then it is understandable that nobody will want to vaccinate.

2. $r < \pi_0$

If $r < \pi_0$, then there exists a p* where $\pi_{p^*} = r$, so that $(\pi_p - r) > 0$ for $p < p^*$, and $(\pi_p - r) < 0$ for $p > p^*$. For any $Q < V$, we get $p = \varepsilon V + (1-\varepsilon)Q < V$, and for $V > Q$ we get $p > V$. Thus for ΔE to be positive, the two terms in the equation must change signs at the same point, which occurs only at p*, so V = p*, where p* is the solution to $\pi_{p^*} = r$, is the Nash equilibrium in this case.

2.3 Modeling changing dynamics with the replicator equation

The Nash equilibrium gives us the steady state solution for our vaccination decision game with mixed strategies, when all parameters are held constant, and we don't consider how the vaccine coverage and infection risk vary over time as a function of one another. In real life though, these two factors specifically will change as the system shifts over the course of the game. To deal with this, we introduce the replicator equation, a differential equation for the change over time in the total number of parents choosing to vaccinate. Because we only have two strategies, vaccinate or don't vaccinate, the replicator equation is simply

$$\dot{x}_v = x_v (1-x_v) (f_v - f_{nv})$$

where x_v is the portion of the parent population that chooses to vaccinate, f_v is the fitness (expected payoff) of vaccinating, and f_{nv} is the fitness of not vaccinating

Substituting in the parameters from our decision matrix, we get

$$\dot{x}_v = x_v (1-x_v) [(-r_v + \alpha + \pi_p(r_i + \pi_c r_c))R + (-r_v + \pi_p(r_i + \pi_c r_c))(1-R)]$$

This simplifies to our final replicator equation:

$$\dot{x}_v = x_v(1-x_v)[\alpha R - r_v + \pi_p(r_i + \pi_c r_c)]$$

2.4 Dividing parents into 3 groups & choosing appropriate parameters

As mentioned earlier, different parents perceive different levels of riskiness of vaccination, and we believe this is the primary indicator of their vaccination decision.[5] To reflect this variation in our model, we chose to split the population of parents of vaccine-eligible adolescents into three groups based on their preconceived opinions on the HPV vaccine. We call these groups Pro (P), Neutral (N), and Anti (A), with respect to their beliefs.

In order to choose appropriate parameters for each of these groups to plug into our decision theory model and replicator equation, we conducted an extensive search of literature relating to parental beliefs and opinions on HPV vaccination. The literature ranged in date from 2007 (when the vaccine had just become available) to 2016, and was mostly comprised of surveys of parents. Through this search, we were able to pick out reasonable parameters for the following three areas: the proportion of parents in each of the three belief groups; the relative risk perceptions for each of these groups; and the percent of doctors who recommend the HPV vaccine to their patients.

1. *Proportion of parents in Pro, Neutral, and Anti groups*

In Rosenthal et al, 153 mothers of vaccine-eligible girls were asked whether or not they believed "the HPV vaccine will be safe for my daughter to get." 57% agreed, 10% disagreed, and 33% were neutral [18].

In Fang et al, data from the 2007 HINTS (Health Information National Trends Survey) was filtered to include responses from only parents of girls under the age of 18 (n=1,383). They were asked whether or not they would choose to get the HPV vaccine for their daughters. 58% said yes they would vaccinate, 18% said no they would not vaccinate, and 25% were unsure [19].

Wilson et al focused on young people rather than parents, asking 1,600 college students at two large Texas universities whether or not they would vaccinate *if* they had a daughter in the future. 55% said yes they would absolutely vaccinate, and 20% absolutely would not[6] [20].

Using the findings from these three studies, we estimate the proportion of parents in each group as follows:

55% Pro HPV vaccine
15% Anti HPV vaccine
30% Neutral

2. *Relative Risk Perceptions*

[5] After an extensive literature review, we conclude that demographic factors such as race, ethnicity, education and income level do not affect parents' vaccination decisions [15, 16, 17]. Factors such as religion and vaccination history that do show a correlation with vaccination/non-vaccination are absorbed into our vaccination opinion categories of pro, neutral, and anti.

[6] We guess that this second statistic (the percentage who "absolutely would not vaccinate") may be higher than the actual national average because of the location of the study.

The main parameter we wish to find for each of the three parent groups is r_v, the perceived cost/risk of HPV vaccination. As mentioned earlier (see footnote on page 4), we can assume that the other risk parameters r_i and r_c do not vary among the groups. However, we must choose these parameters to be reasonable in relation to r_v.

We start by drawing from the study by Reiter et al, a survey of 889 parents of adolescent girls in North Carolina. The parents were asked to rate their beliefs on as number of different issues on a scale of 1-4, with 1 being the least and 4 being the most severe. For the sake of simplicity and consistency, we chose to use this scale of 1-4 for all of the risk parameters in our model.

For the belief of the "perceived severity of cervical cancer if daughter got it," the mean score was 3.7. This value was the same for both parents who decided to vaccinate their daughters and parents who decided not to vaccinate [21]. Thus we take r_c, the perceived cost of HPV-related cancer, to be 3.7 for all three parent groups.

For our r_v values, we use data from Lindley et al, an analysis of the 2013 NIS-teen data (n=16,937). Parents of girls and boys ages 10-17 were asked to rate on a scale of 1-10, "is the HPV vaccine safe?" Averaging the mean scores for parents of girls and boys, we find that: parents whose child had received the vaccine reported a mean rating of 8.2, parents whose child had not received the vaccine but said they may in the future had a mean rating of 6.8, and parents whose child had not received the vaccine and said they definitely would not in the future had a mean rating of 4.3 [22].

We take the three groups of parents in this study to be in line with our constructed pro, neutral, and anti categories, respectively. Converting the ratings to our 1-4 scale using linear interpolation (see appendix), we get the following parameters:

r_v for Pro group = 1.6
r_v for Neutral group = 2.1
r_v for Anti group = 2.9

When we looked into r_i, the risk of HPV infection alone (ignoring any possible subsequent cancers), we found that this risk/cost is perceived to be minimal among all parents [23]. Thus we set r_i, for now, to the minimum value of 1 for all three parent groups.

Lastly we calculate π_c, the probability that an HPV infection will later develop into cancer, as follows: The CDC reports that HPV causes 31,500 cancers per year [24]. The total number of Americans infected with HPV is currently about 79 million. Thus the probability of an HPV-positive individual developing HPV-caused cancer *in any given year* is 31,500 / 79 million = 0.04%. Multiplying this by 25 years (an approximate window of likelihood for developing HPV-related cancer), we get 1%, our π_c parameter.

3. *Percent of Doctors Recommending HPV Vaccine*

To obtain the approximate percent of doctors who recommend the vaccine to their patients, we used data from two studies:

The first is a 2014 national survey of 776 physicians by Gilkey et al. Of these doctors, 73% reported that they believe the HPV vaccine to be either “very” or “extremely important” [25].

The second study was a follow-up to the 2010 National Immunization Survey-Teen, a random survey of parents of adolescents ages 13-17. Of the 3,496 respondents, 73% reported that they had received a doctor’s recommendation for the HPV vaccine [26].

We conclude that approximately 73% of doctors recommend the HPV vaccine, while the remaining 27% do not.

2.5 Parent group simulations & the trade-off between r_i and alpha

In order to sufficiently exaggerate the difference between the costs of vaccination, r_v, and the cost of cancer, r_c (which should be quite a bit higher, but is diminished by the low probability π_c of actually developing cancer) so that we obtain realistic results, we scale all the costs r_v, r_c, and r_i from a 1-4 scale to a 1-1,000 scale, again using linear interpolation.

Thus our parameters become:

r_V Pro = 1.6 → 201
r_V Neutral = 2.1 → 367.67
r_V Anti = 2.9 → 634.33
r_c (all groups) = 3.7 → 901
r_i (all groups) = 1 → 1

Just as we chose r_i, we must also choose our alpha value, the “boost” that comes from agreeing with a doctor who recommends vaccination. We set it to 2 (using our 1-4 scale), so that it is slightly higher than the cost of infection but significantly lower than the cost of cancer. When scaled to 1,000 scale, it becomes 334.33.

Plugging these parameters into our replicator equation, we can run the simulation for each of our three parent groups and see the results. We keep π_p, the probability of infection, constant at 0.8, and we start each parent group at 1% vaccination (starting at 0 would not allow our simulations to take off). As we expect, the pro group quickly goes to 1, so that all parents are choosing to vaccinate. Conversely, the anti group goes to 0, so that no parents are vaccinating.

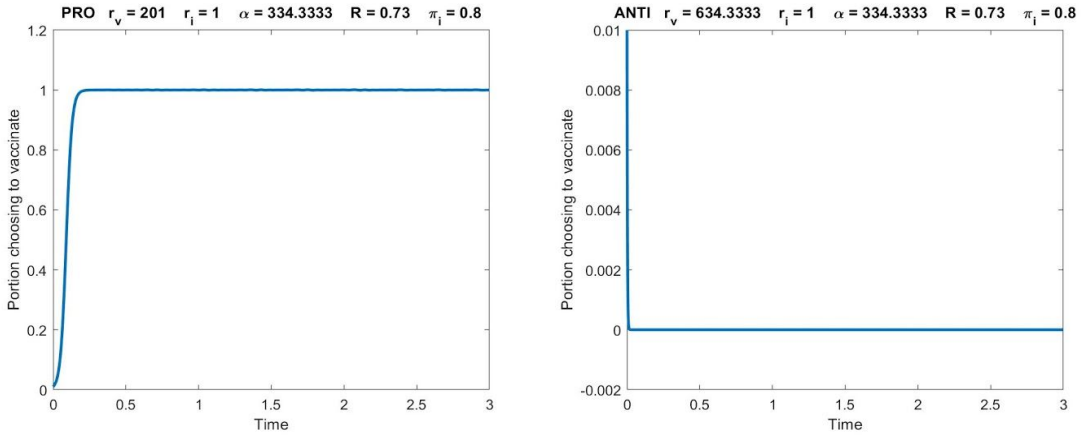

Figure 2: Simulation of parents' vaccination decisions in the pro-group (left) and the anti-group (right), with parameters as listed above

The more interesting simulation happens in our neutral group. In this case, the behavior of the graph and whether it goes to one or to zero depends largely on the values we chose for r_i and alpha. With the parameters as above, the neutral group goes to 0 (no vaccination). However if we increase alpha to 3, the neutral group goes to 1 (complete vaccination). The model thus exhibits a bifurcation as we change alpha, switching from one equilibrium to another.

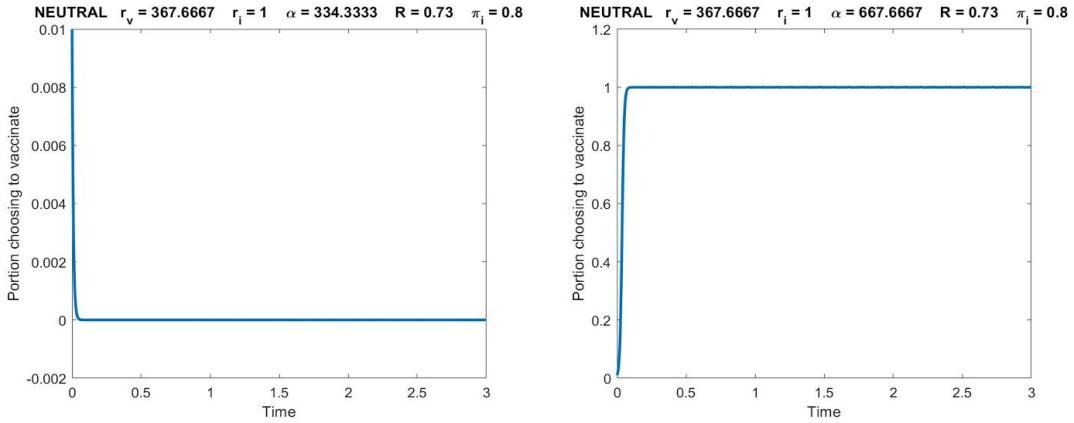

Figure 3: Simulation of vaccination decisions in neutral group for α = 2 (left) and α = 3 (right)

Thus we see that there is some value of alpha between 2 and 3 where this bifurcation happens. It turns out to be approximately 2.479. Below (left) we plot the simulation for the neutral group with an alpha of 2.47 and an alpha of 2.48, and see this divergence occur.

This change doesn't solely depend on alpha, though; we can also play with the value of r_i and see a similar effect. For the neutral group, then, we conclude that there is an important trade-off between r_i and alpha which determines the final behavior of the group, specifically if they will choose to vaccinate or not. We plot a heat map (right) to illustrate this trade-off.

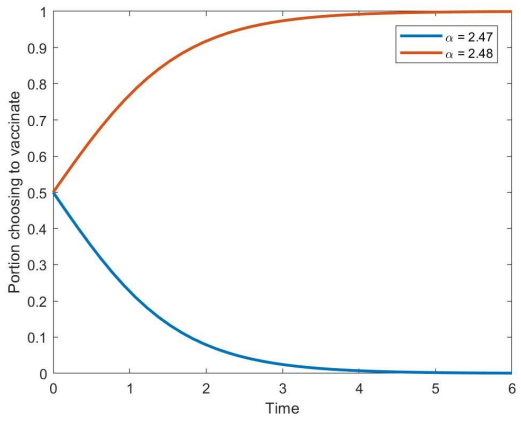

Figure 4: Neutral group vaccination at α tipping point

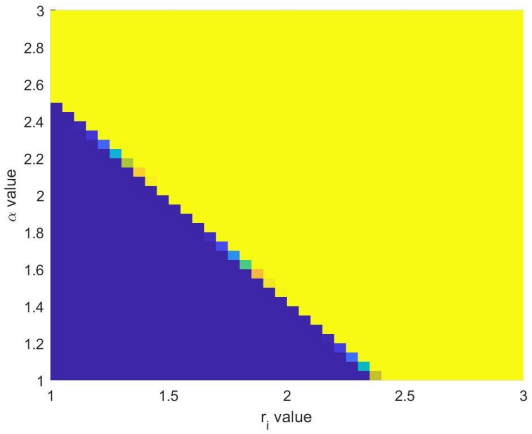

Figure 5: Heat map of vaccination behavior for r_i vs. α

The blue section represents the combination of the two values where the neutral group goes to 0 (no vaccination), while the yellow represents the values where they go to 1 (all vaccination).

This heat map demonstrates the importance of the strength of a doctor's recommendation, since the alpha value must be high enough to overcome a low perceived cost of infection, as in the case of our original choice of $r_i = 1$.

3. Incorporating our decision model into a compartmental disease model

Now that we have our three replicator equations complete with appropriate parameters, we can use them with a basic SIS model with vaccination to study the disease dynamics as parents' vaccination decisions evolve. Our SIS model is

$$\dot{S} = \mu\,(1 - (x_{VP} + x_{VN} + x_{VA})) - \beta SI + \gamma I - \mu S$$
$$\dot{I} = \beta SI - \gamma I - \mu I$$
$$V = \mu\,(x_{VP} + x_{VN} + x_{VA}) - \mu V$$

where μ is the rate of entry and exit from a sexually active state (so that $1/\mu$ is the average lifetime duration of sexual activity), γ is the clearance rate of an infection, β is the contact rate, and the x_V's are the percent who choose to vaccinate in each parent group (so $\max(x_{VP}) = 0.55$, $\max(x_{VN}) = 0.3$, and $\max(x_{VA}) = 0.15$).

We combine these with our three replicator equations, each modified to include the appropriate percentage of the total parent population:

$$\dot{x}_{VP} = x_{VP}\,(0.55 - x_{VP})[\,\alpha R - r_{VP} + \pi_p(r_i + \pi_c r_c)]$$
$$\dot{x}_{VN} = x_{VN}\,(0.3 - x_{VN})[\,\alpha R - r_{VN} + \pi_p(r_i + \pi_c r_c)]$$
$$\dot{x}_{VA} = x_{VA}\,(0.15 - x_{VA})[\,\alpha R - r_{VA} + \pi_p(r_i + \pi_c r_c)]$$

We let π_p = I / (S+I+V), a rough approximation of the probability of infection.[7]

Using the next generation matrix method, we find that $R_0 = \beta / (\gamma + \mu)$.

According to the study by Peyton et al, before the vaccine was available the overall HPV prevalence in the US was about 40% [27]. We start our simulation at these endemic values of $S_0 = 0.6$, $I_0 = 0.4$, and $V_0 = 0$.

We can easily work out that the endemic equilibrium (with no vaccination) occurs when
$S = (\gamma + \mu) / \beta = 1 / R_0$
$I = 1 - (\gamma + \mu) / \beta = 1 - (1 / R_0)$
Using these equations and the baseline prevalence of 40%, we can backwards calculate our R_0 to be approximately 1.67.

We now determine our remaining parameters as follows:

μ: We assume that the average age of sexual debut is roughly 15, and the age of sexual exit (often via marriage/monogamy) is 55, giving an approximate duration of sexual susceptibility & transmission of 40 years. Thus μ = 1/40.

γ: The clearance rate for HPV is typically between 6 months to a year [28], giving us values of 1/0.5 = 2 and 1/1 = 1. We take the midpoint of these values, so γ = 1.5

β: Lastly, we can calculate the contact rate β from our R_0 equation with R_0 = 1.67 and μ and γ as above. We find that β = 2.54.

We run the simulation for alpha values of 3, 2, and 2.5

[7] Although this is not the most technically accurate way to calculate π_p, we believe it is actually a better reflection of a parent's *perception* of the probability of infection, which is what our model really refers to.

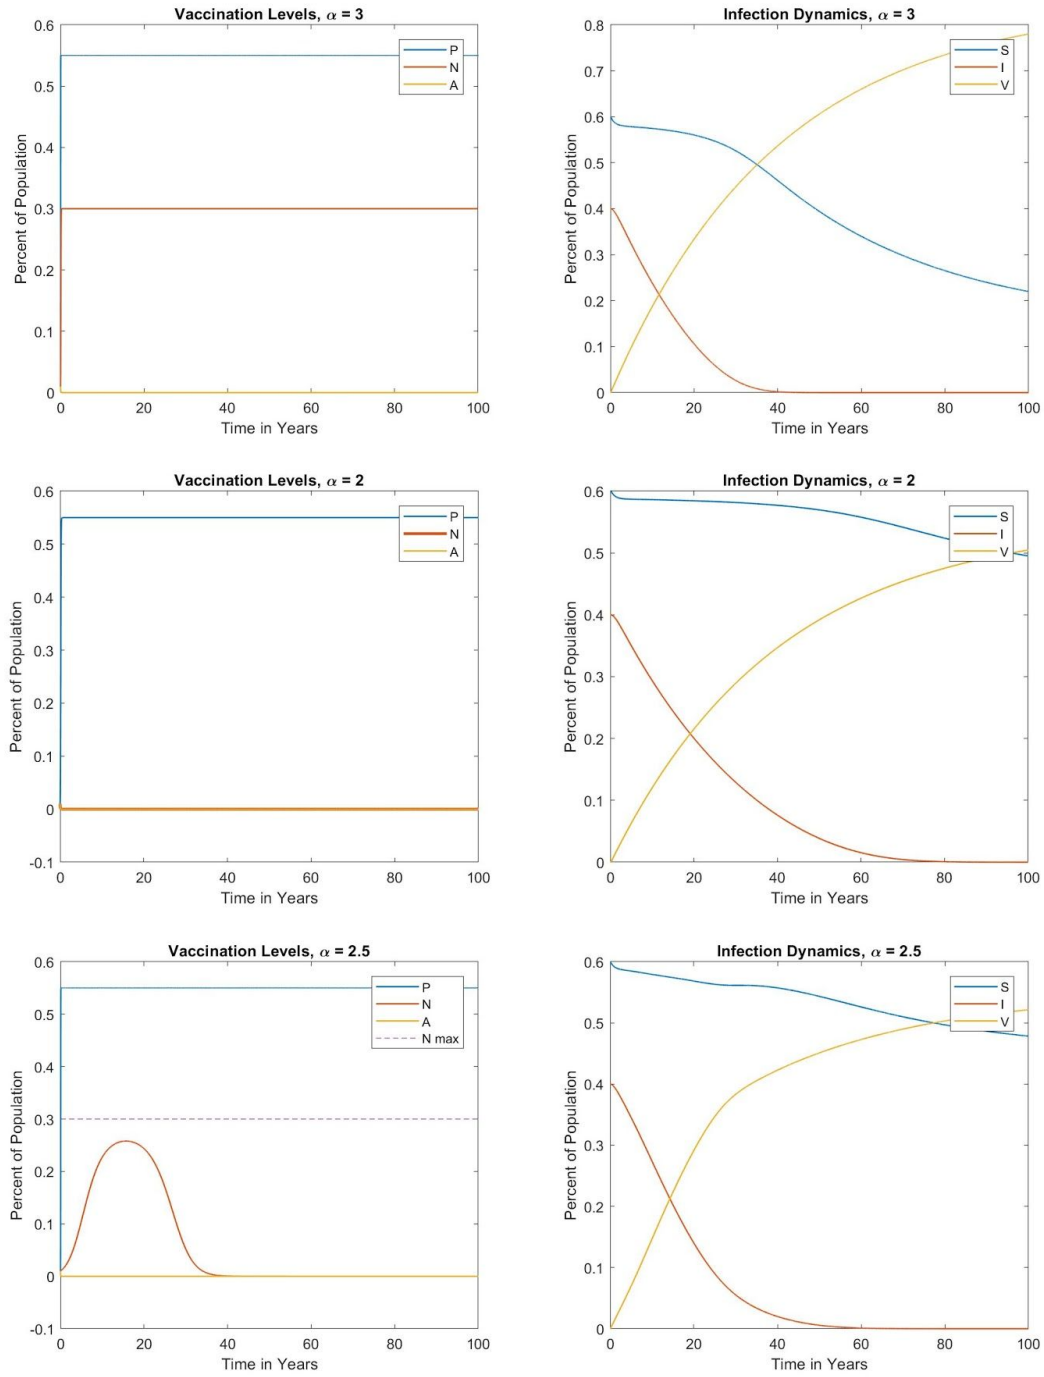

Figure 6: Top rox: vaccination levels (left) and disease progression (right) for α = 3. Middle row: for α = 2. Bottom row: for α = 2.5

We can see from these simulations that the behavior of the neutral parent group depends heavily on the value of alpha in relation to the other parameters. On the other hand, for most reasonable values of alpha, the pro group will always vaccinate and the anti group will not.

However, if we drastically increase alpha, say α = 4, then we can get the anti group to choose to vaccinate.

In addition to varying our alpha parameter, we can also vary beta to study the effects that higher or lower contact rates would have on the disease progression. Below we plot a heat map of alpha[8] vs. beta for values ranging from 0-8, with the color of each square indicating the disease prevalence after 50 years.

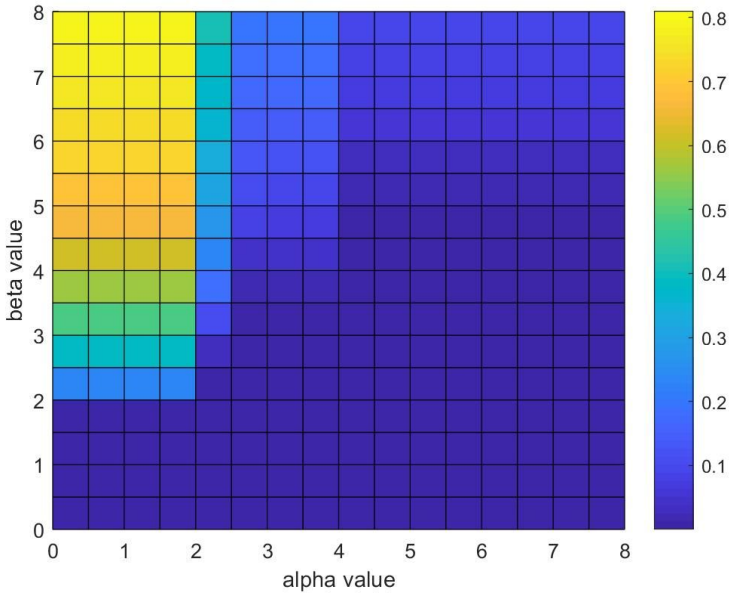

Figure 7: Heat map of α vs. β. Decimal values displayed in the legend refer to the percentage of the total population in the infected compartment after 50 years

As we would expect, higher values of beta with reasonably low (0 - 2.5) values of alpha cause the disease to persist in a state of epidemic even after 50 years. As alpha increases, though, more people (i.e. the neutral group parents) choose to vaccinate and the disease dies out, even for extremely high values of beta.

4. Conclusion

In this paper, we constructed a model of parents' HPV vaccine decision-making, dependent upon whether or not their doctor recommends the vaccine. We divided the parent population into three categories (pro, neutral, and anti) and chose appropriate risk parameters for each. As we expected, for all reasonable values of alpha and r_i, the pro group will choose to vaccinate and the anti group will choose not to, given a fixed disease prevalence of 80%. The neutral group's behavior is dependent upon the balance between these two parameters, with α = 2.479 being the "tipping point" value when $r_i = 1$.

[8] Values on heatmap refer to value of alpha before it is scaled to the 1-1,000 scale we have been using

When we combined this decision model with an SIS disease model, we found again that the neutral group's behavior depends heavily on the exact value of alpha, and that their vaccination choice determines the length of time it takes for the disease to die out. For example, with $\alpha = 3$, the neutral group will vaccinate and infection reaches zero after about 40 years. With a slightly lower value of $\alpha = 2$, however, the neutral group will not vaccinate and it takes close to 80 years for the disease to die out. The effect of alpha on the disease model gives us insight into how the strength of provider recommendation can greatly impact the course of an epidemic, specifically by influencing the neutral parents one way or another.

We concluded our analysis by looking at how varying the value of beta (the contact rate) can also affect our model. We created a heatmap of alpha vs. beta, which illustrates the interplay between the two parameters.

Our model and subsequent analyses demonstrate the nuanced effects of both alpha and beta, a product of the back-and-forth interaction between the disease and decision models. For example, a higher beta value means that the disease will initially spread more rapidly, but it will also encourage more parents to vaccinate due to the higher perceived probability of infection. Conversely, a higher alpha value will increase the payoff of vaccination (given a doctor recommends it), but greater vaccination rates will decrease the overall disease prevalence so that then the neutral group may eventually stop vaccinating (as is the case in the bottom left pane of Figure 6). In our simulations, we were unable to find a set of parameter values for which we could induce a second wave of infection, but we believe it is certainly possible that this could occur if vaccination wanes in response to a drop in perceived susceptibility. Thus it is crucial that doctors recommend the vaccine strongly and consistently, so that we can eventually reach the point of near or total eradication in the population and maintain this even if a new infective is introduced.

Appendix

Using linear interpolation to convert the 1-10 scale to a 1-4 scale, where 1 maps to 4 and 10 maps to 1, we get the equation $y = -\frac{1}{3}x + 4\frac{1}{3}$, which we use to convert the mean values from Lindley et al into our r_v values for each of the three parent groups.

References

1. CDC. (2017). Genital HPV Infection - Fact Sheet. Retrieved from https://www.cdc.gov/std/hpv/stdfact-hpv.htm
2. American Cancer Society. (2017). HPV and HPV Testing. Retrieved from https://www.cancer.org/cancer/cancer-causes/infectious-agents/hpv/hpv-and-hpv-testing.html
3. CDC. (2017). How Many Cancers Are Linked with HPV Each Year?. Retrieved from https://www.cdc.gov/cancer/hpv/statistics/cases.htm
4. CDC. (2018). Human Papillomavirus (HPV) and Cancer. Retrieved from https://www.cdc.gov/cancer/hpv/
5. Meites, E., Kempe, A., & Markowitz, L. E. (2016, December 16). Use of a 2-Dose Schedule for Human Papillomavirus Vaccination — Updated Recommendations of the Advisory Committee on Immunization Practices. *MMWR, 65*(49), 1405-1408. Retrieved from https://www.cdc.gov/mmwr/volumes/65/wr/mm6549a5.htm
6. Merck & Co. What is HPV?. Retrieved from https://www.gardasil9.com/hpv-facts/what-is-hpv/?HPVTypes=true
7. Walker, T. Y., Elam-Evans, L. D., Singleton, J. A., Yankey, D., Markowitz, L. E., Fredua, B., Williams, C. L., Meyer, S. A., & Stokley, S. (2017, August 25). National, Regional, State, and Selected Local Area Vaccination Coverage Among Adolescents Aged 13–17 Years — United States, 2016. *MMWR, 66*(33), 874-882. Retrieved from https://www.cdc.gov/mmwr/volumes/66/wr/mm6633a2.htm#T1_down
8. Holman, D. M., Benard, V., Roland, K. B., Watson, M., Liddon, N., & Stokley, S. (2014). Barriers to human papillomavirus vaccination among US adolescents: a systematic review of the literature. *JAMA Pediatrics, 168*(1), 76-82. doi:10.1001/jamapediatrics.2013.2752
9. Bauch, C. T., & Earn, D. J. (2004). Vaccination and the theory of games. *PNAS, 101*(36), 13391-13394. doi:10.1073/pnas.0403823101
10. CDC. (2018). Frequently asked questions about HPV Vaccine Safety. Retrieved from https://www.cdc.gov/vaccinesafety/vaccines/hpv/hpv-safety-faqs.html#
11. Mayhew, A., Kowalczyk Mullins, T. L., Ding, L., Rosenthal, S. L., Zimet, G. D., Morrow, C., & Kahn, J. A. (2014). Risk perceptions and subsequent sexual behaviors after HPV vaccination in adolescents. *Pediatrics, 133*(3), 404-411. doi:10.1542/peds.2013-2822
12. Kahn, J. A., Rosenthal, S. L., Jin, Y., Huang, B., Namakydoust, A., & Zimet, G. D. (2008). Rates of human papillomavirus vaccination, attitudes about vaccination, and human papillomavirus prevalence in young women. *Obstetrics & Gynecology, 111*(5), 1103-1110. doi:10.1097/AOG.0b013e31817051fa
13. Reiter, P. L., Brewer, N. T., Gottlieb, S. L., McRee, A. L., & Smith, J. S. (2009). Parents' health beliefs and HPV vaccination of their adolescent daughters. *Social Science & Medicine, 69*, 475-480. doi:10.1016/j.socscimed.2009.05.024
14. Bauch, C. T., & Earn, D. J. (2004). Vaccination and the theory of games. *PNAS, 101*(36), 13391-13394. doi:10.1073/pnas.0403823101

15. Dempsey, A., Cohn, L., Dalton, V., & Ruffin, M. (2010). Patient and clinic factors associated with adolescent human papillomavirus vaccine utilization within a university-based health system. *Vaccine, 28*(4), 989-995. doi:10.1016/j.vaccine.2009.10.133
16. Donahue, K. L., Hendrix, K. S., Sturm, L. A., & Zimet, G. D. (2015). Human papillomavirus vaccine initiation among 9–13-year-olds in the United States. *Preventive Medicine Reports, 2*, 892-898. https://doi.org/10.1016/j.pmedr.2015.10.003
17. Reiter, P. L., Cates, J. R., McRee, A. L., Gottlieb, S. L., Shafer, A., Smith, J. S., & Brewer, N. T. (2010). Statewide HPV vaccine initiation among adolescent females in North Carolina. *Sexually Transmitted Diseases, 37*(9), 549-556. doi:10.1097/OLQ.0b013e3181d73bf8
18. Rosenthal, S. L., Rupp, R., Zimet, G. D., Meza, H. M., Loza, M. L., Short, M. B., & Succop, P. A. (2008). Uptake of HPV vaccine: demographics, sexual history and values, parenting style, and vaccine attitudes. *Journal of Adolescent Health, 43*, 239-245. doi:10.1016/j.jadohealth.2008.06.009
19. Fang, C. Y., Coups, E. J., & Heckman, C. J. (2010). Behavioral correlates of HPV vaccine acceptability in the 2007 health information national trends survey (HINTS). *Cancer Epidemiology, Biomarkers & Prevention, 19*(2), 319-326. doi: 10.1158/1055-9965.EPI-09-0918
20. Wilson, K. L., White, A., Rosen, B. L., Chiappone, A., Pulczinski, J. C., Ory, M. G., & Smith, M. L. (2016). Factors associated with college students' intention to vaccinate their daughters against HPV: protecting the next generation. *The Journal of Community Health, 41*(5), 1078-1089. doi:10.1007/s10900-016-0192-8.
21. Reiter, P. L., Brewer, N. T., Gottlieb, S. L., McRee, A. L., & Smith, J. S. (2009). Parents' health beliefs and HPV vaccination of their adolescent daughters. *Social Science & Medicine, 69*(3), 475-480. doi:10.1016/j.socscimed.2009.05.024
22. Lindley, M. C., Jeyarajah, J., Yankey, D., Robinette Curtis, C., Markowitz, L. E., & Stokley, S. (2016). Comparing human papillomavirus vaccine knowledge and intentions among parents of boys and girls. *Human Vaccines & Immunotherapeutics, 12*(6), 1519-1527. http://dx.doi.org/10.1080/21645515.2016.1157673
23. Kahn, J. A., Rosenthal, S. L., Jin, Y., Huang, B., Namakydoust, A., & Zimet, G. D. (2008). Rates of human papillomavirus vaccination, attitudes about vaccination, and human papillomavirus prevalence in young women. *Obstetrics & Gynecology, 111*(5), 1103-1110. doi:10.1097/AOG.0b013e31817051fa
24. CDC. (2017). How Many Cancers Are Linked with HPV Each Year?. Retrieved from https://www.cdc.gov/cancer/hpv/statistics/cases.htm
25. Gilkey, M. B., Malo, T. L., Shah, P. D., Hall, M. E., & Brewer, N. T. (2015). Quality of physician communication about human papillomavirus vaccine: findings from a national survey. *Cancer Epidemiology, Biomarkers & Prevention, 24*(11), 1673-1679. doi: 10.1158/1055-9965.EPI-15-0326
26. Smith, P. J., Stokley, S., Bednarczyk, R. A., Orenstein, W. A., & Omer, S. B. (2016). HPV vaccination coverage of teen girls: the influence of health care providers. *Vaccine, 34*(13), 1604-1610. http://dx.doi.org/10.1016/j.vaccine.2016.01.061
27. Peyton, C. L., Gravitt, P. E., Hunt, W. C., Hundley, R. S., Zhao, M., Apple, R. J., & Wheeler, C. M. (2001). Determinants of genital human papillomavirus detection in a US population. *The Journal of Infectious Diseases, 183*(11), 1554-1564. doi:10.1086/320696
28. Koshiol, J. E., Schroeder, J. C., Jamieson, D. J., Marshall, S. W., Duerr, A., Heilig, C.M., Shah, K. V., Klein, R. S., Cu-Uvin, S., Schuman, P., Celentano, D., & Smith, J. S. (2006). Time to clearance of human papillomavirus infection by type and human immunodeficiency virus serostatus. *International Journal of Cancer, 119*(7), 1623-1629. doi:10.1002/ijc.22015