

University of New Hampshire

University of New Hampshire Scholars' Repository

Doctoral Dissertations

Student Scholarship

Summer 2022

Essays on the Economics of Vaccination

Andrew Evans Whitaker
University of New Hampshire

Follow this and additional works at: <https://scholars.unh.edu/dissertation>

Recommended Citation

Whitaker, Andrew Evans, "Essays on the Economics of Vaccination" (2022). *Doctoral Dissertations*. 2701.
<https://scholars.unh.edu/dissertation/2701>

This Dissertation is brought to you for free and open access by the Student Scholarship at University of New Hampshire Scholars' Repository. It has been accepted for inclusion in Doctoral Dissertations by an authorized administrator of University of New Hampshire Scholars' Repository. For more information, please contact Scholarly.Communication@unh.edu.

Essays on the Economics of Vaccination

By

Andrew Whitaker

DISSERTATION

Submitted to the University of New Hampshire

in Partial Fulfillment of

the Requirements for the Degree of

Doctor of Philosophy

in

Economics

September, 2022

ALL RIGHTS RESERVED

©2022

Andrew Whitaker

This dissertation was examined and approved in partial fulfillment of the requirements
for a Doctor of Philosophy in Economics by:

Co-Chair, Stephanie Brockmann, Assistant Professor of Economics

Co-Chair, Shadi Atallah, Associate Professor of Agricultural and
Consumer Economics, University of Illinois

Bradley Herring, Professor of Economics

Ju-Chin Huang, Professor of Economics

Robert Mohr, Associate Professor of Economics

Approval signatures are on file with the University of New Hampshire Graduate School

Acknowledgments

As with any major undertaking, this dissertation is the culmination of work from many quarters, and would have been impossible without tremendous help and support. In particular, I'd like to express my deepest gratitude Professor Shadi Atallah, whose teachings set me on this path of research and gave me the tools and encouragement to follow it. I am also deeply indebted to Professor Bradley Herring for the extensive work he did with me to refine my research question, models, and presentation.

I would like to extend my sincere thanks to Stephanie Brockmann who helped me navigate the entire process, and provided invaluable guidance on how to avoid burnout and keep soldiering on. I am also extremely grateful to my defense committee, who provided outstanding feedback and showed incredible flexibility.

Many thanks go to my parents Dr. Evans Whitaker and Dr. Deborah Bronstein, whose support let me focus on my studies and whose guidance helped me focus on the goal. Finally, a special thanks goes out to my wife, Lili, who has supported me through the entire process and gave me the strength to see it through.

Table of Contents

Acknowledgments.....	iv
Table of Figures.....	viii
Index of Tables.....	ix
Abstract.....	x
1. Introduction.....	1
2. Literature Review.....	7
2.1. Vaccination Decisions.....	8
2.1.1. Dynamic Demand.....	9
2.1.2. Heuristics.....	9
2.1.3. Misinformation.....	10
2.1.4. Social Networks.....	11
2.2. Disease Spread.....	11
3. Theoretical Model.....	12
3.1. Benchmark of Rational Behavior with Full Information.....	13
3.1.1. Modeling Vaccination Decisions.....	14
3.1.2. Modeling Disease Spread.....	19
3.1.3. Results for Rational Behavior with Full Information.....	21
3.2. Incorporating Heuristic Behavior.....	23
3.2.1 Modeling Heuristics.....	23
3.2.2 Results Incorporating Heuristics.....	24
3.3. Incorporating Autism Misinformation.....	26
3.3.1 Modeling Misinformation.....	26
3.3.2 Results with Misinformation.....	27
3.3.3 Results Combining Both Heuristics and Misinformation.....	29
3.4. Incorporating Social Networks.....	31
3.4.1 Modeling Social Networks.....	31
3.4.2 Results with Social Networks.....	35
3.5. Interventions.....	36
3.5.1. Subsidies.....	37
3.5.2. Mass Vaccination.....	38
3.5.3. Ring Vaccination.....	38
3.5.4. Education.....	39
3.5.5. Vaccine Ambassadors.....	41
3.5.6. Intervention Combinations.....	42
5. Intervention Results.....	43
5.1 Mandated Vaccines.....	44
5.2 Subsidized Vaccines.....	45
5.3 Education.....	48
5.4 Vaccine Ambassadors.....	49

5.5 Combined Interventions.....	50
5.6 Cost-Effectiveness of Interventions.....	53
6. Select Sensitivity Analyses.....	54
6.1. ASD Expectation.....	55
6.1.1. Mean ASD Expectation (μ_{ASD}).....	55
6.1.2. ASD Expectation Variance (σ_{ASD}).....	56
6.2. Social Media Campaign Effectiveness (α).....	57
7. Conclusion.....	58
References.....	64
APPENDICES.....	72
Appendix A: Modeling Specification.....	73
A.2 Theoretical Model.....	74
A.2.1 Rational Behavior.....	75
A.2.2 Heuristics and Misinformation.....	77
A.2.3 Disease Spread.....	77
A.3 Vaccination Decision.....	78
A.3.1 Rational Vaccination Decision.....	78
A.3.2 Heuristic Vaccination Decision.....	79
A.3.3 Misinformed Vaccination Decision.....	79
A.4 Disease Spread.....	80
A.5 Social Networks.....	81
A.5.1 Social Network Specification.....	81
A.5.2 Disease Incidence in Social Networks.....	82
A.5.3 Misinformation in Social Networks.....	82
A.6. Spatial Distribution.....	83
A.7. Stochasticity.....	83
Appendix B: Parameterization.....	84
B.1 Measles.....	84
B.2 Vaccination.....	87
B.3 Vaccination Adverse Reactions.....	90
B.4 Vaccination Decision.....	90
Appendix C: Calibration and Verification.....	92
C.1 Calibration.....	92
C.1.1 Infection Spread Distance.....	92
C.1.2 Background Risk.....	93
C.1.3 Agent Error.....	94
C.1.4. Vaccine Ambassador Impact.....	94
C.2 Verification.....	95
Appendix D: Sensitivity Analyses.....	97
D.1 Rational Agent Parameters.....	97
D.1.1 Incidence Multiplier (I_{MULT}).....	97
D.1.2 Agent Decision Error (ϵ_i).....	98
D.1.3 Social Network Incidence Weighting (w).....	99
D.2 Misinformation Parameters.....	101

D.2.1 ASD Expectation Mean (μ_{ASD}).....	101
D.2.2 ASD Expectation Standard Deviation (σ_{ASD}).....	102
D.3 Education Backlash.....	104
Appendix E: Supplementary Tables.....	106

Table of Figures

Figure 1: The Average Cost of Measles Case by Vaccination Status.....	16
Figure 2: Expected Costs of Measles by Vaccination Status and Risk.....	17
Figure 3: The Susceptible-Exposed-Infected-Recovered (SEIR) Model.....	20
Figure 4: Vaccination Path over Time for Rational Behavior with Full Information.....	22
Figure 5: Vaccination Path over time for Some Heuristic Users and Full Information....	24
Figure 6: Expected Cost of Measles Given Vaccination Status, Risk, and ASD Belief....	27
Figure 7: Vaccination Path over time for Fully Rational Users with Some Misinformation	28
Figure 8: Social Cost per Capita for Heuristic and Misinformation Permutations.....	30
Figure 9: Visualization of a Social Network.....	33
Figure 10: Social Cost per Capita with Social Networks.....	35
Figure 11: Results for Mandated Vaccines.....	44
Figure 12: Results for Subsidies.....	46
Figure 13: Results of Education.....	48
Figure 14: Results for Vaccine Ambassadors.....	49
Figure 15: Detail of Vaccine Ambassador Interventions.....	50
Figure 16: Combining Vaccine Ambassadors with Ring Vaccination.....	52
Figure 17: Combining Education and Rebates.....	52
Figure 18: Impact of ASD Mean on Total Societal Cost.....	56
Figure 19: Impact of ASD Expectation Distribution on Total Societal Cost.....	57
Figure 20: Impact of <i>Vaccine Ambassador</i> Effectiveness on Total Societal Cost.....	58
Figure 21: Transition Matrix.....	78
Figure 22: Model Verification.....	95
Figure 23: Impact of Incidence Perception Multiplier on Total Societal Cost.....	97
Figure 24: Impact of Agent Error Term on Total Societal Cost.....	98
Figure 25: Impact of Increasing Social Network Weight on <i>Total Societal Cost</i>	99
Figure 26: Impact of Mean ASD Expectation on Total Societal Cost.....	101
Figure 27: Impact of the Distribution of ASD Expectation on Total Societal Cost.....	102
Figure 28: Impact of ASD Expectation Distribution Without Social Networks.....	103
Figure 29: Impact of Education Backlash.....	104

Index of Tables

Table 1: Final Outcomes Assuming Rational Behavior with Full Information.....	23
Table 2: Final Outcomes with Some Heuristics.....	25
Table 3: Final Outcomes with Some Misinformation.....	29
Table 4: Final Outcomes With Heuristics and Misinformation.....	31
Table 5: Scenario Summary.....	42
Table 6: Cost-effectiveness of Interventions.....	54
Table 7: Measles Progression Parameters and Sources.....	85
Table 8: Measles Pathology Parameters and Sources.....	86
Table 9: Vaccine Parameters and Sources.....	89
Table 10: Cost Parameters and Sources.....	91
Table 11: Results of Single Interventions.....	105
Table 12: Results of Combined Interventions.....	106

Abstract

I examine vaccination behavior during a measles outbreak. By abandoning the rational expectations framework, I develop a model of vaccine behavior which recreates empirically observed vaccine hesitancy, as well as vaccination spikes during an outbreak. I use an agent-based model to simulate disease spread and agent behavior in a measles outbreak, in which rational agents minimize their expected costs by choosing their vaccination state. I allow some agents to instead use a heuristic, and others to have misinformation regarding vaccine risks, and finds that both reduce welfare. Including a social network has an ambiguous effect, as using more relevant local data allows agents to better estimate their risk from disease, but the same social network amplifies the impact of misinformation. I then examine a series of regulator interventions, and find that using a social media campaign to change agent's perceptions of their peers' views is the most cost-effective intervention. This presents regulators with a new framework with which to understand vaccine hesitancy, and an expanded menu of options to employ in the event of an outbreak.

1. Introduction

Rational expectations theory does not seem to apply well to the problem of vaccination. Vaccination is a very inexpensive procedure which averts very expensive outcomes, and so we would expect individuals to vaccinate in advance of a disease outbreak. Yet when we observe vaccination behavior empirically, we see that vaccination rates remain well below recommended levels until an outbreak happens. Then a spike in vaccinations is observed, but too late to prevent the outbreak in the first place. By departing from an assumption of rational expectations, I seek to reconcile this delayed vaccination rate we observe with a theoretical model of vaccination behavior in a measles outbreak, which allows for the imperfect information that rational expectation models struggle with (Colander et al., 2009, p. 7).

Measles is a deadly disease which spreads very quickly, making it a major concern for public health officials. The Measles-Mumps-Rubella (MMR) vaccine provides prophylactic protection, but its uptake is well below the near-universal adoption recommended by the U.S. Centers for Disease Control and Prevention (CDC) to obtain herd immunity. Measles spreads exceptionally fast, so even a modest drop in vaccination rates can lead to outbreaks. For example, New York City had only a 73% vaccination rate in 2019, and its number of measles cases grew from 1 to 650 despite extensive interventions by the Department of Health and Mental Hygiene (Zucker et al., 2020).

Relatively low vaccination rates are puzzling because vaccines provide large social benefits, in terms of millions of lives and billions of dollars saved over the past two decades (Ozawa et al., 2017). Vaccines are one of the few medical interventions which can reduce total medical costs by preventing disease and expensive treatments later (Armstrong, 2007). This alone should, in principle, be a compelling

argument for everyone to vaccinate, but there are additional benefits which can accrue: once a certain percent of the population is immune, that illness can no longer spread and eventually dies out. This represents a large positive externality, in which even individuals who do not vaccinate benefit from the vaccination of others.

With modest production costs for vaccines and large health benefits from vaccination, persistently low vaccination rates (relative to public health recommendations) suggest that most agents do not behave according to rational expectations theory, due at least in part to individuals having an inaccurate perception of risk (Brahmbhatt and Dutta, 2008). Furthermore, models which incorporate agents that do not have rational expectations, and may instead rely on a rule of thumb, may be good at predicting behavior in response to disease outbreaks (Lee and McKibbin, 2004).

There are factors to explain such low vaccination rates beyond the simple observation that many people do not incorporate positive externalities into their self-interested decision-making. First, when disease prevalence is low enough (perhaps due to relatively-high vaccination rates), the marginal benefits from vaccination may fall below the vaccine's price; Geoffard and Philipson (1997) refer to this endogeneity as "prevalence-dependent demand" and discuss how it prevents eradication. This is examined more recently by Manfredi et al. (2009). Second, some people may not be equipped to make rational decisions which compare complex marginal costs to marginal benefits and instead rely on heuristics (or simpler "rules of thumb") when making their vaccination decisions (MacDonald et al., 2012; Omer et al., 2017). Third, even rational behavior to compare marginal costs to marginal benefits may be hampered by misinformation about the debunked link of Autism Spectrum Disorder (ASD) with vaccines (Fombonne et al., 2020; Himelboim et al., 2020; Taylor et al., 2014). Fourth, there is growing interest in how social networks may impact these different issues (Betsch et al., 2013; Milani et al., 2020; Smith and Graham, 2017).

This dissertation incorporates each of these factors incrementally into an economic model of decision-making in order to disentangle their independent effects. I use an agent-based model (ABM) to simulate a hypothetical outbreak of measles within a heterogeneous population. There are two primary reasons to examine a hypothetical outbreak. First, it allows the likelihood of catching measles to vary over time; this, in turn, affects the benefits from vaccinating over time, so that I may incorporate how vaccination decisions are affected by changes in risk. Second, simulating an outbreak allows us to estimate important policy-relevant outcomes (i.e., numbers infected, healthcare costs, disability costs, and deaths) and, critically, how those final outcomes from otherwise-identical initial outbreaks vary across different models for economic behavior. Although many aspects of this model are informed by empirical observations, the underlying objective of my analysis is largely theoretical in nature: I seek to demonstrate how agents' vaccination decisions, and the outcomes that follow from them, are impacted by incorporating heuristic behavior, ASD misinformation, and social networks into the model.

Once I explain why individuals may be vaccine hesitant, the question becomes what the government can do about it. Typically policy instruments such as subsidies and mass education have seen only modest increases in vaccination rates. The usual instruments employed are public education measures to provide better information, and subsidies to change the price.

There are advantages to using an ABM for my analysis, given three important factors in disease spread. First, the location of the agents is important for determining infections. Second, there is an asymptomatic period in which an agent with measles is infecting others but not displaying symptoms. As a result, agents are consistently working with dated information about observable symptom status; incorporating this feature is critical. Third, vaccination decisions evolve over time endogenously. Disease spread is affected by vaccination status, and vaccination decisions are affected by disease

incidence, leading to an endogenous feedback loop that can be difficult to examine with other empirical models. But ABMs incorporate such endogeneity with relative ease.

I begin with a model of approximately rational behavior with full information. In this benchmark model, rational agents examine the expected costs of measles in a vaccinated state versus an unvaccinated state; if the vaccinated state's net costs are lower, they choose to vaccinate¹. When I introduce some heuristic users, those agents instead use a simple decision rule that they vaccinate only if the observed disease incidence rises above a certain threshold. When misinformation is introduced, I allow some agents to have a false belief that vaccines may cause ASD in a small portion of cases (and incorporate that "cost" into their expected costs of vaccinating). Finally, when I add social networks, all agents can use more relevant local data about disease incidence, but the social networks may also exacerbate the effects of misinformation. With the initial benchmark case of rational behavior with full information, incorporating these three additional dimensions (i.e., heuristic users, misinformation, social networks) generates eight different scenarios through their permutations.

The spatial and timing aspects of the ABM are important; even if individuals made decisions with a perfect process, a delay exists between when data is collected (e.g., by the CDC) and when it is published. Similarly, there is necessarily limited information on local areas to protect patients' private information. Even an ideally behaved *homo economicus* would be acting to counter an outbreak as it occurred last week, and on average. Given the exponential spread of diseases, this means the delayed information provision could lead to undervaccination. Thus, we can think of outbreaks as happening in the time between when an individual becomes infected, and when they begin to show symptoms; even with otherwise perfect information, that delay permits diseases to spread.

¹ By computationally rational, I refer to a model in which agents accurately tally costs, but may or may not accurately judge probability of infection. Details follow in subsequent sections.

By using a spatio-temporal agent-based model, I incorporate the cost, not only of the vaccine, but also of not vaccinating. That is, the agent models the spread of measles based on cases only on the cases at the present time, drawing from both global statistics and local information from their social network. Thus, because the distribution of ill individuals changes with time, and some agents may hold the false but earnest belief in long-term and delayed reactions to vaccination, a spatially implicit, inter-temporal model is necessary.

The agent is a child-parent pair, aged 18 months to 17 years. The parent observes data, receives education, communicates with other parents, and makes decisions on behalf of the child, while the child experiences the impact of health, vaccination and infection, and spreads the disease to children in other child-parent pairings. The vaccination decision is assumed to be made each day.²

Within the model, the agent chooses at each time step whether to vaccinate or not, based on which action minimizes their expected costs. This approach has two benefits: firstly, its simplicity lends it well to modeling. More importantly, it reflects a bias discussed in literature (Gershoff and Koehler, 2011) that people tend to gravitate toward simple calculations and solutions, (so-called “Enumeration Bias”) and toward (*ceteris paribus*) risks resulting from inaction instead of from action (“Omission Bias”). The cost of not vaccinating is a cost of omission, playing into the Omission Bias. These costs, while large, have a relatively small number of negative health outcomes, while the misinformed dangers of vaccinating have an ever-growing and changing list of negative outcomes, which plays into the Enumeration Bias.

Most famously these misinformed dangers include Autism Spectrum Disorders, but also includes many other factors such as death, diabetes, etc. (Dubé et al., 2015). This suggests that a bias exists toward not

² Two factors are abstracted out by this decision: first, while the majority of vaccine hesitant individuals tend to be female (Smith and Graham, 2017), this does not need to be included, as my model has the parents making the decision for the children. As such, gender demographics are assumed to balance out. Second, this model does not make allowances for siblings; while this may be an interesting extension in the future, it is not modeled here for the sake of expediency.

vaccinating because its dangers are better known. Contributing to this is that the unsubstantiated claims of vaccine dangers have been proliferated by social media (Callender, 2016). The result is that the public seems to be paying too much attention to the dangers of vaccinating, and too little to the dangers of not vaccinating.

A variety of interventions will be examined: subsidizing the cost of the vaccines to the agents; educating agents on the risks of the disease; implementing a vaccine ambassador program where members of the public are trained to educate their peers; mass vaccination where all agents are vaccinated before an outbreak; and targeted vaccination, where vaccines are administered to anyone who may have been in contact with an infected individual. A feature worthy of note is that, for education, a backlash effect is modeled; that is, for every 100 people educated, 23 tend to become more entrenched in their incorrect views (Attwell and Freeman, 2015). As such, education has the potential to make some individuals' assessment of costs *more* incorrect.

I find that agents using misinformation unambiguously reduces societal welfare, and can capture the empirically observed vaccination behavior in vaccine-hesitant communities. Heuristics as modeled also show a decrease in welfare. Incorporating social networks has an ambiguous effect: the gains from improved agent forecasting using local information are counterbalanced by losses due to exacerbated misinformation.

With respect to government intervention, I find that social nudges are the most effective. However, ring vaccination provides a second-best alternative. Ring vaccination is defined here as the process by which, when an individual begins to show symptoms, a prophylactic vaccine is mandated for anyone potentially exposed. Policy interventions are discussed in more detail in Section 3.5. Interventions (p. 36).

The remainder of the dissertation is organized as follows: Review of previous work in the areas of vaccination decisions and disease spread is provided in Section 2. Literature Review (p. 7). Section 3. Theoretical Model (p. 12) describes the model in detail and compares the performance of the model over the three characteristics of heuristics, misinformation, and social networks. It also defines the interventions the regulator is able to perform. Section 5. Intervention Results (p. 43) discusses the outcomes of the various regulator interventions. Section 6. Select Sensitivity Analyses (p. 55) shows how model outcomes change as certain parameters change. The paper concludes in Section 7. Conclusion (p. 59). Details of the exact model specification can be found in Appendix A: Modeling Specification (p. 74).

2. Literature Review

The beneficial impacts of vaccines are well-studied (Armstrong, 2007; Ozawa et al., 2017; Seither et al., 2014; White et al., 1985). In addition to benefiting society by preventing extremely costly outbreaks, they further reduce the healthcare costs *an individual* may expect to face. The risks associated with vaccinating are nearly entirely mild, and uniformly nonfatal (Govaert et al., 1993). Despite the known benefits, vaccines have regularly been underutilized by the general public (Callender, 2016; Fine and Clarkson, 1986).

The general theme of the literature to date is extensive, but specialized research. As shall be elaborated on, there are many important factors which determine vaccination behavior, but the literature has focused on the individual factors; few papers attempt to combine them. As far as I am aware, none have performed an economic analysis which incorporates the observed imperfect information that

individuals use. As a result the literature is full of incremental progress, but fails to comprehensively explain the phenomenon of vaccine hesitancy and examine the impact to society.

2.1. Vaccination Decisions

As noted above, the benefits of vaccines have been extensively studied (Armstrong, 2007; Ozawa et al., 2017; Seither et al., 2014; White et al., 1985). These benefits come in two forms, the individual and the collective. Individually, being vaccinated against measles lowers the probability that one will fall ill, and even if one does become sick, it reduces the severity of the disease, which means vaccines can reduce the overall medical expenditures of an individual. Collectively, they can prevent large outbreaks altogether by inducing herd immunity, the state where enough of a population is immune that the disease cannot propagate.

Despite these benefits, vaccines have been regularly underutilized by the public (Callender, 2016; Fine and Clarkson, 1986). Given that the price of the MMR vaccine is relatively low (see, for example, “Current CDC Vaccine Price List” 2021), and that vaccine-adverse reactions are nearly always mild and uniformly nonfatal (Govaert et al., 1993), it is not clear why vaccine uptake is not higher.

Economic factors such as vaccine price, demographic characteristics, implicit costs, and time (e.g., combining several inoculations into a single injection) are contributing factors but generally have minimal effects (Chen et al., 2011; Davis, 2010; Jacobson and Sewell, 2002; Kondo et al., 2009; Kulpeng et al., 2013; Manfredi et al., 2009).³ Giving vaccines away for free does not result in near-universal uptake (Fisher, 2012); this suggests that there are more complexities to consider in understanding vaccination decisions.

³ The psychology literature has examined cognitive biases (Gigerenzer and Edwards, 2003; Hotez et al., 2020; MacDonald et al., 2012). The sociology literature has examined the impacts of peer pressure (Day, 2021; Schoeppe et al., 2017; Sobkowicz and Sobkowicz, 2021).

2.1.1. Dynamic Demand

One complexity is that the demand for vaccinations is unlikely to be static over time. Instead, the demand for a vaccine is likely to be dependent on how changes in the risk of catching the disease impact the marginal benefits of vaccinating. This is the notion of “prevalence-dependent demand” introduced by Geoffard and Philipson (1997). Because the demand for vaccinations decreases as the disease prevalence decreases, disease eradication is elusive; so when vaccination rates decline, the prevalence can resurface, leading to cycles over time. Dror et al. (2020) and Amirthalingam et al. (2013) indeed find that vaccination decisions are dependent on variation in prevalence, but these studies do not allow for the disease spread to also be impacted by vaccination decisions.

2.1.2. Heuristics

A second complexity is that these individual cost-benefit calculations are themselves quite complicated for many people, so the assumption of rational behavior may not hold for everyone. Healthcare costs are often obscure, and the probabilities of disease, disability, and death are often difficult to grasp. Many people, including trained professionals, have trouble estimating exponential functions (Gigerenzer and Edwards, 2003). Given that disease spreads exponentially, it is reasonable to expect that at least some agents would be unable or unwilling to devote the resources to accurately gauge risk from a disease outbreak and may instead use a simple heuristic (or “rule of thumb”). In particular, people may choose to vaccinate only if they believe that a disease is a threat (Jacobson et al., 2015).⁴

⁴ Other heuristics have been proposed, including peer opinion (Metzger et al., 2010), institutional trust (Razai et al., 2021), and a plethora of cognitive biases (Omer, Amin, and Limaye 2017; Salmon et al. 2015).

2.1.3. Misinformation

A third complexity is the potential for misinformation. In a 1998 study published in the *Lancet*, Andrew Wakefield and colleagues presented results suggesting that there was a connection between the MMR vaccine and pervasive developmental disorder in children (Wakefield et al., 1998). Despite the retraction of this Wakefield et al. study by the *Lancet* editors (Horton, 2010), the narrative of vaccine-induced ASD remains; 16% of caregivers believed in some link between vaccines and ASD (Fombonne et al., 2020). This misinformation has since spread to incorporate a large range of ailments, including minor ones like rash (Callender, 2016), Sudden Infant Death Syndrome (Yang and Shaw, 2018), and long-term complications such as ADHD (ibid.), and overloaded immune system (Berman, 2020). These fears of nonexistent outcomes affect vaccination decisions (Imhoff and Lamberty, 2020; Romer and Jamieson, 2020; Roozenbeek et al., 2020) with very real consequences (Salmon et al., 2015). Such a drop in vaccination has been associated with new outbreaks of vaccine-preventable diseases (Siddiqui et al., 2013).

Solutions to this wrinkle have been explored. In small groups, it's been shown that an education initiative helps raise the intent to vaccinate (Valdez et al., 2015). However, when actually attempted (Attwell and Freeman, 2015), the results were mixed; while 77% of the participants did have a positive response, the remainder was more polarized in their rejection of vaccines.

I incorporate this education backlash effect, betrayal aversion, and economic considerations within a bounded rationality framework, all while modeling how the spread of the disease is impacted by and impacts people's vaccination decisions.

2.1.4. Social Networks

A fourth complexity is how social networks affect each of these considerations. On the one hand, social networks can provide more-localized information about disease prevalence in a way that is not feasible in the absence of a social network. In this context, a parent may be better able to make a vaccination decision based on more-relevant local data for the prevalence of measles (as opposed to information only about the global prevalence of measles); this could benefit both those making rational decisions and those using heuristics. On the other hand, though, social networks have the ability to alter an individual's misinformation about the risks of vaccination-induced ASD; this has been shown recently by Hotez et al. (2020),⁵ though Lord et al. (1979) presented early evidence of confirmation bias being reinforced by our social networks. Such networks can moderate misinformation; a single vaccine-hesitant individual in a sea of vaccinating agents is more likely to vaccinate (Attwell and Freeman, 2015). But multiple vaccine-hesitant individuals who become better able to interact with each other can form ideological echo chambers, as social networks tend to asymmetrically reinforce opinions people already hold (Wojcieszak et al., 2021). This tendency has been emphasized in the context of vaccination decisions (Sobkowicz and Sobkowicz, 2021; Wu and Zhang, 2013), though has not yet been examined in an economic context.

2.2. Disease Spread

A second relevant area is the epidemiological literature on modeling the spread of disease, namely the "Susceptible-Exposed-Infectious-Recovered" (SEIR) model. Such models have been extensively used to model outbreaks (Biswas et al., 2014; Li et al., 1999), including models which include a vaccination

⁵ The connection between misinformation and social networks is not limited to ASD. Other related issues include "immune overload" (i.e., a concern about the simultaneous administration of three attenuated viruses) (Hulsey and Bland, 2015) and misinformation understating the effectiveness of vaccines (Valdez et al., 2015). However, for the purposes of this paper, I will examine only the false link between vaccines and ASD.

component (d’Onofrio, 2002; De la Sen et al., 2012). However, these epidemiological SEIR models have only incorporated vaccine use exogenously, either through an assumption of an incomplete take-up rate or perhaps a mandate on the population, which can be ineffective when put into practice (Richwine et al., 2019). Jaharuddin and Bakhtiar (2020) present a model which does allow vaccine uptake to vary, but is still a form of regulator intervention, not spontaneous agent behavior. A solid theoretical understanding of vaccination behaviors in an outbreak context should therefore allow for heterogeneity in individual behavior, especially in response to changing conditions.

3. Theoretical Model

The theoretical model employed is an agent-based model, simulating 500 agents over 150 time periods. The model simulates both the spread of measles in an outbreak, as well as agent’s response to the outbreak via vaccination decisions, and tracks outcomes such as social cost per capita and vaccination rate. This is discussed in more detail in subsection 3.1. Benchmark of Rational Behavior with Full Information (p. 13), which discusses how we expect a fully informed agent, computing costs properly using a reasonable approximation of disease spread, to act. Subsection 3.2. Incorporating Heuristic Behavior (p. 23) examines the impact of allowing some agents to use a “rule of thumb” instead of performing a cost-benefit analysis, while subsection 3.3. Incorporating Autism Misinformation (p. 26) details how a misinformed agent would act. Subsection 3.4. Incorporating Social Networks (p. 31) details how social networks are implemented and their impact on our model. Subsection 3.5. Interventions (p. 36) details the various interventions regulators may perform, and how they are implemented. A complete specification for the model can be found in Appendix A: Modeling Specification (p. 74).

3.1. Benchmark of Rational Behavior with Full Information

I begin by first examining a model which assumes rational behavior with full information. I use this as a benchmark to generate baseline outcomes, against which I compare the results of other model specifications (i.e., incorporating heuristics, misinformation, and social networks). This section also serves to introduce the general modeling framework (which will carry over into those additional models). I discuss the model in broad terms here in the text; for a full description of its more technical mechanisms, see Appendix A: Modeling Specification (p. 74).

The model consists of 500 agents randomly distributed in a geometric space⁶. Each agent represents a parent-child pair, in which the parent gathers information and makes decisions, while the child is subjected to disease spread. I assume the parent has fully internalized the utility of the child, and each parent-child pair is treated as a single agent.

The simulation has 150 time periods⁷, with each period representing one day. At the start of the simulation, one agent is infected with measles, which then spreads spatially throughout the population over time (The SEIR modeling for disease spread is described in Section 3.2. further below). The agents will, in every period, decide whether to get vaccinated based on the expected costs of vaccinating versus the expected costs of remaining unvaccinated (The economic modeling for this vaccination decision is described in Section 3.1.1. below). In my model, they make this decision in the absence of any vaccine mandates and consider only their individual benefit from vaccination (i.e., they do not consider the positive externalities from vaccinating).

Importantly, the expected costs of measles (whether vaccinated or unvaccinated) are dependent on the probability of being infected by measles. When individuals gauge their risk of infection, they do so by

⁶ For details, see Appendix A.6. Spatial Distribution (p. 84).

⁷ This number is chosen for completeness; in practice, the simulation reaches equilibrium after around 60 time periods.

observing the percent of the population that is currently showing symptoms and assuming that the probability of infection is directly proportional to that incidence of symptoms⁸. This is not a perfectly accurate forecast of future events; because such future events are unknown, it represents an approximation used by the agents for the purposes of their decision making⁹. One could imagine a constant of proportionality such that an agent's forecast approximates the true outcome. A detailed analysis of how this constant of proportionality impacts model performance can be found in Appendix D.1.1 Incidence Multiplier (IMULT) (p. 97).

As noted above, the percent symptomatic is different than the percent of the population truly infected because being infected is unobservable during a three-to-five day asymptomatic period. Nevertheless, I refer to this as the “full information” scenario because even accomplished epidemiologists have trouble accurately modeling the true spread of an outbreak in real time (Daunizeau et al., 2020).

After 150 days, the simulation ends, and data is gathered for four key outcomes: vaccination rates, attack rates (i.e., what percent of the population became ill), number of deaths, and total social costs (described in more detail below). To generate confidence intervals, 1,000 simulations for each of the eight scenarios are run.¹⁰

3.1.1. Modeling Vaccination Decisions

In general, economists assume agents seek to maximize utility. In this model, I assume agents minimize expected costs and damages from the disease and its vaccine. This assumption is made for two reasons.

First, it gives great gains in simplicity. Costs can be directly observed, and thus there is ample data to

8 The perceived probability of infection is capped at 100%.

9 Because agents make the vaccination decision each time period, the farthest they need to forecast is two weeks. If they forecast two weeks in advance, they can vaccinate in time for the vaccine to become effective. Forecasting 15 days in advance gives no added benefit; they will face the same outcome as if they wait a day, forecast 14 days ahead, and end up in the same immune state on day 15 as if they had a longer forecast.

10 In the context of this model, the 95% confidence interval comes only from variance in how the disease spreads through the population. All characteristics are the same through simulation runs, with the only variation coming from the path the disease takes. See Appendix A.7. Stochasticity (p. 84) for a detailed discussion.

inform my choices for parameters regarding costs. Furthermore, by minimizing expected costs, one only has to compute the probabilities and magnitudes of each cost and add them together. An expected utility framework would require computing every combination of costs, finding the corresponding utility level, adjusting for the probability of those outcomes, and summing them. Given the number of outcomes involved in a measles case, this would quickly become prohibitively computationally intensive.

Second, as alluded to above, using an expected utility model would require defining the utility function, including the level of risk aversion. Without defining such a utility function, there would be no way for the model to compare different outcomes. If such a function were defined, its choice would be completely arbitrary, and I expect the model would be highly sensitive to the specification of that utility function.¹¹ By assuming risk-neutral agents, maximizing expected utility becomes equivalent to maximizing expected wealth, or minimizing expected costs.

For the purpose of this model, agents decide on their vaccination behavior by comparing their expected costs if they are vaccinated $E[C_{VAC}]$ against the expected costs if they are not vaccinated $E[C_{UNV}]$ and choosing the lower expected cost.¹² Moreover, I incorporate additional heterogeneity in utility costs across agents by adding a normally distributed error term with mean zero to the following decision rule: *vaccinate* in period t if $E[C_{UNV,t}] - E[C_{VAC,t}] > \varepsilon$.

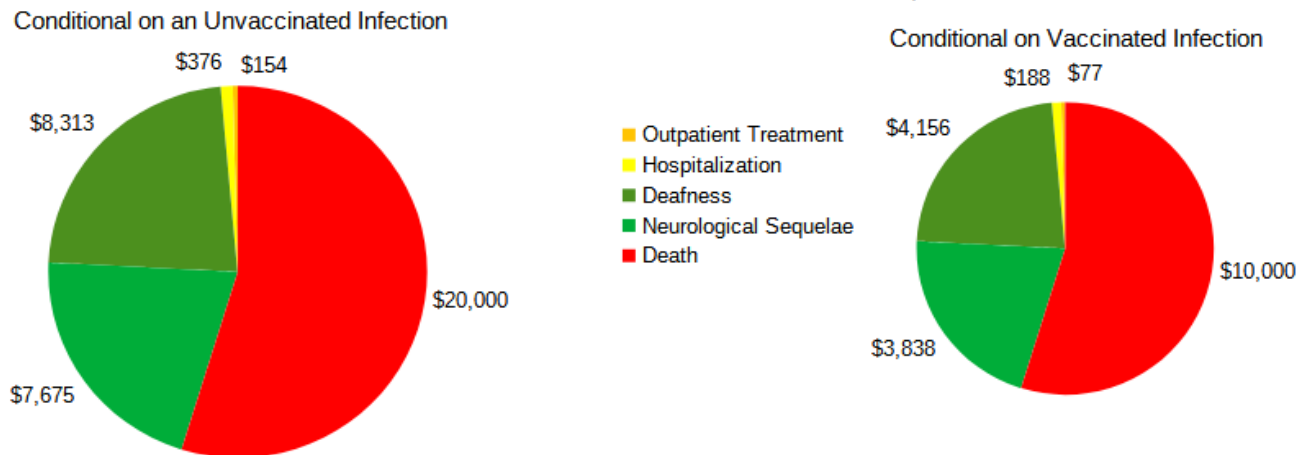
The costs conditional on becoming sick with measles are shown in Figure 1: The Average Cost of Measles Case by Vaccination Status (p. 16). They are grouped into three broad categories: the treatment costs, disability costs, and mortality. Treatment costs, defined as T , are the short-term outpatient and hospitalization costs for a case of measles. The disability costs, defined as D , are the long-term

11 While this would be an interesting extension to the model, such an expected utility framework, including sensitivity to the functional form of the utility function, is not explored here.

12 These costs are dollar valuations, which include dollar values for willingness to pay to avoid death.

disutility from potential complications caused by measles, namely deafness and neurological sequelae. Finally, the mortality cost, defined as M , arises because a small percentage of measles cases lead to death (with a cost equal to that probability of death times the Value of a Statistical Life). Note that the left side of this figure shows the expected costs of a measles case for someone who is unvaccinated, while the right side shows the expected costs of a measles case for someone who is vaccinated. While vaccines lower the probability of infection (albeit not down to a zero probability), an additional benefit is that vaccines also reduce the severity of infection, thereby lowering the average conditional costs of treatment, disability, and mortality. I define this reduction in severity as $0 < b < 1$, such that $T_{VAC} = bT_{UNV}$, $D_{VAC} = bD_{UNV}$, and $M_{VAC} = bM_{UNV}$. For a detailed discussion on where these values come from, see Appendix B: Parameterization (p. 85).

Figure 1: The Average Cost of Measles Case by Vaccination Status



The expected costs of measles are produced by taking the above cost of a case of measles, and multiplying them by the risk of being infected, $r(s_t)$, where this probability is a function of the percent of the population in period t who are showing symptoms, s_t ; individuals who are infected, but latent, cannot be counted in one's measure of risk. As noted above, the primary benefit from vaccinations is that they reduce the probability of infection, holding the percent symptomatic constant:

$p_{VAC}(s_t) = kp_{UNV}(s_t)$ where $0 < k < 1$. We can therefore express the expected costs of measles if unvaccinated and the expected costs of measles if vaccinated as:

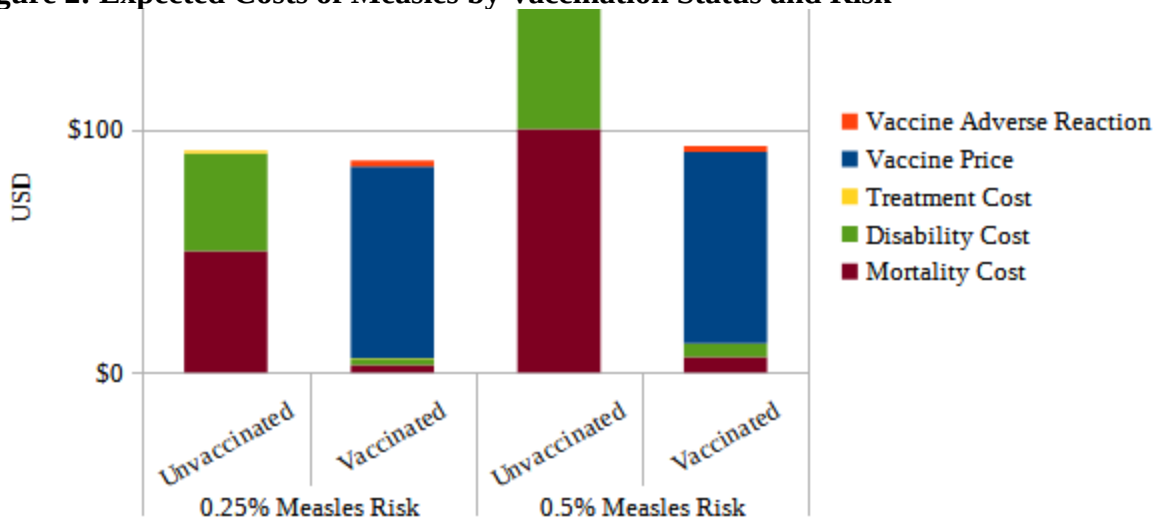
$$E[C_{UNV,t}] = p_{UNV}(s_t) \times (T_{UNV} + D_{UNV} + M_{UNV})$$

$$E[C_{VAC,t}] = p_{VAC}(s_t) \times (T_{VAC} + D_{VAC} + M_{VAC}) + P_V + A_V$$

$$= kb \times p_{UNV}(s_t) \times (T_{UNV} + D_{UNV} + M_{UNV}) + P_V + A_V$$

where P_V is the vaccine's price and A_V is the expected disutility costs of vaccine adverse reactions (VARs). The adverse reaction I am modeling is aseptic meningitis¹³. These adverse reactions are rare and uniformly nonfatal. (To be clear, these are not the debunked links to autism.)

Figure 2: Expected Costs of Measles by Vaccination Status and Risk



Estimates for these two expected costs are shown in Figure 2: Expected Costs of Measles by Vaccination Status and Risk (p. 17), where the left shows them for a relatively low probability of being infected and the right shows them for a relatively high probability of being infected. At this relatively

13 This VAR is chosen because it is the only one with substantial attached costs. More common VARs include soreness, fatigue, and even anaphylaxis or seizures. However, all of those either require no medical intervention (soreness, fatigue) or else temporary intervention (e.g., anaphylaxis, seizures) with relatively low costs. Aseptic meningitis requires several days of hospitalization, with substantial cost. Compared to this, the other costs are considered negligible.

low level of risk (i.e., 0.25% risk), an agent would be roughly indifferent between remaining unvaccinated and becoming vaccinated. As risk rises, both unvaccinated costs and vaccinated costs rise, but the unvaccinated cost rises much more quickly due to the multiplicative effect of kb . As a result, at the relatively high level of risk (i.e., 0.5% risk), an agent's expected costs of remaining unvaccinated are roughly twice as large as the costs of becoming vaccinated, and thus the agent would decide to vaccinate.

More formally, this vaccination decision rule of $E[C_{UNV,t}] - E[C_{VAC,t}] > \varepsilon$ can then be expressed as:

$$\text{vaccinate if } p_{UNV}(s_t) \times (T_{UNV} + D_{UNV} + M_{UNV}) \times (1-kb) - P_V - A_V > \varepsilon$$

Each of these factors influencing the vaccination decision is relatively intuitive. Regarding the underlying disease's characteristics, an agent is more likely to vaccinate when the risk of catching measles, $p_{UNV}(s_t)$, is higher or the cost of a case of measles, $T_{UNV} + D_{UNV} + M_{UNV}$, is higher. Regarding the vaccine's characteristics, an agent is more likely to vaccinate when its price, P_V , is lower, when its adverse reactions, A_V , are smaller, or when its benefits, $1-kb$, are greater (recall that $b \rightarrow 0$ for reduced severity and $k \rightarrow 0$ for reduced infections each imply that the vaccine's effectiveness is better). As noted above, I assume that agents do not consider the positive externalities of vaccination in their own decisions to vaccinate, and there are no public mandates to vaccinate. Moreover, I assume that agents do not have private or public insurance to cover the costs of vaccines or treatment.

In the model, agents decide whether to vaccinate in each period t . The only parameter in the above decision rule which varies over time is the risk of infection; all the other parameters are static. Recall that the risk of infection is based on the proportion of the population that shows symptoms of measles. That said, agents are assumed to know that there is a nonzero baseline risk of infection even if none of the 500 agents in the model are actually showing symptoms. I therefore define $p_0 \equiv p_{UNV}(s=0)$ and have

$p_{UNV}(s_t) \geq p_0$ for all $s_t > 0$. Put differently, agents update their perception of risk in each period by selecting the maximum of baseline risk or the level of risk based on observable symptoms in the population. This allows some agents to start off vaccinated in the model despite no one actually showing symptoms (yet); I choose the value for this baseline risk so that the vaccination rate in the first period is 50%.

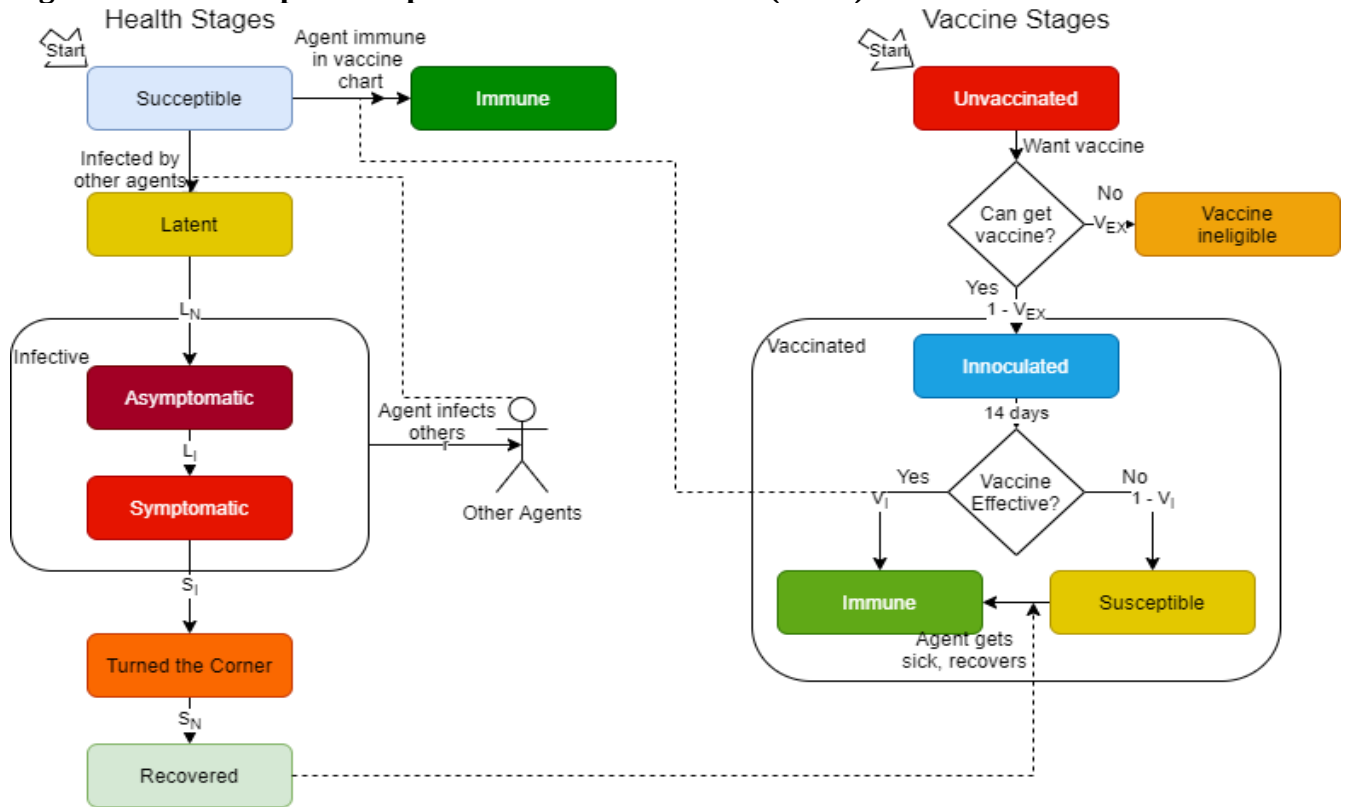
3.1.2. Modeling Disease Spread

The disease begins spreading by infecting a random agent in the first period of the simulation. It spreads through the population following a modified Susceptible-Exposed-Infected-Recovered (SEIR) model. All agents start out in a susceptible state, where they are not infected but vulnerable. If an infected neighbor has contact with the agent, they become exposed. In this state, they are incubating the disease, but not showing symptoms or infecting others.

After an amount of time, which is heterogeneous across agents, the infected agent begins infecting others, but is still not showing symptoms. After another heterogeneous amount of time, the agent begins showing symptoms and continues to be infectious. After three to five days, an agent continues to show symptoms, but ceases to infect others, a state commonly called “turning the corner.” Finally, an agent will recover, and is immune to further infection.¹⁴ This SEIR process is shown in Figure 3: The Susceptible-Exposed-Infected-Recovered (SEIR) Model (p. 20) , and discussed in more detail in Appendix A: Modeling Specification (p. 74).

¹⁴ In the time horizon of this model, waning immunity from contracting the disease or from vaccines are not considered significant. Thus, in this model, immunity is considered permanent.

Figure 3: The Susceptible-Exposed-Infected-Recovered (SEIR) Model



Vaccinating puts an agent into an inoculated state. Once in this state, the severity of sickness is reduced should they become ill. After two weeks in this inoculated state, an agent will become immune with a probability of $1-k$. An immune agent will never become infected. Agents that are not immune remain inoculated. If an agent becomes ill before that two-week inoculation period elapses, or because the vaccine was not fully effective, they face the reduced costs of being ill. After recovering, they become immune.

The disease spreads through the population in my ABM in the following manner. An infective agent will attempt to infect one random agent who is located nearby.¹⁵ If that agent is susceptible, they become exposed, and enter the latent state. If that agent is immune or already infected, there is no

¹⁵ For a discussion of where these rates, lengths, and other parameters come from, see Appendix B: Parameterization (p. 85).

effect. This behavior simulates the fact that a sick individual in a healthy population will spread the disease more quickly than in an already ill population. It also captures the impact of herd immunity.

There are two additional salient features to the model. First, individuals in an asymptomatic period are contributing to the disease spread but are not counted in the observable risk measure, while individuals who have “turned the corner” are counted in the observable risk measure but are not actually spreading the disease. This means that when agents make their vaccination decisions based on data for the observed incidence of symptoms, they are using dated information about the true rate of infection (which is unobservable). In essence, an outbreak spreads during the time between when an agent starts to infect others and starts to show symptoms, and this gap can lead to suboptimal decisions *ex post*, even when agents are rational and use all of the observable information available to them.

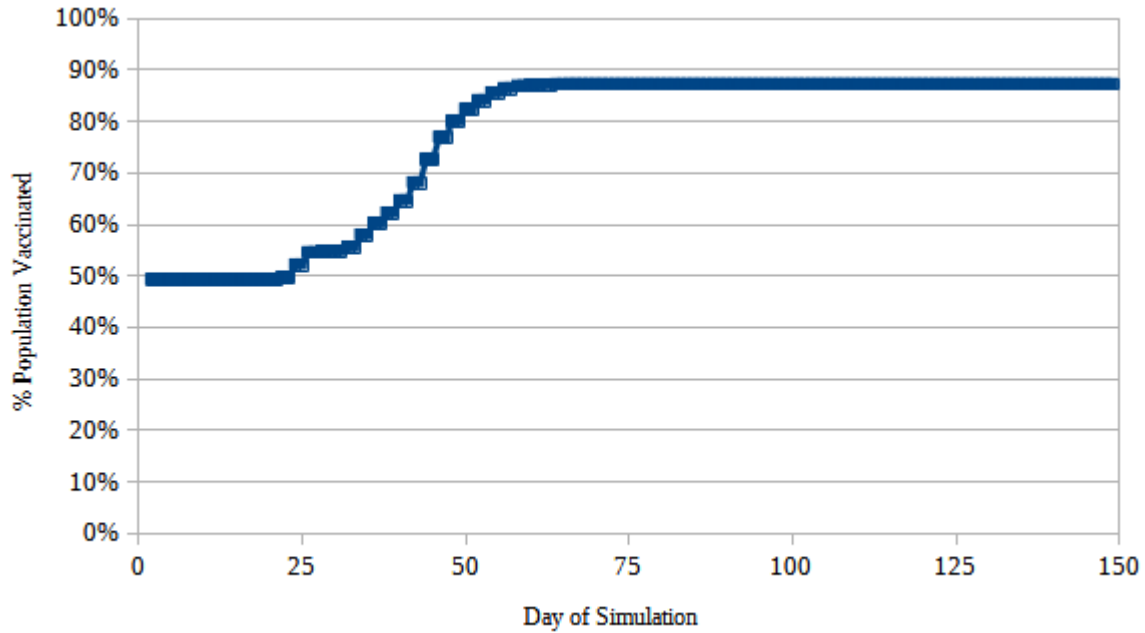
The second feature is the inoculation time of vaccines. Because it takes two weeks to be rendered fully immune, an agent who decides to vaccinate due to an imminent risk will not reap the full benefits of vaccination for two weeks. Because of how quickly measles spread, this may be much too late. For instance, an agent choosing to vaccinate because the population’s infection rate appears high will get the benefit of reduced disease severity, but because it is likely they will become exposed before the two weeks have elapsed, will *not* receive the potential benefit of immunity.

3.1.3. Results for Rational Behavior with Full Information

Figure 4 shows the results for vaccination rates over time from this benchmark model assuming rational behavior with full information. As noted above, I am primarily interested in the implications of making alternative assumptions about the theoretical model for how agents make their vaccination decisions (i.e., bounded rationality’s heuristics, ASD misinformation, and social networks), so the magnitudes presented here should not be considered to have much meaning outside of comparing them

to the outcomes from the subsequent scenarios I present. The parameters were chosen so that the baseline vaccination rate was 50%, and the final vaccination rate was about 90%. The results represent the average values over 1,000 simulations.

Figure 4: Vaccination Path over Time for Rational Behavior with Full Information



While the first measles infection begins on the first day, the initial baseline vaccination rate of 50% is unaffected until the first individual starts to show symptoms around day 25. Thereafter, agents begin vaccinating more and more quickly, until those increased vaccination decisions eventually taper off around day 60.

While most of the action surrounding vaccination happens between days 30 and 60, the disease is spreading through the population before day 30; after day 60, the disease has generally run its course. Despite the occurrence of a deadly measles outbreak, the final vaccination rate of under 90% does not achieve a higher rate due to the heterogeneity of agents (i.e., some have a higher threshold of expected utility costs that they require before vaccinating).

Table 1: Final Outcomes (p. 23) shows the results for the four main outcomes I use in my comparison of this benchmark case of rational behavior with full information to the other scenarios (incorporating heuristics, misinformation, and social networks). Vaccination Rate is the percentage of agents who received a vaccine by the time the simulation ends; it is 87.2% here. Attack Rate is the percentage of agents who ever became ill during the simulation; it is 47.1% here. Mortality Rate is the expected number of deaths; it is 1.83 deaths per 10,000 people here. Social Cost per Capita is the total cost to society, as derived from the price paid for the vaccine, the costs of adverse reactions, the cost of measles treatment (i.e., outpatient and inpatient care), the willingness-to-pay measure to avoid disability, and the willingness-to-pay measure to avoid death shown in Figure 1. Based on results not shown, 54.8% of the \$4,082 social cost is from mortality, and 43.8% is from disability.

Table 1: Final Outcomes Assuming Rational Behavior with Full Information

Vaccination Rate	Attack Rate	Mortality Rate Per 10,000	Social Cost per Capita
87.2%	47.1%	1.83	\$4,082
(0.16%)	(0.16%)	(0.005)	(\$10)

*Numbers in parentheses represent 95% confidence intervals.

3.2. Incorporating Heuristic Behavior

3.2.1 Modeling Heuristics

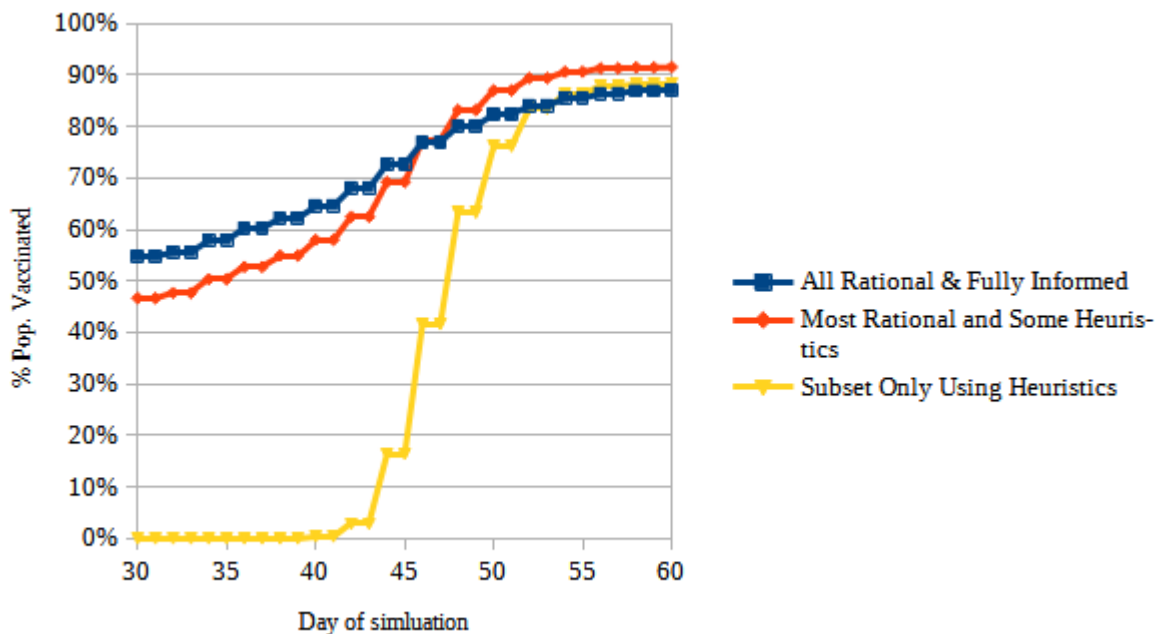
I now examine the effects of allowing some agents to use a simple heuristic instead of making the calculation of expected costs consistent with rational behavior. I assume that 15% of agents use a heuristic to decide whether to vaccinate, while the remaining 85% decide whether to vaccinate if $E[C_{UNV,t}] - E[C_{VAC,t}] > \varepsilon$ as outlined in Section 3.1. Specifically, I assume that heuristic users vaccinate if

the observed symptom incidence rises above a certain threshold, \bar{s} ; i.e., their decision rule is: *vaccinate* if $s_t > \bar{s}$.¹⁶ The remaining model for the disease spread functions as described above in Subsection 3.1.

3.2.2 Results Incorporating Heuristics

Figure 5: Vaccination Path over time for Some Heuristic Users and Full Information (p. 24) shows the vaccination path over time during a key period of the simulation. The blue line represents the benchmark results presented in Section 3.3 from assuming the entire population exhibits rational behavior with full information. The yellow line shows the vaccination rate for the 15% subset who use a heuristic (i.e., they do not vaccinate until a critical threshold is reached), and the orange line shows the average vaccination rate across the combined sample of mostly rational behavior and some heuristic users. As observed in the previous section, the changes in vaccination rates generally occur between days 30 and 60 of the simulation, so my figure here (and similar figures below) focuses on this window.

Figure 5: Vaccination Path over time for Some Heuristic Users and Full Information



¹⁶ This value is assumed. I am not aware of any literature that examines what heuristics the general population uses to judge their risk from disease. As such, given the lack of data, a value is assumed.

The value for the heuristic threshold parameter, \bar{s} , is chosen so that the vaccination rate's average over the simulation's 150 days is approximately the same as the average vaccination rate in the benchmark simulation. However, the nature of the heuristic decision rule implies that the vaccination rate across the entire population is relatively lower than the benchmark during the initial periods, but it is relatively higher than the benchmark during the later periods.¹⁷

Table 2: Final Outcomes with Some Heuristics (p. 25) shows the results for the four main outcomes considered. The observation that the final vaccination rate is higher for the assumption of some heuristic users (i.e., 91.5%) than for the benchmark assumption of fully rational agents (87.2%), but the attack rate is also higher (i.e., 55.1% versus 47.1%) may seem counter intuitive at first (since one would generally expect higher vaccination rates to lower the attack rate); but this increased attack rate occurs because the heuristic agents essentially waited too long to vaccinate. The relatively higher attack rate is consistent with the higher mortality rate (i.e., 2.04 versus 1.83 deaths per 10,000) and the higher social costs (i.e., \$4,537 versus \$4,082).

Table 2: Final Outcomes with Some Heuristics

	Vaccination Rate	Attack Rate	Mortality Rate Per 10,000	Social Cost per Capita
All Rational & Fully Informed	87.2% (0.16%)	47.1% (0.16%)	1.83 (0.005)	\$4,082 (\$10)
Some Heuristics & Fully Informed	91.5% (0.16%)	55.1% (0.17%)	2.04 (0.005)	\$4,537 (\$10)

*Numbers in parentheses represent 95% confidence intervals.

I note that, while these results seem to indicate that heuristics are unambiguously bad for society, one could imagine a parameter set in which this is not the case. For a sufficiently low value of \bar{s} , we would expect heuristics user behavior to be equal to the average rational agent behavior, with the only

¹⁷ Note that this figure shows the average vaccination rates by period across 1,000 simulations. However, in each individual simulation, the change in the vaccination rate is more monolithic as the incidence passes the threshold. That threshold is reached on a different day for each simulation. For instance, on day 45, roughly 20% of simulations had all heuristics users vaccinated, and 80% of simulations had no heuristic users vaccinated.

difference being vaccinating *en bloc* rather than piecemeal¹⁸. If we made the value of \bar{s} heterogeneous across agents, and distributed sufficiently, the two populations could be made indistinguishable. The threshold value of 7.5% was chosen to illustrate the nature of the model, not for any significance.

3.3. Incorporating Autism Misinformation

3.3.1 Modeling Misinformation

I now examine the effects of allowing some agents to be affected by misinformation about the debunked link between vaccines and ASD. I compare this case here to my initial assumption of rational behavior with full information. In this case here, I assume that everyone makes a rational decision, albeit a potentially misinformed rational decision. (In the initial version of this model incorporating misinformation, no agents use a heuristic.)

A recent study found that 16% of caregivers of a child with ASD believe that vaccines contributed to causing the disease (Fombonne et al., 2020). I therefore assume that 16% of agents in my model believe that there is some probability y that receiving a vaccine will induce ASD. Moreover, among those 16%, I assume that agents have beliefs of varying strength, as represented by different y_i drawn from a distribution. Each agent with misinformation now faces an additional perceived cost; they now treat their perceived cost of vaccination as $A_V^* = A_V + y \times ASD$. (For the remaining 84%, $y_i = 0$.)

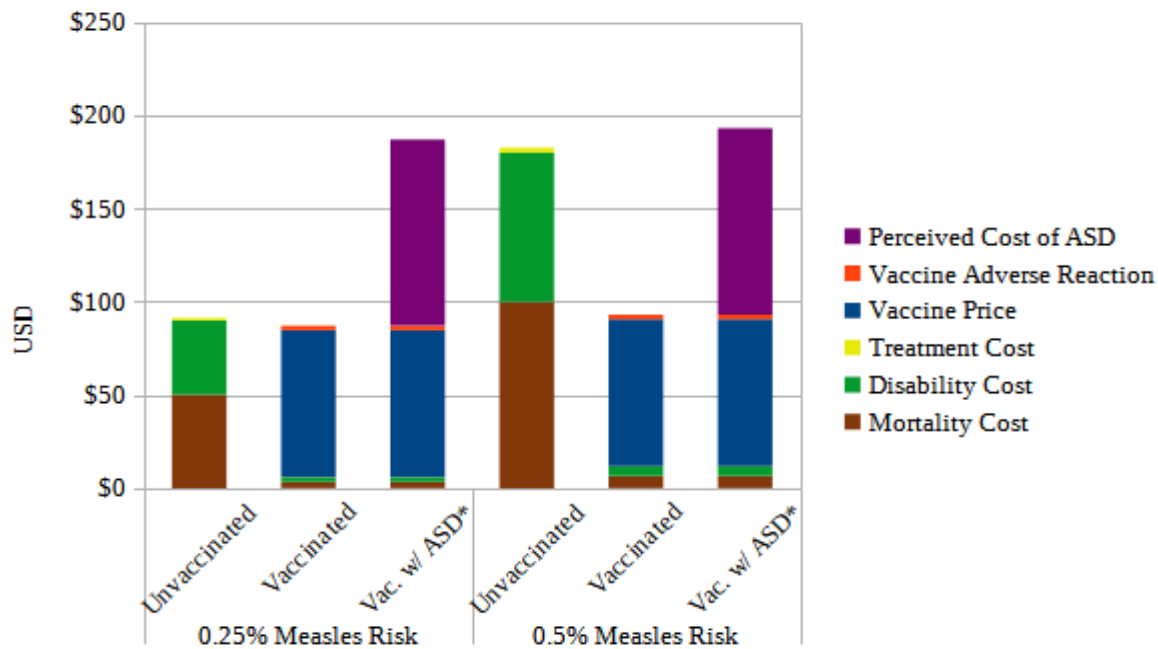
Both the well-informed and the misinformed exhibit rational behavior specified by the vaccination decision rule: $E[C_{UNV,t}] - E[C_{VAC,t}^*] > \epsilon$. Note that while actual costs are unchanged, misinformation about A_V^* does enter into their perception of costs and thus their decision rule. A visualization of this change to incorporate misinformation can be seen in Figure 6: Expected Cost of Measles Given

18 Indeed, for a value of \bar{s} lower than the level of risk that makes a rational agent indifferent to vaccination, I would expect to see *improvement* to society from using heuristics.

Vaccination Status, Risk, and ASD Belief (p. 27). Note that for even a very small belief in vaccine-induced ASD risk, a rational agent would not vaccinate until the measles risk became extremely high (i.e., higher than the 0.5% shown here).

(The remainder of the SEIR model for disease spread follows as in Section 3.1.2. Modeling Disease Spread (p. 19).)

Figure 6: Expected Cost of Measles Given Vaccination Status, Risk, and ASD Belief



* Belief in 0.01% chance of Vaccine Induced ASD

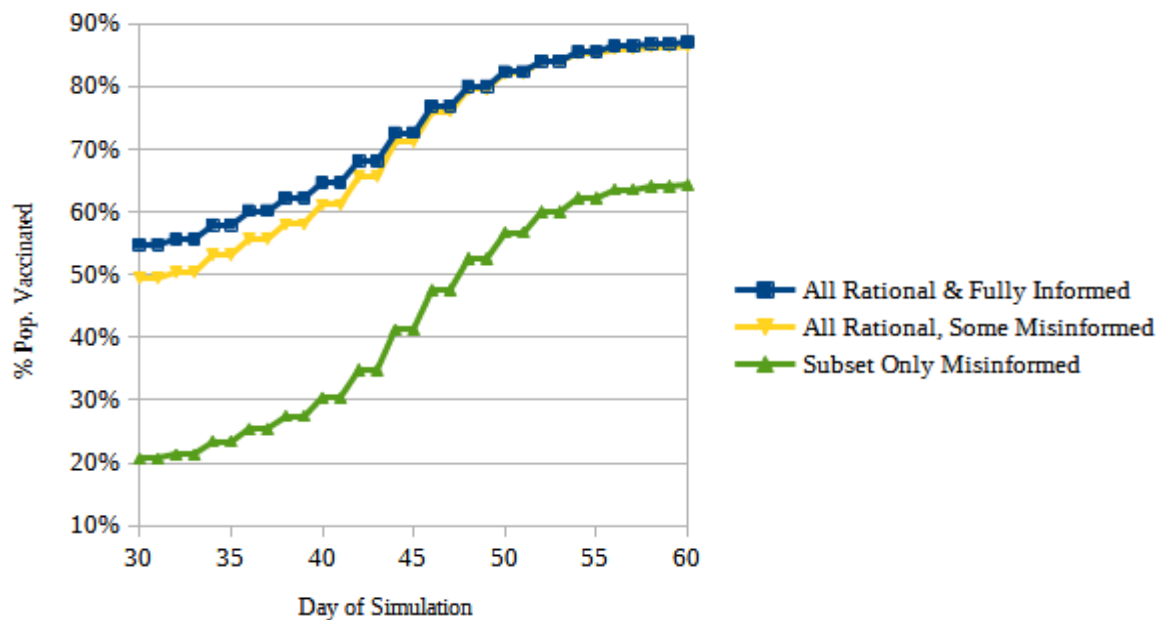
3.3.2 Results with Misinformation

Figure 7: Vaccination Path over time for Fully Rational Users with Some Misinformation (p. 28) shows the vaccination path over time during the simulation. As before, the blue line represents the benchmark results presented in Section 3.1.3. Results for Rational Behavior with Full Information (p. 21) with rational agents with full information. The green line shows the vaccination rate for the 16% portion of

the population which is rational but misinformed¹⁹ about the link to ASD ($y \neq 0$), and the yellow line shows the average vaccination rate across the combined sample of some fully-informed and some misinformed users. As discussed previously, the figure focuses on days 30 to 60 of the simulation.

As expected, misinformed agents vaccinate at a much lower rate than fully informed ones. Despite this difference, some misinformed agents believe the link is weak enough to justify vaccinating before the outbreak, and the majority of the misinformed agents are convinced to vaccinate by the end. However, across the entire population, vaccinations occur later than one composed only of fully informed individuals.

Figure 7: Vaccination Path over time for Fully Rational Users with Some Misinformation



19 I use this term to differentiate between two causes of sub-optimal decisions. Agents can arrive at incorrect conclusions either through a rational decision rule with bad information, or by using an irrational decision rule. A heuristic user represents this latter case, where the method they use to decide is flawed. Such an agent’s actions are inconsistent with maximizing their utility. An agent misinformed about ASD uses a sound method of decision making, but mistakenly incorporates misinformation regarding ASD into their otherwise correct cost-benefit analysis. This method, while misinformed, is a sound one and, if the agent takes the misinformation as true, is consistent with minimizing their costs.

Table 3: Final Outcomes with Some Misinformation (p. 29) shows the results for the four main outcomes. The final vaccination rate has dropped slightly (86.5% versus 87.2%), which is consistent with a minority of the population viewing vaccination as more costly. The attack rate is much larger (53.0% versus 47.1%), due to both the late vaccination from vaccine hesitancy, and the fact that some individuals never vaccinate. The higher attack rate results in higher mortality (2.06 versus 1.83) and social cost per capita (\$4,573 vs \$4,082).

Table 3: Final Outcomes with Some Misinformation

	Vaccination Rate	Attack Rate	Mortality Rate Per 10,000	Social Cost per Capita
All Rational & Fully Informed	87.2% (0.16%)	47.1% (0.16%)	1.83 (0.005)	\$4,082 (\$10)
All Rational & Some Misinformation	86.5% (0.16%)	53.0% (0.16%)	2.06 (0.005)	\$4,573 (\$11)

*Numbers in parentheses represent 95% confidence intervals.

3.3.3 Results Combining Both Heuristics and Misinformation

What then happens when I combine both heuristics and misinformation? To do this I have four groups of agents, for whom their rationality and information status are simultaneously determined. Recall that 85% of agents are fully rational, and 84% are fully informed. Therefore, 71.4% (i.e., 0.85×0.84) are fully rational and fully informed, 13.6% are rational but misinformed, 12.6% use a heuristic yet have full information, and 2.4% use a heuristic and are misinformed²⁰.

Table 4: Final Outcomes With Heuristics and Misinformation (p. 31) presents the results for all four outcomes across the four permutations of heuristics and misinformation, and Figure 8: Social Cost per Capita for Heuristic and Misinformation Permutations (p. 30) focuses on the Social Costs Per Capita

²⁰ Note that because the cost of ASD would not factor into a heuristic decision rule, the misinformed heuristic users will act identically to the fully informed heuristic users and not ultimately be affected by misinformation.

outcomes across those four scenarios, as that one measure summarizes the general pattern of the results quite well. The error bars in the figure represent the 95% confidence intervals.

As expected, because heuristics and misinformation each result in worse outcomes independently (i.e. \$4,537 and \$4,573 versus \$4,082, respectively), combining the two yields even worse outcomes (i.e., \$4,931 versus \$4,082). Interestingly though, the two effects are not additive. This is because the heuristic user who misbelieves there is a vaccine-induced ASD link does not actually incorporate this misinformation into their non-rational decision making; their decision rule is based only on symptomatic incidence.

Figure 8: Social Cost per Capita for Heuristic and Misinformation Permutations




Table 4: Final Outcomes With Heuristics and Misinformation

	Vaccination Rate	Attack Rate	Mortality Rate Per 10,000	Social Cost per Capita
All Rational & Fully Informed	87.2% (0.16%)	47.1% (0.16%)	1.83 (0.005)	\$4,082 (\$10)
Some Heuristics & Fully Informed	91.5% (0.16%)	55.1% (0.17%)	2.04 (0.005)	\$4,537 (\$10)
All Rational & Some Misinformation	86.5% (0.16%)	53.0% (0.16%)	2.06 (0.005)	\$4,573 (\$11)
Some Heuristics & Some Misinformation	90.8% (0.11%)	60.3% (0.17%)	2.22 (0.005)	\$4,931 (\$11)

*Numbers in parentheses represent 95% confidence intervals.

3.4. Incorporating Social Networks

3.4.1 Modeling Social Networks

I now examine the effects of incorporating social networks into the models for vaccination decisions considered in the prior three sections. As indicated above, I expect that social networks can have two distinct effects on vaccination decisions.

First, they can provide more accurate information about the risk of catching measles from agents. For the models of vaccination behavior in the prior sections, I assume that agents based their measure of risk based on the percent of all 500 agents showing symptoms (i.e., a “global” average incidence).

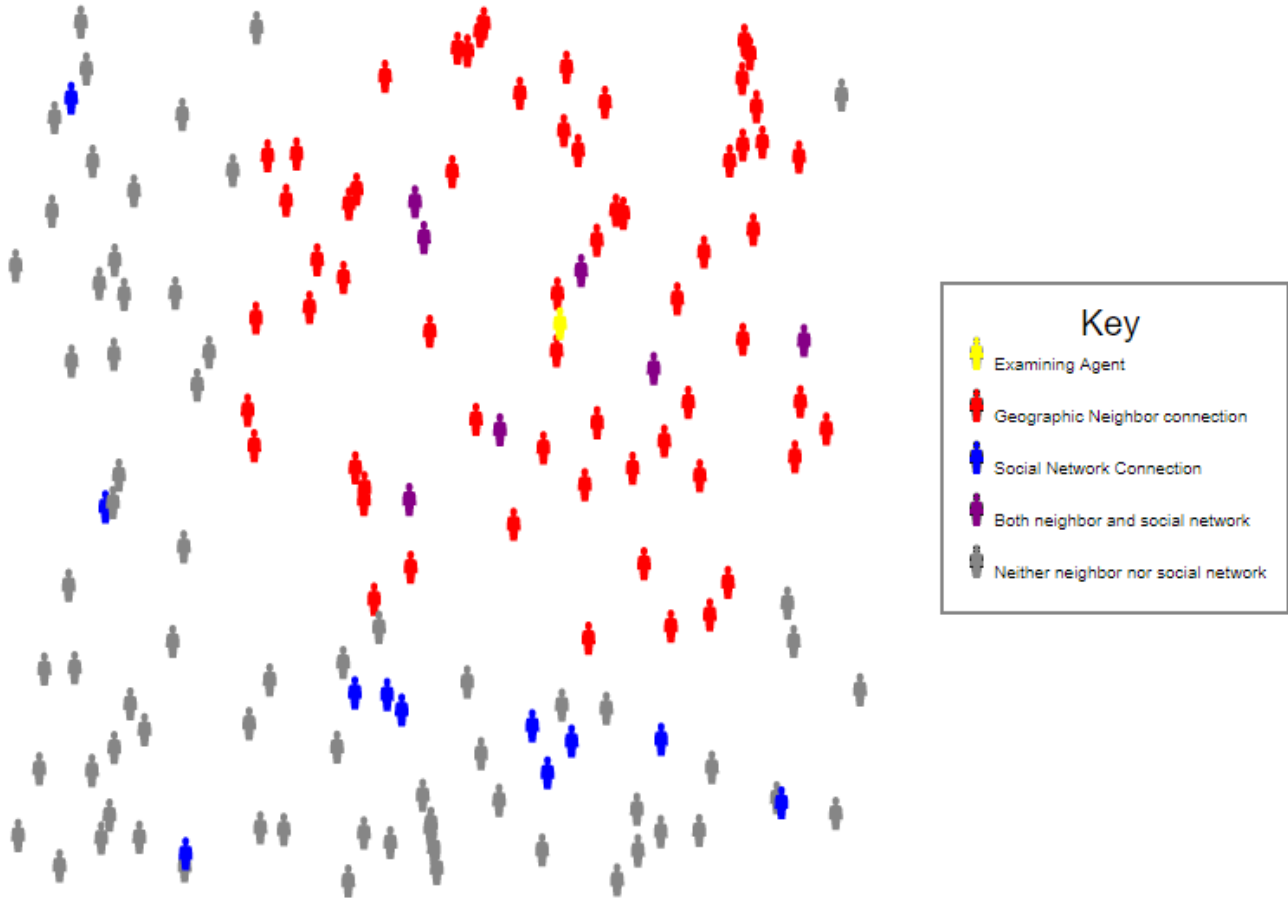
However, an agent’s social network likely includes neighbors who pose a more immediate threat of infection if ill due to their geographic proximity, so an agent who incorporates the percent of agents in their social network showing symptoms (i.e., a “local” average incidence) into their measure of risk will have improved information about the likelihood of catching measles.

Second, social networks can, in principle, provide more accurate information about the “risk” of vaccine-induced ASD. For one of the 16% of agents misinformed about the link to ASD, a social

network with relatively few agents sharing that misbelief might temper their strength of belief (which I model as the percent of ASD cases “caused” by vaccines) while a social network with relatively more agents sharing that misbelief might increase their strength of belief. Moreover, I expect the effect of social networks on one’s belief about vaccine risks to be asymmetric; i.e., decreasing misinformation will have a smaller impact than exacerbating misinformation.

I therefore incorporate both of these aspects of a social network into my model for vaccination decisions. For each agent i , there is one set of agents who are geographically close, whom I label “neighbors.” The neighbors represent any agent with whom an individual has regular in-person contact. Each agent i also has a second set of agents that are randomly distributed throughout the geometric space, whom I label social media “friends.” Every other agent belongs to one, both, or neither of these two categories for agent i . Figure 9: Visualization of a Social Network (p. 33). Details of how this network is formed can be found in Appendix A: Modeling Specification (p. 74).

Figure 9: Visualization of a Social Network



Each individual now has two data sets to draw from: they can observe the global symptomatic rate S_t , as would be reported by a governing agency, and can also calculate the symptomatic rate within their network, $s_{i,t}$.²¹ When performing vaccination decisions, agents use a weighted average of the two, such that $s_t = (S_t + w s_{i,t}) / (1 + w)$. I assume a $w > 1$ to reflect that an agent is likely to weigh local information more heavily than a global average.

21 Agents draw from a pool of both neighbors and friends to determine $s_{i,t}$. I argue that this represents people using social media to gather information on their local area at a very low cost, as well as keeping tabs on the wider world via more distant friends.

For misinformation about vaccine-induced ASD, I model the effect of social networks using the following approach. Recall that I assume that 16% of agents believe that there is some nonzero probability, y_i , that receiving a vaccine will induce ASD. I now allow that agent i 's belief to increase or decrease based on whether their social network's average belief, Y_i , is higher or lower than the average belief across the entire population, \bar{y} . (This includes all the $y=0$ values from agents who do not believe in any such link.) If $Y_i < \bar{y}$, one's network is relatively well-informed about vaccine-induced ASD, and i 's belief used to compute costs is reduced by an amount proportional to that ratio. If $Y_i > \bar{y}$, one's network is relatively poorly-informed, and i 's belief used to compute costs is increased by an amount proportional to that ratio. I assume that the magnitude of these changes is asymmetric (i.e., the increase is larger in magnitude than the decrease) to reflect the ideological "echo chamber" described above.

3.4.2 Results with Social Networks

Figure 10: Social Cost per Capita with Social Networks




Figure 10: Social Cost per Capita with Social Networks (p. 35) shows the results of incorporating social networks into the models presented above. As expected, giving agents access to better information in the fully-informed scenarios results in better outcomes, as agents can vaccinate sooner when an outbreak appears to threaten them personally. For the benchmark model of all rational agents with full information, the presence of social networks decreases social costs per capita from \$4,082 to \$3,950. For the second model of some heuristic users, the presence of social networks decreases social costs per capita from \$4,537 to \$4,232. Even though some agents use an overly-simple decision rule on whether to vaccinate, basing that decision on the local incidence of symptoms rising above some particular threshold is more beneficial than waiting for the global incidence to rise above that threshold. However, when misinformation is present, these gains from the improved information about the risk of catching measles can be outweighed by the losses due to exacerbating misinformed agents' beliefs

about vaccine-induced ASD. For the third case of all rational agents facing some misinformation, the presence of social networks increases social costs per capita from \$4,573 to \$4,796. For the fourth case of mostly rational agents facing some misinformation, the presence of social networks increases social costs per capita from \$4,931 to \$4,975.

3.5. Interventions

This subsection explores the following interventions: subsidies, education, mass vaccination, ring vaccination, and vaccine ambassadors. Subsidies were examined in the previous chapter, and mass vaccination is a simple mandate from the regulator that all agents vaccinate immediately. Ring vaccination is the process where, once an agent begins showing symptoms, any individuals who have had contact with the agent are immediately required to vaccinate. Education in this context will refer to an intervention designed to change each agent's individual belief that vaccines cause ASD. Social media campaigns instead target an agent's socially-adjusted belief that vaccines cause ASD.

Furthermore, I examine combinations of these interventions, as discussed in 3.5.6. Intervention Combinations (p. 42).

For the purposes of this subsection, the regulator is assumed to be the local government. Its objective is to minimize total damage to society, as measured by total societal cost per capita. It is assumed to be capable of swift action, and to expend whatever funds are required to accomplish its goals. When analyzing the results of these interventions, we use a model with social networks and misinformation, but no heuristics users. In order to better ground these results in empirical fact, for any analysis of regulator interventions, I changed the value for the confirmation bias factor ξ to 3.5, to reflect the data in Gilbey and Hill (2012). All other parameters are as found in Appendix B: Parameterization (p. 85).

3.5.1. Subsidies

I examine how government subsidies might be used to increase vaccination rates. In all my models thus far, agents pay the full price for the MMR vaccine. The government can subsidize vaccines by covering some of the MMR price or provide the MMR vaccine for free. The government could also, in principle, provide individuals with a free vaccine and an additional cash payment in exchange for them vaccinating. Similarly, the government could fine individuals who do not vaccinate; these cases function identically in terms of the difference in expected costs for vaccinating, so I consider the case of the government providing subsidies.

Specifically, I assume that agents face a one-time cash transfer of $T \geq 0$ from the government only if they become vaccinated. This changes their cost of being vaccinated to:

$$E[C_{VAC,t}] = kb \times p_{UNV}(s_t) \times (T_{UNV} + D_{UNV} + M_{UNV}) + P_V - T + A_V$$

I illustrate the results by focusing on the social costs per capita. When the government provides such a cash payment, the money is presumed to come from taxes. Thus, each payment can represent a simple transfer of funds (without a direct social cost per se). These subsidies are assumed to be financed through taxes, and represent transfers of wealth. These transfers are not without loss, however; it is assumed that 32% of each transfer is lost to dead weight loss per Feldstein (1999). These losses are counted as losses to society as a whole, but are not internalized by individual agents.

I examine five subsidy levels: No subsidy (\$0), a subsidy for the full price of the vaccine (P_V), and a full subsidy plus rebates for \$100, \$500, and \$1,000. These correspond to P_V^* values of P_V , 0, -100, -500, and -1000 respectively²².

22 Aside from the fully subsidized costs, these figures have no particular significance, and are chosen strictly for the sake of being round numbers.

$$p_{UNV}(s_t) \times (T_{UNV} + D_{UNV} + M_{UNV}) \times (1 - kb) - P_V^* - A_V > \epsilon_i$$

Equation 1: Agent Decision Rule

3.5.2. Mass Vaccination

Mass vaccination will be a mandate by the regulator that all individuals not medically exempt immediately vaccinate. The price of the vaccine is paid for directly by the regulator. This option is included to illustrate what some public health officials consider to be a best-case response. However, such a solution is seldom politically feasible, so it will be examined only to provide comparison to the others.

This intervention can either be active or inactive. If active, the regulator requires mass vaccination immediately. If inactive, they never mandate a population-wide vaccine. Note that this intervention bypasses the agent's decision rules entirely, as vaccination becomes a requirement imposed externally rather than decided internally.

3.5.3. Ring Vaccination

Ring Vaccination is a process designed to slow the spread of a disease with fewer mandated vaccinations than the mass vaccination scenario. In this process, once an agent begins to show symptoms, any individual they could have had contact with are immediately instructed to vaccinate. This presents several advantages and disadvantages.

The chief disadvantage is that this is necessarily a reactive strategy; a ring vaccination only triggers when an agent has begun to show symptoms, which means that agent has already been spreading the disease for several days. As a result, it is likely that the disease will continue to spread.

The main advantage is that it requires individuals most at risk of exposure to vaccinate immediately. This can dramatically slow the spread of a disease, and even individuals who have been exposed gain a benefit. Although the vaccine's main benefit is prophylactic, it also has a therapeutic component. If a patient receives a vaccine after being infected, but before they show symptoms, they will still face the disease at a reduced severity (Gershon, 2015).

A secondary advantage is that such a strategy is more politically feasible. Because it responds to a clear and proximate danger, and because it doesn't overtly require mass participation, such a solution is more likely to be more politically palatable than mass vaccination.

This intervention can either be active or inactive; if active, it triggers every time an agent begins to show symptoms. If inactive, it never triggers. As above, note that this intervention bypasses the decision rule of affected agents.

3.5.4. Education

The regulator may attempt to educate the masses. The idea is simple: if agents have bad information, get them better information. However, this is the most complicated intervention from a mechanistic viewpoint: what exactly does education do? For the purposes of this paper, I assume it changes behavior in one way only: it will change the agent's expectation that vaccines cause ASD (hereafter called "ASD belief").

Recall that each agent has a heterogeneous belief that vaccines cause ASD at a probability of y_i . For 84% of agents, $y_i=0$ (Fombonne et al., 2020), and for the remainder it is assumed that $y_i \sim N(0.25\%, 0.1\%)$.

As observed by Attwell & Freeman (2015), the impact of education on attitude is not monotonic; while 77% of those educated reported improved attitudes toward vaccination, the remaining 23% became more set in their views. I model this as follows.

When an individual is subjected to a public education program, they receive a draw from a random exponential distribution ($\lambda=0.25$)²³, denoted D , which is used to modify their perceived costs to the vaccine. This will do two things: first, it will influence their beliefs with respect to ASD; second, it can make them accept vaccination in a similar way to a neighborhood effect. The shape of the distribution is chosen to be consistent with Sobkowicz (2016), though a truncated normal distribution and beta distribution have also been employed in this role.

$$\hat{y}_i^* = y_i^* / D$$

Equation 2: Impact of Education on ASD Belief

With respect to the ASD outcome, the draw from the exponential distribution will flatly divide their expected outcome of ASD. This is shown in Equation 2: Impact of Education on ASD Belief (p. 40), where ASD_{post} is the post-education expectation of ASD from a vaccine, ASD_{pre} is the initial (pre-education) expectation, and D is the random draw as described above. Note that this means that some individuals will believe ASD is more likely because of the program, perhaps believing in a government or big pharma cover-up. This is deliberate, and consistent with Attwell & Freeman (2015). The results of the model are insensitive to incorporating this backlash effect, as shown in Appendix D.3 Education Backlash (p. 104).

23 This λ is chosen such that 23% of the draws lie below 1, and the remainder above, meaning 77% of the draws show a reduction in perceived probability, while 23% see an increase, in order to stay consistent with Attwell and Freeman.

This intervention has three levels, defined by how much of the population (selected at random) is subjected to a public education program: 0%, meaning no education is in place, 50% which implies half the population is educated, and 100% which implies the entire population is exposed to a public education program. An education program faces costs of \$21 per person educated (Karanth et al., 2017)²⁴. This intervention may be used in concert with others.

3.5.5. Vaccine Ambassadors

This intervention is a social media intervention, designed not to change agent's inherent beliefs y_i , but rather their socially adjusted belief y_i^* , using peer effects to change behavior rather than pure information. They do this by educating agents on vaccine safety and giving them a media platform, and so “dilute” the impact of the misinformed by increasing the quantity of informed discussion.

I call these selected agents “vaccine ambassadors”. Such an approach has been shown to produce a 2.7 percentage point increase in vaccine uptake (Attwell and Freeman, 2015). For the purposes of this model, the regulator will disseminate materials featuring the vaccine ambassador at the same cost as the education intervention, above. The impact is described below.

Recall that each agent has a socially-adjusted ASD belief of y_i^* . This belief stems from the average belief level in their social network, Y_i , and to what degree their network differs from the global mean \bar{y} , as defined by a percent difference between the two y_{dif} . If their network is, on average, well-informed relative to the global mean, (\bar{y}) then their own belief is weakened. That is, if $Y_i < \bar{y}$, then $y_i^* = (1 - y_{dif}) y_i$. Conversely, if their network is worse informed relative to the mean, it amplifies their misinformed belief. Furthermore, it does so asymmetrically. So if $Y_i > \bar{y}$, then $y_i^* = (1 + \xi y_{dif}) y_i$ for some $\xi > 0$.

24 This study gives a range of possible values. As the size of the targeted population increases, the per-person cost decreases. I use the cost for lowest population size given in the study, as it most closely matches my model. However, education may prove more cost-effective as the targeted population increases.

Previously, Y_i was defined by $Y_i = 1/n \times \sum y_j$, where n is the total number of individual's in the agent's social network. Each agent who is targeted by a vaccine ambassador intervention has that $y_j = 0$ voice "counted" α extra times. Thus, for an agent i who has received a vaccine ambassador intervention, $Y_i = 1/(n + \alpha) \sum y_j$. The value for α is determined by calibration. (See Appendix C.1 Calibration, p. 93)

This intervention is investigated at five levels, each corresponding to what percentage of the population the regulator exposes to the vaccine ambassador materials: 0%, 5%, 25%, 50%, and 100%. This intervention may be attempted in concert with other interventions.

3.5.6. Intervention Combinations

In addition to investigating how each intervention performs in isolation, I also examine select combinations. A summary of all investigated scenarios are described in Table 5: Scenario Summary (p. 42). The leftmost column is the name by which I refer to the scenario, and the remaining columns indicate which interventions are active, and to what level. The subsidy level indicates a dollar value for a payment given to agents who vaccinate. The education and vaccine ambassador entries indicate what percentage of the populace is subjected to that intervention. Targeted and mass vaccination are either "yes" if the intervention is active, or "no" if it is not.

Table 5: Scenario Summary

Scenario Name	Subsidy	Education	Vaccine Ambassadors	Ring Vaccination	Mass Vaccination
Baseline	0	0	0	No	No
Full Subsidy	78.68	0	0	No	No
\$100 Rebate	178.68	0	0	No	No
\$500 Rebate	578.68	0	0	No	No
\$1,000 Rebate	1078.68	0	0	No	No
50% Education	0	50	0	No	No

Scenario Name	Subsidy	Education	Vaccine Ambassadors	Ring Vaccination	Mass Vaccination
100% Education	0	100	0	No	No
5% Vaccine Ambassador	5	0	0	No	No
25% Vaccine Ambassador	25	0	0	No	No
50% Vaccine Ambassador	50	0	0	No	No
100% Vaccine Ambassador	100	0	0	No	No
Ring Vaccination	0	0	0	Yes	No
Mass Vaccination	0	0	0	No	Yes
Rebate + Edu	578.68	100	0	No	No
Rebate + Vac. Amb.	578.68	0	25	No	No
Rebate + Ring	578.68	0	0	Yes	No
Edu + Vac. Amb.	0	100	25	No	No
Edu + Ring	0	100	0	Yes	No
Vac. Amb. + Ring	0	0	25	Yes	No
Rebate + Edu + Vac. Amb.	578.68	100	25	No	No
Rebate + Edu + Ring	578.68	100	0	Yes	No
Rebate + Vac. Amb. + Ring	578.68	0	25	Yes	No
Edu + Vac. Amb. + Ring	0	100	25	Yes	No
All Interventions	578.68	100	25	Yes	No

5. Intervention Results

To determine the results of the interventions, I examine the per capita social cost for each intervention described in Section 3.5.6. Intervention Combinations (p. 42). For each figure, the error bars represent a 95% confidence interval for the 1,000 simulations run²⁵. Select summaries are shown in this section.

This section only shows results for total societal costs, for as described above, the other metrics tend to

²⁵ See Appendix A.7. Stochasticity (p. 84) for a detailed discussion on the origin of this variance.

be reflected in total societal cost. For a complete table of results, including the performance of other metrics, see Appendix E: Supplementary Tables (p. 106).

5.1 Mandated Vaccines

I first examine the impact of mandating vaccines. This includes mass vaccination, for which all but the medically exempt agents are required to vaccinate before the outbreak, and ring vaccination, in which anyone potentially exposed to a newly symptomatic individual is required to vaccinate.

Figure 11: Results for Mandated Vaccines




Figure 11: Results for Mandated Vaccines (p. 44) shows that both ring vaccination (\$1,628) and mass vaccination (\$95) show a marked improvement over the baseline no intervention case (\$4,764). This is hardly surprising; by bypassing the agent's decision making process, any flaws in that process are similarly avoided. This can be seen by how the mass vaccination costs is only slightly above the individual expected costs of vaccinating; this shows that it stops an outbreak very quickly, and very few

cases of measles are experienced. While both interventions result in high vaccination rates, the attack rate for the mass vaccination case is almost zero, while the attack rate for the ring vaccination intervention is 23.7%. (See Table 11: Results of Single Interventions, p. 106 in Appendix E: Supplementary Tables for details.) Two further details are worthy of note.

First, ring vaccination represents a marked improvement over the no intervention case, meaning that if mass vaccination is not feasible, ring vaccination may be a highly effective strategy. However, the model does not incorporate the costs of contact tracing and coordination, which may be substantial. Armbruster and Brandeau (2007) find a minimum annual cost of maintaining such a capacity to be \$18,000²⁶.

Second, the mass vaccination is assumed to occur before the outbreak. However, this model assumes that an outbreak is inevitable. The ring vaccination strategy has zero costs if an outbreak does not occur²⁷, while the mass vaccination continues to have the vaccination costs accrue even in the absence of an outbreak. A real-world regulator must weigh the probable savings during an outbreak against the constant costs incurred regardless of disease spread.

5.2 Subsidized Vaccines

I now examine the impact of subsidizing the price of vaccines. The levels included are a full price subsidy, making vaccines free for agents, and providing free vaccines as well as a \$100, \$500, and \$1,000 incentive to agents to vaccinate. Costs to society include the deadweight loss from these transfers.

26 This value is particularly portable as Armbruster and Brandeau (2007) also use a population of 500 agents for their study. Furthermore, this value maintains a contact tracing capacity only of 3; during peak disease spread, this would be insufficient in my model to adequately apply ring vaccinations.

27 Recall that this model does not implement the \$18,000 annual cost of contact tracing infrastructure described above.

Figure 12: Results for Subsidies




Figure 12: Results for Subsidies (p. 46) shows that while simply making the vaccines free has no impact on total costs to society, larger rebates have progressively higher gains. In principle, there is a larger subsidy amount at which a marginal increase in that subsidy instead *increases* total societal cost (because the additional deadweight loss from those subsidies is larger than the marginal improvements in measles disability and mortality). However, even very large subsidy levels (such as \$10,000) continue to lead to decreases in societal costs in this model. Because of the political impracticability of such large subsidies (or penalties)²⁸, for the remainder of the paper, I acknowledge that such a minimum social cost theoretically exists, but only examine the \$500 rebate level for the remainder of the dissertation.

²⁸ Larger subsidies up to \$2,000 have been studied (Robertson et al., 2021). However, given that the largest rebate I could find record of was \$100 (Fischels, 2021), I limit myself to values at or below \$1,000 as feasible.

Noteworthy is that these changes in cost do not represent changes in final vaccine uptake, which is nearly identical across all cases, but only *when* agents choose to take the vaccine. If we examine only the final vaccination rates, we find a completely inelastic price elasticity of demand for prices. This is consistent with Kondo, Hoshi, and Okubo (2009), who empirically find nearly zero elasticities for the influenza vaccine. However, this may not tell the whole story. It is more informative to examine how rebates change demand *before* an outbreak.


By selecting the vaccination rates on day 10, before the first agent begins to show symptoms, I can estimate the effect of only price on the demand for vaccines. I find a 50.50% ($\pm 2.27\%$) vaccination rate for the \$1,000 rebate case on day 10, and find only a 46.91% ($\pm 2.38\%$) vaccination rate for the \$500 rebate case on the same day. This 3.59% difference in quantity consumed, divided by the 100% increase in price, gives an elasticity of -0.0359. This is within the figures found by Kondo, Hoshi and Okubo when they restrict their analysis by season. This suggests that these findings that rebates must be large are consistent with empirical fact.

A limitation of this model is that the figures above assume that an outbreak is inevitable. In reality, outbreaks are an uncertain occurrence. Subsidies, such as those detailed above, would cause deadweight loss in both the outbreak and the no-outbreak states. The benefits, in the form of reduced outbreak severity, would only accrue in an outbreak state. While a \$1,000 rebate may be cost-effective in an outbreak state, it may not be when an outbreak is uncertain, depending on the frequency of outbreaks²⁹.

29 This problem could be circumvented by using a system wherein rebates are inactive unless a public health organization declares an outbreak. However, two problems arise with this possibility. First, not all cases are reported, and even if they are, exactly when a public health official declares an outbreak may be arbitrary. Second, it creates an incentive for agents to *not* vaccinate until there is an outbreak, then vaccinate immediately, in order to avoid loss of income via the rebates. Thus, I assume here that rebates apply in non-outbreak as well as outbreak states for simplicity.

5.3 Education

Figure 13: Results of Education



For this model, an education intervention is designed to change how likely individual agents are to erroneously believe vaccines are to cause ASD. Figure 13: Results of Education (p. 48) shows that education does not give any gains. There appears to be a slight upward trend in costs to society, presumably driven by the costs of education, but such differences are not statistically significant.

This implies that an education campaign which targets only ASD misinformation is likely to fail. While it improves the views of some agents, the variance in these improvements means some are necessarily not convinced, and these unconvinced agents still fail to vaccinate. This is consistent with findings in the literature (Nyhan et al., 2014; Sadaf et al., 2013).

However, it should be noted that other types of campaigns are possible: campaigns better educating individuals about the risk of not vaccinating could provide another avenue for future study. This would have to be done by providing agents with a better estimate of the probability of infection $p_{UNV}(s_i)$, as the costs of the illness are assumed to be entirely and accurately internalized.

5.4 Vaccine Ambassadors

A vaccine ambassador program targets the social aspect of ASD misinformation. By putting out an advertising campaign showing real individuals in the vaccine hesitant community, a regulator can artificially inflate the perception that one's peers are not worried about ASD. I examine four campaign levels, which seek to reach 5%, 25%, 50%, and 100% of the population. Figure 14: Results for Vaccine Ambassadors (p. 49) shows that such interventions are highly effective. Even the lowest level of a vaccine ambassador program shows significant gains (\$1,228 vs \$4,764).

Figure 14: Results for Vaccine Ambassadors





Figure 15: Detail of Vaccine Ambassador Interventions (p. 50) shows how the interventions compare to each other. While moving from a 5% to a 25% level shows small but statistically significant improvement (\$1,193 vs \$1,228), any higher levels are not significantly different. If anything, they may be slightly higher, owing to the increased cost of the program. This would imply that a 25% level is

enough to begin developing herd immunity, and any further efforts do not result in improved disease outcomes.

The gains from such a program may seem to be out of proportion to the 2.7 pp increase in vaccine uptake shown in Attwell and Freeman (2015). The key difference is that Attwell and Freeman observed such a difference in vaccination in the absence of an outbreak. When an outbreak is present, even that relatively modest decrease in perceived cost of vaccinating results in agents vaccinating sooner.

Because they vaccinate sooner, they incur fewer costs, hence the relatively large difference in behavior from a no-outbreak scenario.

Figure 15: Detail of Vaccine Ambassador Interventions



5.5 Combined Interventions

Due to the large number of intervention combinations, I will only discuss the noteworthy results. A full set of tables for all results can be found in Appendix E: Results Tables.

First, I combine a 25% vaccine ambassador campaign with ring vaccination. Figure 16: Combining Vaccine Ambassadors with Ring Vaccination (p. 52) shows that there are small but statistically significant gains from combining these two interventions. Given that they target different mechanisms, this makes sense; the vaccine ambassador intervention reduces the number of agents hesitant to vaccinate for fear of ASD, while the ring vaccination intervention reduces the impact of the time lag between when agents begin infecting others and when they start showing symptoms. Thus, these two interventions can be considered to be compliments.

Interestingly, this is the only intervention that shows gains to society. All other combinations tested either did not result in a significant change, or else resulted in a *worsening* of outcomes. For instance, Figure 17: Combining Education and Rebates (p. 53) shows that if I combine the ineffective education intervention with a rebate, I see a *higher* cost to society relative to rebates only (\$4,630 versus \$4,602). This cost is driven entirely by the cost of education; attack and vaccination rates as well as mortality remain unchanged. (Details can be found in Table 12: Results of Combined Interventions, p. 107, in Appendix E: Supplementary Tables.) Viewed in this light, we can consider all other interventions to be substitutes to one another.

This further suggests that a \$500 rebate or 25% vaccine ambassador campaign each individually is sufficient to quickly get the population up to herd immunity. Combining either of these with each other or education is a duplication of efforts, which leads to a small but statistically significant increase in costs for no appreciable gain.

Any combination of three or more interventions results in a small rise in total costs. All interventions at once does not appear to be statistically different from any of the three-intervention combinations.

Figure 16: Combining Vaccine Ambassadors with Ring Vaccination





Figure 17: Combining Education and Rebates



5.6 Cost-Effectiveness of Interventions

An important metric for regulators is how cost-effective an intervention is. I define cost-effectiveness here as the ratio of loss averted (measured in dollars) over the amount of expense required by the intervention. Note that this is distinct from simply minimizing total societal cost; here the regulator also has an interest in minimizing its own expenditures.

Table 6: Cost-effectiveness of Interventions (p. 55) shows that by far the most cost-effective approach is using the vaccine ambassador intervention. However, the way costs for the ad campaign scale in this model³⁰ may not be realistic; instead, one could imagine that an advertising campaign carries substantial fixed costs. Thus, even using the most conservative value for the 100% vaccine ambassador intervention, it remains superior even to mass vaccination. Mass vaccination itself remains the next

³⁰ This model assumes that all advertising costs are variable costs.

most cost-effective, followed by ring vaccination. All other interventions have substantially lower ratios.

This picture, however, is incomplete, as it assumes an outbreak is inevitable. Several interventions accrue costs even in a no-outbreak case. For instance, mass vaccination faces the same costs in either case, as it is assumed to happen prior to any outbreak. These per-capita costs are shown in the “No-outbreak Cost” column. They show that while mass vaccination may be preferable in the event of an outbreak, the lower cost in a no-outbreak scenario may make ring vaccinations more appealing if the chance of an outbreak is sufficiently low.

Table 6: Cost-effectiveness of Interventions

Scenario	Expenses	Savings	Ratio	No-outbreak Cost	No-outbreak uptake:
Baseline	\$0.00	\$0.00		\$0.00	
Mass Vaccination	\$78.52	\$4,669.76	59.47	\$78.52	
Ring Vaccination	\$114.44	\$3,136.52	27.41	\$52	
Full Price Subsidy	\$20.56	\$3.93	0.19	\$10.78	42.82%
\$100 Rebate	\$46.65	\$53.76	1.15	\$25.21	44.10%
\$500 Rebate	\$150.65	\$162.55	1.08	\$86.88	46.91%
\$1000 Rebate	\$279.84	\$291.43	1.04	\$174.30	50.50%
5% Vaccine Ambassador	\$1.05	\$3,536.49	3,368.09	\$0.00	
25% Vaccine Ambassador	\$5.25	\$3,571.24	680.24	\$0.00	
50% Vaccine Ambassador	\$10.50	\$3,551.92	338.28	\$0.00	
100% Vaccine Ambassador	\$21.00	\$3,550.43	169.07	\$0.00	
50% Education	\$10.32	-\$3.32	-0.32	\$0.00	
100% Education	\$20.64	-\$11.45	-0.55	\$0.00	

Expenses represents the per-capita expenses inured by the regulator for the intervention

Savings represents the per capita value of averted loss by the intervention

Ratio is computed as Savings/Expenses.

No-outbreak Cost represents the cost of the intervention if there is no outbreak.

No-outbreak uptake shows vaccine uptake driven by rebates in a no-outbreak scenario.

6. Select Sensitivity Analyses

Some sensitivity analyses are performed on the base model in Appendix D: Sensitivity Analyses (p.

97), but a few select ones are performed again with the new model parameters here. All sensitivity

analyses change a single parameter of the model, and run 1,000 simulations for each parameter value.

Unless otherwise specified, all analyses are performed with no heuristics, some misinformation, and social networks active, with no interventions active.


6.1. ASD Expectation

6.1.1. Mean ASD Expectation (μ_{ASD})

I vary the average misinformed belief in the probability that vaccines cause ASD (μ_{ASD}) from 0% to 1% in 0.1% increments. Figure 18: Impact of ASD Mean on Total Societal Cost (p. 57) shows the same relationship discussed in Appendix D: Sensitivity Analyses (p. 97). This implies that the model is fairly sensitive to this value.

Given the scarcity of information on how much of the population is misinformed, it may be prohibitively difficult to find data on the strength of that belief. If such data were found, it could be used to refine this model further. However, even in the absence of hard data on the strength of ASD belief, the range of results varies only by 20% from the most extreme values. Thus, while the model is sensitive to average ASD belief strength, it provides a good upper and lower bound on expected costs.


Figure 18: Impact of ASD Mean on Total Societal Cost



6.1.2. ASD Expectation Variance (σ_{ASD})

As above, the variance of the ASD belief was varied from 0% to 0.5% in 0.05% increments. Figure 19: Impact of ASD Expectation Distribution on Total Societal Cost (p. 58) shows the same relationship observed in Appendix D: Sensitivity Analyses (p. 97).

Figure 19: Impact of ASD Expectation Distribution on Total Societal Cost



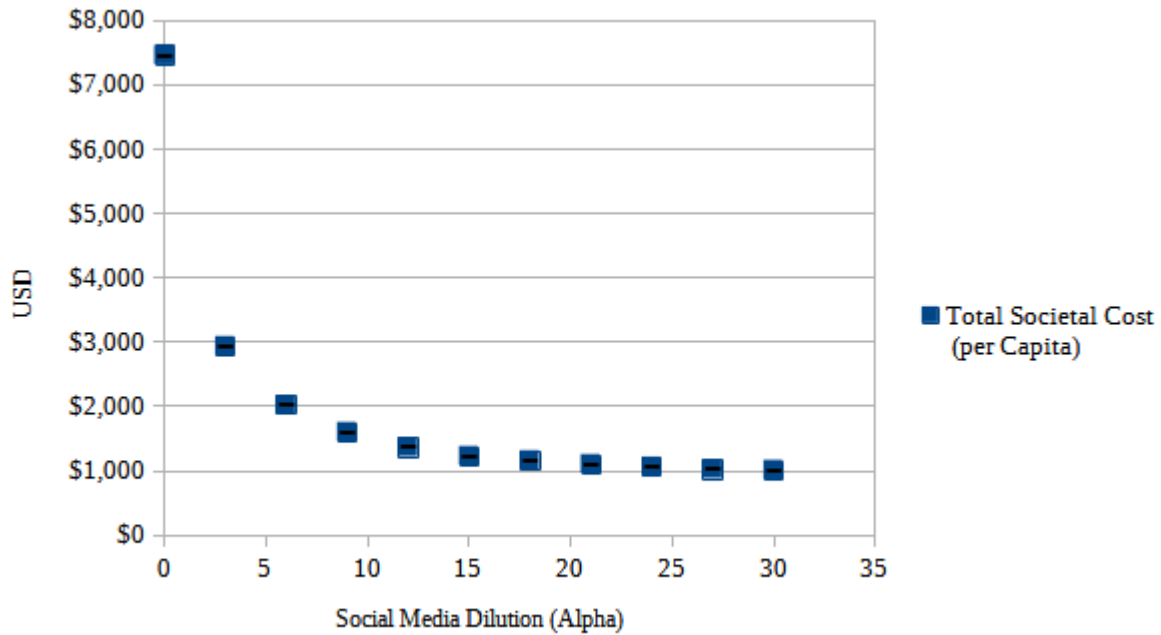
6.2. Social Media Campaign Effectiveness (α)

For this intervention, I set the vaccine ambassador intervention to the 25% level. All other settings are as described above. I do this by changing the effectiveness of the campaign (α). Recall that the vaccine ambassador intervention modifies affected agent’s socially adjusted ASD Expectation by using social media to “dilute” any misinformed agents with the perception of a higher quantity of informed ones. (See 3.5.5. Vaccine Ambassadors, p. 41, for details.)

I vary α from 0 to 30 in increments of 3. Figure 20: Impact of Vaccine Ambassador Effectiveness on Total Societal Cost (p. 59) shows that the initial gains from even a weak campaign are huge, but an increasingly effective campaign quickly sees a reduction in marginal gains. As the campaign becomes stronger, it quickly reaches a point where all but the agents with the highest ASD expectation are

choosing to vaccinate. This intervention is assumed to apply to all agents equally. The more agents that have already been convinced, the more additional effort is wasted as it impacts those who have already decided to vaccinate.

Figure 20: Impact of Vaccine Ambassador Effectiveness on Total Societal Cost



The model’s sensitivity to the effectiveness of the campaign applies mostly to lower levels of effectiveness. Given the calibrated value of 16.2, the empirical value lies firmly in the more effective range, for which the model is insensitive. Thus, I conclude that for likely ranges of campaign effectiveness, the model is insensitive to the effectiveness of the campaign.

7. Conclusion

This dissertation illustrates how vaccine hesitancy can be explained in an economic context, in a way that a rational expectations model would be unable to do. It also examines interventions available to the

regulator and finds that the second-best but politically feasible interventions of vaccine ambassadors and ring vaccination are the most cost effective.

The theoretical model examines self-interested agents making vaccination decisions in the context of a measles outbreak, and finds that three factors can contribute to suboptimal vaccination rates.

Misinformation has an unambiguously negative impact on vaccination rates and subsequent outcomes in the event of an outbreak. Heuristics also shows in my model a similarly negative outcome. While it is possible for a heuristic to be sensitive enough to disease to result in an outcome comparable to a fully rational population, the lack of any data on if or how the population defines its heuristic renders the argument academic, and thus I do not examine heuristics further. Social networks have an ambiguous impact on outcomes. The more relevant local information they provide helps agents make a better estimate of their risk of infection, but in the presence of misinformation that same network's capacity to amplify incorrect beliefs represents a substantial liability, which can outweigh any benefits from more relevant incidence data.

This model provides a framework to understand vaccination behavior which has been poorly explained to date. For instance, the relatively low price elasticity of demand for vaccines seems puzzling at first glance. This model shows that such changes in price fail to account for the substantial indirect costs of vaccination. These costs grow very large, indeed, if they incorporate misinformation regarding vaccine health impacts. I model them here as misinformation regarding ASD, but the model can easily be adapted to include other feared outcomes such as sudden infant death syndrome, multiple sclerosis, or any other ailment which the vaccine hesitant fear. Accounting for these costs shows that even an agent performing a correct cost-benefit analysis on vaccination may refuse to vaccinate if they operate on a false believe in risk of significant side effects, such as ASD.

Having established a theoretical model, I investigate the efficacy of various interventions in the context of a model of rational but potentially misinformed agents³¹. That mass vaccination results in the best outcome is unsurprising; public health officials often advocate vaccinating early and completely. However, if I assume this option is not available to the regulator for whatever reason, I seek to find a second-best solution.

I find that interventions targeting vaccine cost directly are feasible, but only for high levels of subsidy that essentially pay patients to vaccinate. Beyond that, interventions that seek to change agent's inherent beliefs of vaccine safety appear to be of limited use. Campaigns that seek to change the social context of those beliefs, however, appear to have great promise, especially if used in conjunction with ring vaccination techniques.

These results come with a series of caveats. First, the education investigated here specifically targeted agent's belief that vaccines cause ASD; another interesting regime to investigate is one that tries to better inform agents of the probability of infection, $p_{UNV}(S_i)$, which has the potential to change behavior through another channel. Similarly, it may be possible that the way I modeled education is inaccurate; further research into how education changes probability beliefs may shed more light on the subject.

Second, I reiterate that while the mass vaccination option appears cheapest, this is only if I assume that an outbreak will happen. Such an intervention would accrue costs even if an outbreak never happened. To a lesser extent, the rebate intervention would likewise have some costs if implemented before an outbreak. All the other interventions are assumed to be implemented on a reactive basis; when an infection is observed, the campaigns are launched. If no outbreak occurs, the campaigns are not launched, and thus incur no costs. This may explain the lack of support for mass vaccination programs;

31 Due to the lack of data on if or how heuristics are employed, contrasted with reasonable data on ASD misinformation, heuristics are disabled for all intervention analyses.

a topic for future research might be to examine the probability of an outbreak as compared to the expected costs within and outside of an outbreak state.

Third, the Ring Vaccination intervention has costs not shown in this analysis. As mentioned previously, it is estimated that maintaining a contact tracing infrastructure for a group of 500 individuals costs \$18,000 annually, or a per capita annual cost of \$36. This cost is sufficiently low in an outbreak state as to not change our analysis, but if such costs also accrue in the non-outbreak state, then it faces a similar problem to the mass vaccination policy above.

Fourth, my modeling of the vaccine ambassador intervention suggests it as a promising avenue, but there is no literature on the exact mechanism by which the “I Immunize” program, which inspired this intervention, functions. As a result, if the I Immunize program functions through another channel, then my vaccine ambassador intervention may require a different sort of implementation.

Despite these reservations, this paper presents some promising interventions for a regulator to attempt. By accounting for how agents make their decisions, I offer a model consistent with the failure of traditional educational campaigns on vaccine safety, and potential alternatives which may offer effective alternatives to encourage the public to vaccinate. In particular, two policies are worthy of note to regulators.

First, the vaccine ambassador intervention can be implemented in an entirely reactionary context. This means that, in addition to being the most cost-effective intervention from the regulator’s perspective, it also does not have any costs in a no-outbreak state. This makes it a very promising intervention, worthy of further study.

Second, ring vaccination represents a solid compromise between cost-effectiveness, cost in a no-outbreak state, and political feasibility. Unlike mass vaccination, ring vaccination may have an easier

time gaining traction politically, as it does not require participation of the population until an outbreak is in progress, at which point the need will be clear. Its cost effectiveness ratio is smaller, but still an order of magnitude greater than subsidies. Finally, its annual cost per capita is relatively small.

Whatever implementation is chosen, this model provides an explanation for vaccine hesitancy in an economic context that the theory of rational expectations has been unable to provide. In doing so, it provides regulators not only with a menu of options in the event of an outbreak, but also with a framework through which to examine the problem of vaccine hesitancy and better craft and target policy interventions.

Future work can progress around three lines of research. First, the model can be further refined. In particular, building a framework for risk aversion into the model could be relevant, especially given the large costs which occur only at a low order of probability. This could have a significant impact on model outcomes if we assume, as much literature does, that agents are risk-averse. It could also be enlightening to develop a method whereby agents have an even more accurate forecast of disease spread, and allowing their perception of risk to be an exponential function of incidence rather than a linear one.

Second, the model could be adapted to examine other diseases. The model could easily be adapted to examine the impacts of other diseases. Because measles is a deadly and fast-spreading disease, I expect that examining it gives us stark results. However, the framework could also apply to less dangerous diseases such as chickenpox, influenza, or even COVID. By examining whether the model applies to behavior in other diseases, with their own costs and epidemiological characteristics, the model could be validated more thoroughly, and its predictive power better tested.

Finally, this model assumes that vaccine mandates are effective. In the United States, nearly every state mandates the MMR vaccine in order to attend public school, but all such states also have exemptions to

the mandate. A significant literature exists which examines how restricting these exemptions lowers the incidence of such diseases. This model has shown that, even in the absence of mandates, diseases tend to result in cyclical behavior: an outbreak happens, agents vaccinate, and then incidence falls. If removing vaccine mandate exemptions is caused by such outbreaks, then perhaps the resulting decrease in incidence is not a product of removing exemptions, but rather a natural consequence of the outbreak. If disease outbreaks do truly cause the removal exemptions, then the studies examining the effectiveness of those exemptions may not have accounted for that endogeneity, and overstated the impact of exemptions. The CDC began tracking state-level measles data in 2019, so in the coming years enough data should be available to run a more thorough analysis.

References

- Aimone, J.A., Houser, D., 2013. Harnessing the benefits of betrayal aversion. *Journal of Economic Behavior & Organization* 89, 1–8. <https://doi.org/10.1016/j.jebo.2013.02.001>
- Aimone, J.A., Houser, D., 2012. What you don't know won't hurt you: a laboratory analysis of betrayal aversion. *Exp Econ* 15, 571–588. <https://doi.org/10.1007/s10683-012-9314-z>
- Amirthalingam, G., Gupta, S., Campbell, H., 2013. Pertussis immunisation and control in England and Wales, 1957 to 2012: a historical review. *Eurosurveillance* 18, 20587. <https://doi.org/10.2807/1560-7917.ES2013.18.38.20587>
- Armbruster, B., Brandeau, M.L., 2007. Contact tracing to control infectious disease: when enough is enough. *Health Care Manage Sci* 10, 341–355. <https://doi.org/10.1007/s10729-007-9027-6>
- Armstrong, E.P., 2007. Economic Benefits and Costs Associated With Target Vaccinations. *JMCP* 13, 12–15. <https://doi.org/10.18553/jmcp.2007.13.s7-b.12>
- Attwell, K., Freeman, M., 2015. I Immunise: An evaluation of a values-based campaign to change attitudes and beliefs. *Vaccine* 33, 6235–6240. <https://doi.org/10.1016/j.vaccine.2015.09.092>
- Berman, J.M., 2020. *Anti-vaxxers: How to Challenge a Misinformed Movement*. MIT Press.
- Bester, J.C., 2016. Measles and Measles Vaccination: A Review. *JAMA Pediatr* 170, 1209. <https://doi.org/10.1001/jamapediatrics.2016.1787>
- Betsch, C., Renkewitz, F., Haase, N., 2013. Effect of Narrative Reports about Vaccine Adverse Events and Bias-Awareness Disclaimers on Vaccine Decisions: A Simulation of an Online Patient Social Network. *Med Decis Making* 33, 14–25. <https://doi.org/10.1177/0272989X12452342>
- Biellik, R.J., Clements, C.J., 1997. Strategies for minimizing nosocomial measles transmission. *Bull World Health Organ* 75, 367–375.
- Biswas, M.H.A., Paiva, L.T., de Pinho, M., 2014. A SEIR model for control of infectious diseases with constraints. *Mathematical Biosciences & Engineering* 11, 761. <http://dx.doi.org/10.3934/mbe.2014.11.761>
- Brahmbhatt, M., Dutta, A., 2008. On SARS type economic effects during infectious disease outbreaks. *World Bank Policy Research Working Paper*.
- Buchanan, R., Bonthius, D.J., 2012. Measles Virus and Associated Central Nervous System Sequelae. *Seminars in Pediatric Neurology, Viral Infections of the Pediatric Nervous System* 19, 107–114. <https://doi.org/10.1016/j.spen.2012.02.003>

- Callender, D., 2016. Vaccine hesitancy: More than a movement. *Human Vaccines & Immunotherapeutics* 12, 2464–2468. <https://doi.org/10.1080/21645515.2016.1178434>
- Carabin, H., Edmunds, W.J., Gyldmark, M., Beutels, P., Lévy-Bruhl, D., Salo, H., Griffiths, U.K., 2003. The cost of measles in industrialised countries. *Vaccine* 21, 4167–4177. [https://doi.org/10.1016/S0264-410X\(03\)00534-6](https://doi.org/10.1016/S0264-410X(03)00534-6)
- Carabin, H., John Edmunds, W., Kou, U., van den Hof, S., Hung Nguyen, V., 2002. The average cost of measles cases and adverse events following vaccination in industrialised countries. *BMC Public Health* 2, 22. <https://doi.org/10.1186/1471-2458-2-22>
- Chen, M.-F., Wang, R.-H., Schneider, J.K., Tsai, C.-T., Jiang, D.D.-S., Hung, M.-N., Lin, L.-J., 2011. Using the health belief model to understand caregiver factors influencing childhood influenza vaccinations. *Journal of Community Health Nursing* 28, 29–40.
- Christandl, F., 2008. How economic laypeople perceive economic growth and inflation (PhD Thesis). Universität zu Köln.
- Cohen, B.E., Durstenfeld, A., Roehm, P.C., 2014. Viral Causes of Hearing Loss: A Review for Hearing Health Professionals. *Trends in Hearing* 18, 2331216514541361. <https://doi.org/10.1177/2331216514541361>
- Colander, D., Föllmer, H., Haas, A., Goldberg, M.D., Juselius, K., Kirman, A., Lux, T., Sloth, B., 2009. The financial crisis and the systemic failure of academic economics. Univ. of Copenhagen Dept. of Economics Discussion Paper.
- Current CDC Vaccine Price List (Price List), 2021. , Vaccines for Children Program. Centers for Disease Control and Prevention.
- d’Onofrio, A., 2002. Stability properties of pulse vaccination strategy in SEIR epidemic model. *Mathematical Biosciences* 179, 57–72. [https://doi.org/10.1016/S0025-5564\(02\)00095-0](https://doi.org/10.1016/S0025-5564(02)00095-0)
- Daunizeau, J., Moran, R., Mattout, J., Friston, K., 2020. On the reliability of model-based predictions in the context of the current COVID epidemic event: impact of outbreak peak phase and data paucity. *MedRxiv*.
- Davis, M.M., 2010. Price as a double-edged sword in the golden era of vaccines. *Human Vaccines* 6, 689–693. <https://doi.org/10.4161/hv.6.9.13496>
- Day, A., 2021. Jonathan M. Berman, *Anti-Vaxxers: How to Challenge a Misinformed Movement*. *Social History of Medicine*. <https://doi.org/10.1093/shm/hkab003>
- Dayan, G.H., Rubin, S., 2008. Mumps Outbreaks in Vaccinated Populations: Are Available Mumps Vaccines Effective Enough to Prevent Outbreaks? *CLIN INFECT DIS* 47, 1458–1467. <https://doi.org/10.1086/591196>
- De la Sen, M., Ibeas, A., Alonso-Quesada, S., 2012. On vaccination controls for the SEIR epidemic model. *Communications in Nonlinear Science and Numerical Simulation* 17, 2637–2658. <https://doi.org/10.1016/j.cnsns.2011.10.012>

- Dror, A.A., Eisenbach, N., Taiber, S., Morozov, N.G., Mizrachi, M., Zigron, A., Srouji, S., Sela, E., 2020. Vaccine hesitancy: the next challenge in the fight against COVID-19. *Eur J Epidemiol* 35, 775–779. <https://doi.org/10.1007/s10654-020-00671-y>
- Dubé, E., Vivion, M., MacDonald, N.E., 2015. Vaccine hesitancy, vaccine refusal and the anti-vaccine movement: influence, impact and implications. *Expert Review of Vaccines* 14, 99–117. <https://doi.org/10.1586/14760584.2015.964212>
- Edmunds, W.J., Gay, N.J., Kretzschmar, M., Pebody, R.G., Wachmann, H., 1994. The pre-vaccination epidemiology of measles, mumps and rubella in Europe: implications for modelling studies 16.
- Fanoy, E.B., Cremer, J., Ferreira, J.A., Dittrich, S., van Lier, A., Hahné, S.J.H., Boot, H.J., van Binnendijk, R.S., 2011. Transmission of mumps virus from mumps-vaccinated individuals to close contacts. *Vaccine* 29, 9551–9556. <https://doi.org/10.1016/j.vaccine.2011.09.100>
- Farrington, C.P., 1991. Subacute sclerosing panencephalitis in England and Wales: Transient effects and risk estimates. *Statistics in Medicine* 10, 1733–1744. <https://doi.org/10.1002/sim.4780101111>
- Feldstein, M., 1999. Tax Avoidance and the Deadweight Loss of the Income Tax. *The Review of Economics and Statistics* 81, 674–680. <https://doi.org/10.1162/003465399558391>
- Fine, P.E.M., Clarkson, J.A., 1986. Individual versus public priorities in the determination of optimal vaccination policies. *American Journal of Epidemiology* 124, 1012–1020. <https://doi.org/10.1093/oxfordjournals.aje.a114471>
- Fischels, J., 2021. Get \$100 For A Vaccine? Cash Incentives Work For Some, Others Not So Much. NPR.
- Fisher, W.A., 2012. Understanding Human Papillomavirus Vaccine Uptake. *Vaccine, Comprehensive Control of HPV Infections and Related Diseases* 30, F149–F156. <https://doi.org/10.1016/j.vaccine.2012.04.107>
- Fombonne, E., Goin-Kochel, R.P., O’Roak, B.J., 2020. Beliefs in vaccine as causes of autism among SPARK cohort caregivers. *Vaccine* 38, 1794–1803. <https://doi.org/10.1016/j.vaccine.2019.12.026>
- Geoffard, P.-Y., Philipson, T., 1997. Disease Eradication: Private versus Public Vaccination. *The American Economic Review* 87, 222–230.
- Gershoff, A.D., Koehler, J.J., 2011. Safety First? The Role of Emotion in Safety Product Betrayal Aversion. *J Consum Res* 38, 140–150. <https://doi.org/10.1086/658883>
- Gershon, A.A., 2015. Mandell, Douglas, and Bennett’s principles and practice of infectious diseases, in: *Measles Virus (Rubeola)*. WB Saunders, pp. 1967–1973.
- Gigerenzer, G., Edwards, A., 2003. Simple tools for understanding risks: from innumeracy to insight. *BMJ* 327, 741–744.

- Gilbey, A., Hill, S., 2012. Confirmation bias in general aviation lost procedures. *Applied Cognitive Psychology* 26, 785–795.
- Goldstein, K.P., Philipson, T.J., Joo, H., Daum, R.S., 1996. The Effect of Epidemic Measles on Immunization Rates. *JAMA* 276, 56–58. <https://doi.org/10.1001/jama.1996.03540010058031>
- Govaert, T.M., Dinant, G.J., Aretz, K., Masurel, N., Sprenger, M.J., Knottnerus, J.A., 1993. Adverse reactions to influenza vaccine in elderly people: randomised double blind placebo controlled trial. *BMJ* 307, 988–990. <https://doi.org/10.1136/bmj.307.6910.988>
- Herceg, A., Passaris, I., Mead, C., 1994. An outbreak of measles in a highly immunised population: immunisation status and vaccine efficacy. *Australian Journal of Public Health* 18, 249–252. <https://doi.org/10.1111/j.1753-6405.1994.tb00239.x>
- Himmelboim, I., Xiao, X., Lee, D.K.L., Wang, M.Y., Borah, P., 2020. A Social Networks Approach to Understanding Vaccine Conversations on Twitter: Network Clusters, Sentiment, and Certainty in HPV Social Networks. *Health Communication* 35, 607–615. <https://doi.org/10.1080/10410236.2019.1573446>
- Horton, R., 2010. Retraction--Ileal-lymphoid-nodular hyperplasia, non-specific colitis, and pervasive developmental disorder in children. *Lancet* 375, 445. [https://doi.org/10.1016/S0140-6736\(10\)60175-4](https://doi.org/10.1016/S0140-6736(10)60175-4)
- Hotez, P.J., Nuzhath, T., Colwell, B., 2020. Combating vaccine hesitancy and other 21st century social determinants in the global fight against measles. *Current Opinion in Virology, Special Section: Measles* 41, 1–7. <https://doi.org/10.1016/j.coviro.2020.01.001>
- Hotz, V.J., Miller, R.A., 1993. Conditional Choice Probabilities and the Estimation of Dynamic Models. *The Review of Economic Studies* 60, 497–529. <https://doi.org/10.2307/2298122>
- Hulsey, E., Bland, T., 2015. Immune overload: Parental attitudes toward combination and single antigen vaccines. *Vaccine* 33, 2546–2550. <https://doi.org/10.1016/j.vaccine.2015.04.020>
- Hviid, A., Rubin, S., Mühlemann, K., 2008. Mumps. *The Lancet* 371, 932–944. [https://doi.org/10.1016/S0140-6736\(08\)60419-5](https://doi.org/10.1016/S0140-6736(08)60419-5)
- Imhoff, R., Lamberty, P., 2020. A Bioweapon or a Hoax? The Link Between Distinct Conspiracy Beliefs About the Coronavirus Disease (COVID-19) Outbreak and Pandemic Behavior. *Social Psychological and Personality Science* 11, 1110–1118. <https://doi.org/10.1177/1948550620934692>
- Irani, D.N., 2008. Aseptic Meningitis and Viral Myelitis. *Neurologic Clinics, Current Issues in Clinical Neurovirology: Pathogenesis, Diagnosis and Treatment* 26, 635–655. <https://doi.org/10.1016/j.ncl.2008.03.003>
- Jacobson, R.M., Sauver, J.L.S., Rutten, L.J.F., 2015. Vaccine Hesitancy. *Mayo Clinic Proceedings* 90, 1562–1568. <https://doi.org/10.1016/j.mayocp.2015.09.006>

- Jacobson, S.H., Sewell, E.C., 2002. Using Monte Carlo Simulation to Determine Combination Vaccine Price Distributions for Childhood Diseases. *Health Care Management Science* 5, 11.
- Jaharuddin, Bakhtiar, T., 2020. Control Policy Mix in Measles Transmission Dynamics Using Vaccination, Therapy, and Treatment. *International Journal of Mathematics and Mathematical Sciences* 2020, e1561569. <https://doi.org/10.1155/2020/1561569>
- Karant, S.S., Lairson, D.R., Savas, L.S., Vernon, S.W., Fernández, M.E., 2017. The cost of developing a computerized tailored interactive multimedia intervention vs. a print based Photonovella intervention for HPV vaccine education. *Evaluation and Program Planning* 63, 1–6. <https://doi.org/10.1016/j.evalprogplan.2017.02.009>
- Kim-Farley, R., Bart, S., Stetler, H., Orenstein, W., Bart, K., Sullivan, K., Haplin, T., Sirotkin, B., 1985. Clinical Mumps Vaccine Efficacy. *American Journal of Epidemiology* 121, 593–597. <https://doi.org/10.1093/oxfordjournals.aje.a114037>
- Kniesner, T.J., Viscusi, W.K., 2019. The Value of a Statistical Life, in: *Oxford Research Encyclopedia of Economics and Finance*. Oxford University Press. <https://doi.org/10.1093/acrefore/9780190625979.013.138>
- Kondo, M., Hoshi, S., Okubo, I., 2009. Does subsidy work? Price elasticity of demand for influenza vaccination among the elderly in Japan. *Health Policy* 91, 269–276. <https://doi.org/10.1016/j.healthpol.2008.12.014>
- Kulpeng, W., Leelahavarong, P., RattanaVIPapong, W., Sornsrivichai, V., Baggett, H.C., Meeyai, A., Punpanich, W., Teerawattananon, Y., 2013. Cost-utility analysis of 10- and 13-valent pneumococcal conjugate vaccines: Protection at what price in the Thai context? *Vaccine* 31, 2839–2847. <https://doi.org/10.1016/j.vaccine.2013.03.047>
- Lee, J.-W., McKibbin, W.J., 2004. Globalization and disease: The case of SARS. *Asian economic papers* 3, 113–131.
- Li, M.Y., Graef, J.R., Wang, L., Karsai, J., 1999. Global dynamics of a SEIR model with varying total population size. *Mathematical Biosciences* 160, 191–213. [https://doi.org/10.1016/S0025-5564\(99\)00030-9](https://doi.org/10.1016/S0025-5564(99)00030-9)
- Lisowski, B., Yuwan, S., Bier, M., 2019. Outbreaks of the measles in the Dutch Bible Belt and in other places – New prospects for a 1000 year old virus. *Biosystems* 177, 16–23. <https://doi.org/10.1016/j.biosystems.2019.01.003>
- London, W.P., Yorke, J.A., 1973. Recurrent outbreaks of measles, chickenpox and mumps: I. Seasonal variation in contact rates. *American Journal of Epidemiology* 98, 453–468.
- Lord, C.G., Ross, L., Lepper, M.R., 1979. Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence. *Journal of Personality and Social Psychology* 37, 2098–2109.

- Ma, J., van den Driessche, P., Willeboordse, F.H., 2013. The importance of contact network topology for the success of vaccination strategies. *Journal of Theoretical Biology* 325, 12–21. <https://doi.org/10.1016/j.jtbi.2013.01.006>
- MacDonald, N.E., Smith, J., Appleton, M., 2012. Risk perception, risk management and safety assessment: What can governments do to increase public confidence in their vaccine system? *Biologicals, Barcelona Symposium - Post-Licensure Evaluation of Vaccine Safety* 40, 384–388. <https://doi.org/10.1016/j.biologicals.2011.08.001>
- Manfredi, P., Posta, P.D., d’Onofrio, A., Salinelli, E., Centrone, F., Meo, C., Poletti, P., 2009. Optimal vaccination choice, vaccination games, and rational exemption: An appraisal. *Vaccine* 28, 98–109. <https://doi.org/10.1016/j.vaccine.2009.09.109>
- Metzger, M.J., Flanagan, A.J., Medders, R.B., 2010. Social and Heuristic Approaches to Credibility Evaluation Online. *Journal of Communication* 60, 413–439. <https://doi.org/10.1111/j.1460-2466.2010.01488.x>
- Milani, E., Weitkamp, E., Webb, P., 2020. The Visual Vaccine Debate on Twitter: A Social Network Analysis. *MaC* 8, 364–375. <https://doi.org/10.17645/mac.v8i2.2847>
- Nyhan, B., Reifler, J., Richey, S., Freed, G.L., 2014. Effective Messages in Vaccine Promotion: A Randomized Trial. *Pediatrics* 133, e835–e842. <https://doi.org/10.1542/peds.2013-2365>
- Omer, S.B., Amin, A.B., Limaye, R.J., 2017. Communicating About Vaccines in a Fact-Resistant World. *JAMA Pediatrics* 171, 929–930. <https://doi.org/10.1001/jamapediatrics.2017.2219>
- Ozawa, S., Clark, S., Portnoy, A., Grewal, S., Stack, M.L., Sinha, A., Mirelman, A., Franklin, H., Friberg, I.K., Tam, Y., Walker, N., Clark, A., Ferrari, M., Suraratdecha, C., Sweet, S., Goldie, S.J., Garske, T., Li, M., Hansen, P.M., Johnson, H.L., Walker, D., 2017. Estimated economic impact of vaccinations in 73 low- and middle-income countries, 2001–2020. *Bull World Health Organ* 95, 629–638. <https://doi.org/10.2471/BLT.16.178475>
- Perry, R.T., Halsey, N.A., 2004. The Clinical Significance of Measles: A Review. *The Journal of Infectious Diseases* 189, S4–S16. <https://doi.org/10.1086/377712>
- Publication on Cost-Benefits (Supplement), 2019. , *Vaccines for Children Program*. Centers for Disease Control and Prevention.
- Razai, M.S., Osama, T., McKechnie, D.G.J., Majeed, A., 2021. Covid-19 vaccine hesitancy among ethnic minority groups. *BMJ* 372, n513. <https://doi.org/10.1136/bmj.n513>
- Richwine, C., Dor, A., Moghtaderi, A., 2019. Do Stricter Immunization Laws Improve Coverage? Evidence from the Repeal of Non-medical Exemptions for School Mandated Vaccines (No. w25847). National Bureau of Economic Research, Cambridge, MA. <https://doi.org/10.3386/w25847>
- Robertson, C., Scheitrum, D., Schaefer, A., Malone, T., McFadden, B.R., Messer, K.D., Ferraro, P.J., 2021. Paying Americans to take the vaccine—would it help or backfire? *Journal of Law and the Biosciences* 8, lsab027. <https://doi.org/10.1093/jlb/lsab027>

- Romer, D., Jamieson, K.H., 2020. Conspiracy theories as barriers to controlling the spread of COVID-19 in the U.S. *Social Science & Medicine* 263, 113356. <https://doi.org/10.1016/j.socscimed.2020.113356>
- Roozenbeek, J., Schneider, C.R., Dryhurst, S., Kerr, J., Freeman, A.L.J., Recchia, G., van der Bles, A.M., van der Linden, S., 2020. Susceptibility to misinformation about COVID-19 around the world. *Royal Society Open Science* 7, 201199. <https://doi.org/10.1098/rsos.201199>
- Sadaf, A., Richards, J.L., Glanz, J., Salmon, D.A., Omer, S.B., 2013. A systematic review of interventions for reducing parental vaccine refusal and vaccine hesitancy. *Vaccine* 31, 4293–4304. <https://doi.org/10.1016/j.vaccine.2013.07.013>
- Salmon, D.A., Dudley, M.Z., Glanz, J.M., Omer, S.B., 2015. Vaccine hesitancy: Causes, consequences, and a call to action. *Vaccine, Advancing Vaccines and Immunization* 33, D66–D71. <https://doi.org/10.1016/j.vaccine.2015.09.035>
- Schoeppe, J., Cheadle, A., Melton, M., Faubion, T., Miller, C., Matthys, J., Hsu, C., 2017. The Immunity Community: A Community Engagement Strategy for Reducing Vaccine Hesitancy. *Health Promotion Practice* 18, 654–661. <https://doi.org/10.1177/1524839917697303>
- Seither, R., Masalovich, S., Knighton, C.L., Mellerson, J., Singleton, J.A., Greby, S.M., 2014. Vaccination Coverage Among Children in Kindergarten — United States, 2013–14 School Year (No. 63(41)), *Morbidity and Mortality Weekly Report*. Centers for Disease Control and Prevention.
- Shiel, A., Jorm, L.R., Carruthers, R., Fitzsimmons, G.J., 1998. Cost-effectiveness of measles outbreak intervention strategies. *Australian and New Zealand Journal of Public Health* 22, 126–132. <https://doi.org/10.1111/j.1467-842X.1998.tb01156.x>
- Siddiqui, M., Salmon, D.A., Omer, S.B., 2013. Epidemiology of vaccine hesitancy in the United States. *vaccines* 9, 2643–2648. <https://doi.org/10.4161/hv.27243>
- Smith, N., Graham, T., 2017. Mapping the anti-vaccination movement on Facebook. *Information, Communication & Society* 22, 1–18. <https://doi.org/10.1080/1369118X.2017.1418406>
- Sobkowicz, P., 2016. Quantitative Agent Based Model of Opinion Dynamics: Polish Elections of 2015. *PLOS ONE* 11, e0155098. <https://doi.org/10.1371/journal.pone.0155098>
- Sobkowicz, P., Sobkowicz, A., 2021. Agent Based Model of Anti-Vaccination Movements: Simulations and Comparison with Empirical Data. *Vaccines* 9, 809.
- Taylor, L.E., Swerdfeger, A.L., Eslick, G.D., 2014. Vaccines are not associated with autism: An evidence-based meta-analysis of case-control and cohort studies. *Vaccine* 32.
- Valdez, A., Stewart, S.L., Tanjasiri, S.P., Levy, V., Garza, A., 2015. Design and efficacy of a multilingual, multicultural HPV vaccine education intervention. *J Commun Healthc* 8, 106–118. <https://doi.org/10.1179/1753807615Y.0000000015>

- Wakefield, A.J., Murch, S.H., Anthony, A., Linnell, J., Casson, D.M., Malik, M., Berelowitz, M., Dhillon, A.P., Thomson, M.A., Harvey, P., Valentine, A., Davies, S.E., Walker-Smith, J.A., 1998. Ileal-lymphoid-nodular hyperplasia, non-specific colitis, and pervasive developmental disorder in children. *Lancet* 351, 637–641. [https://doi.org/10.1016/s0140-6736\(97\)11096-0](https://doi.org/10.1016/s0140-6736(97)11096-0)
- Welsh, W.A., 1991. The Economic Impact of Deafness 24, 10.
- White, C.C., Koplan, J.P., Orenstein, W.A., 1985. Benefits, risks and costs of immunization for measles, mumps and rubella. *Am J Public Health* 75, 739–744. <https://doi.org/10.2105/AJPH.75.7.739>
- Wojcieszak, M., Sobkowicz, P., Yu, X., Bulat, B., 2021. What Information Drives Political Polarization? Comparing the Effects of In-group Praise, Out-group Derogation, and Evidence-based Communications on Polarization. *The International Journal of Press/Politics* 19401612211004416. <https://doi.org/10.1177/19401612211004418>
- Woudenberg, T., van Binnendijk, R.S., Sanders, E.A.M., Wallinga, J., de Melker, H.E., Ruijs, W.L.M., Hahné, S.J.M., 2017. Large measles epidemic in the Netherlands, May 2013 to March 2014: changing epidemiology. *Euro Surveill* 22, 30443. <https://doi.org/10.2807/1560-7917.ES.2017.22.3.30443>
- Wu, Z.-X., Zhang, H.-F., 2013. Peer pressure is a double-edged sword in vaccination dynamics. *EPL (Europhysics Letters)* 104, 10002.
- Yang, Y.T., Shaw, J., 2018. Sudden infant death syndrome, attention-deficit/hyperactivity disorder and vaccines: Longitudinal population analyses. *Vaccine* 36, 595–598. <https://doi.org/10.1016/j.vaccine.2017.12.065>
- Zucker, J.R., Rosen, J.B., Iwamoto, M., Arciuolo, R.J., Langdon-Embry, M., Vora, N.M., Rakeman, J.L., Isaac, B.M., Jean, A., Asfaw, M., Hawkins, S.C., Merrill, T.G., Kennelly, M.O., Maldin Morgenthau, B., Daskalakis, D.C., Barbot, O., 2020. Consequences of Undervaccination — Measles Outbreak, New York City, 2018–2019. *New England Journal of Medicine* 382, 1009–1017. <https://doi.org/10.1056/NEJMoa1912514>

APPENDICES

Appendix A: Modeling Specification

This appendix details the model I constructed to explain the low rate of vaccination consistently observed. It does this by modeling agent behavior in the midst of a measles outbreak,³² and presents the findings. It contains a more detailed description than included in the main body, presented for the purpose of greater understanding and predictability.

I use a computer simulation to model child-parent pairs (agents). The children experience all health outcomes and is subject to disease spread, while the parent collects data, makes decisions on behalf of the child, and maintains a social network. These agents seek to minimize their expected lifetime cost by regularly deciding whether or not to vaccinate. They do this by comparing their perceived cost of vaccinating against the cost of not vaccinating. The cost of vaccinating is both in terms of vaccine price and any real or imagined adverse reactions to the vaccine, calculated as an expected cost. The cost of not vaccinating depends on how likely an agent is to get sick, as well as the cost of the various health outcomes of becoming sick with the measles. Exactly how heavily agents weight these outcomes is heterogeneous among agents to account for the substantial observed variation in vaccination beliefs. Agents start unvaccinated and healthy. An agent who becomes infected proceeds undergoes a latency period where they experience no effect. They then experience an asymptomatic but infective period, in which they are infecting other agents in their spatial area. They then begin to show symptoms, and

32 The model is set in an outbreak for two reasons. First, vaccination rates change most quickly during outbreaks, and tend to change slowly otherwise. This provides useful variation in data. Second, an outbreak is a time when it would make the most sense to vaccinate; with the disease on one's doorstep, it would seem rational to vaccinate. That vaccination rates do not spike to 100% during outbreaks seems to indicate an alternative model is required; this model attempts to fill that gap, and so is presented in its context. Furthermore, while outbreaks are rare, they aren't terribly rare; they appear to happen at a rate of roughly one every two years.

continue infecting. Following this, they retain symptoms but cease infecting others, then finally recover and are immune to further infection. The length of these periods is heterogeneous among agents.

Each scenario (i.e., set of parameters) is modeled at 1 day time steps for 150 days³³. Agent position is spatially implicit to account for two factors: first, the disease only spreads within a certain range. This means the natural clustering of individuals simulates the hot-beds of measles outbreaks which I observe in the world. Second, it provides a basis for individuals to make their vaccination decisions in accordance with the model.

The remainder of this chapter is organized as follows: Section A.2 Theoretical Model (p. 75) details the general framework employed. Section A.3 Vaccination Decision (p. 79) examines how agents make a vaccination decision, while Section A.4 Disease Spread (p. 81) describes how the disease spreads within and among agents. Section A.5 Social Networks (p. 82) describes how social networks impact vaccination decisions. All simulations are performed in Anylogic 8 University³⁴.

A.2 Theoretical Model

This subsection outlines the model that describes vaccination behavior. Within the model, each agent represents a child-parent pair. The child is subjected to disease spread, vaccination, and spreads disease, while the parent gathers information and decides whether or not to get vaccinated each time period.

Because the parent is assumed to fully internalize the child's utility, this pair is treated as a single agent. The agent seeks to maximize utility by minimizing expected costs via choosing whether or not to vaccinate.

33 This number is chosen because, after 100 days, the simulation has reached equilibrium, after which no changes are observed. 50 days are added for safety.

34 Version 8.7.12, Build Build: 8.7.12.202205250455 x64

A.2.1 Rational Behavior

Agents face costs during an outbreak. These costs come from the costs associated with measles, as well as costs associated with the vaccine. They do this by, in each period, comparing that period's expected cost in a vaccinated state, $E[C_{VAC,t}]$ against the expected costs if they are unvaccinated, $E[C_{UNV,t}]$. If being vaccinated offers the lower expected cost, the agent chooses to vaccinate. Each agent also possesses a heterogeneous error term ϵ_i , representing the different attitudes towards vaccination³⁵. Thus, the rational decision rule is formally: *Vaccinate* if $E[C_{VAC,t}] - E[C_{UNV,t}] > \epsilon_i$.

The costs of measles depends on one's vaccination status, and the probability of contracting measles depends both on the vaccination status and on how prevalent measles is at that time. The costs associated are direct treatment costs T , disability costs D , and mortality costs M . Treatment includes cost of being hospitalized C_{HOS} , and the cost of outpatient visits, C_{MD} , with their corresponding probabilities. Thus $T = p_{HOS} C_{HOS} + p_{MD} C_{MD}$. For details on the values of these parameters, see Appendix B: Parameterization (p. 85). Similarly, $M = p_{deaf} C_{deaf} + p_{NS} C_{NS}$, where C_{deaf} is the lifetime cost of deafness as measured by lost wages, and C_{NS} is the lifetime cost of Neurological Sequelae. Finally, D includes both the cost of dying directly from the measles infection C_{DEATH} as measured by the Value of Statistical Life, as well as the cost of delayed death due to Subacute sclerosing panencephalitis (SSPE), noted as C_{SSPE} . Thus, we can write $D = p_{DEATH} C_{DEATH} + p_{SSPE} C_{SSPE}$.

As mentioned above, these costs also depend on an agent's vaccination state. A vaccinated patient tends to experience symptoms less severe than an unvaccinated patient. Thus, we can write that for a vaccinated agent, the costs they expect to incur given a sickness is a scalar multiple b of the costs an unvaccinated individual faces ($0 < b < 1$). In mathematical terms, $(T_{VAC} + D_{VAC} + M_{VAC}) = (b T_{UNV} + b D_{UNV} + b M_{UNV}) = b \times (T_{UNV} + D_{UNV} + M_{UNV})$.

³⁵ This term is assumed to be $\sim N(0, \sigma_\epsilon)$.

Agents gauge how likely they are to contract measles $p(s_t)$ based on the portion of the population currently displaying symptoms at that time, denoted by s_t .³⁶ Being vaccinated also reduces the probability that an agent will contract measles. Being vaccinated generates immunity at a predictable rate of $(1 - k)$. Thus, $p_{VAC}(s_t) = kp_{UNV}(s_t)$.

$$E[C_{UNV,t}] = p_{UNV}(s_t) \times (T_{UNV} + D_{UNV} + M_{UNV})$$

$$E[C_{VAC,t}] = kb \times p_{UNV}(s_t) \times (T_{UNV} + D_{UNV} + M_{UNV})$$

Equation 3: Expected Costs of Vaccinated and Unvaccinated States

Agents choosing to vaccinate must also pay the price of the vaccine, P_V , and face a rare and uniformly nonfatal chance of a Vaccine Adverse Reaction (VAR), as denoted by A_V . Such reactions require hospitalization, but have no lasting effects. Thus, $A_V = V_{AR} C_{HOS}$, where V_{AR} is the rate at which such reactions occur. Again I stress that this does not include the debunked links to Autism. An agent's expected costs in both states for a given time period t is therefore:

The rational decision rule can, therefore, be written as *vaccinate* if $p_{UNV}(s_t) \times (T_{UNV} + D_{UNV} + M_{UNV}) \times (1 - kb) - P_V - A_V > \epsilon_i$. Note that, because the agent makes this decision each time period, they do not apply any discounting. This is one of two areas where an agent's projection of costs may be inaccurate; because the vaccine takes two weeks to induce immunity, the actual costs they face may be higher even if vaccinated. However, this is counterbalanced by the fact that agents are using symptomatic incidence of the present s_t to forecast future spread tends to underestimate the resulting spread, and thus cost of being unvaccinated.

36 For the purposes of this paper, let $p(s_t) = \max\{1, I_{MULT} \times s_t, s_{BG}\}$ where s_{BG} is the background level of risk, and I_{MULT} is a scalar multiplier to incidence. It is through I_{MULT} that any agent expectations of future spread are captured.

A.2.2 Heuristics and Misinformation

Agents using a heuristic use a modified decision rule. Instead of vaccinating based on cost, they instead use a simple threshold. Heuristic users will *vaccinate* if $s_t > \bar{s}$.

Misinformation is similarly simple to include. Each agent has a heterogeneous belief that vaccines can cause ASD with a probability of w . For 16% of the population, $w \sim N(\mu_{ASD}, \sigma_{ASD})$. For the remaining 84% of the population, $w=0$. The mean and standard deviation are assumed, and subjected to sensitivity analyses, as shown in Appendix D: Sensitivity Analyses (p. 97).

A.2.3 Disease Spread

The disease itself progresses normally as shown in Figure 21: Transition Matrix. (p. 4) Individuals begin susceptible, and at a rate of r per infected individual per day, are infected by sick others. Once infected, they are latent but noninfective for a period L_N days. Once this passes, as shown by the use of the heavyside step function $\theta(\cdot)$ they immediately progress to being asymptomatic but begin infecting others. This latent infective period lasts L_I days, at which point they begin showing symptoms and keep infecting others. This symptomatic and infective period lasts S_I days, at which point they still show symptoms but cease infecting others. This lasts for S_N days, at which point they are recovered and immune to further infection.

The length of these periods can be different from agent-to-agent; the lengths of each period are further discussed in section Figure 21: Transition Matrix (p. 79). A detailed description of all parameters, and their origins, can be found in Appendix B: Parameterization (p. 85).

$$\begin{bmatrix} 1-r & r & 0 & 0 & 0 & 0 \\ 0 & 1-\theta(t-L_N) & \theta(t-L_N) & 0 & 0 & 0 \\ 0 & 0 & 1-\theta(t-L_N-L_I) & t-L_N-L_I & 0 & 0 \\ 0 & 0 & 0 & 1-\theta(t-L_N-L_I-S_I) & \theta(t-L_N-L_I-S_I) & 0 \\ 0 & 0 & 0 & 0 & 1-\theta(t-L_N-L_I-S_I-S_N) & \theta(t-L_N-L_I-S_I-S_N) \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

Figure 21: Transition Matrix

These periods are heterogeneous across agents.

A.3 Vaccination Decision

As described above, agents decide whether or not to vaccinate according to their decision rule.

Heuristic users employ a simple threshold, while rational users calculate costs and pick the lowest one.

When determining costs, agents include only costs they themselves will face individually; they do not consider any positive externality they may generate by being vaccinated.

I describe the rational decision first, and describe subsequent scenarios as variants of the rational baseline.

A.3.1 Rational Vaccination Decision

Rational agents generate two costs: The costs in a vaccinated state, $E[C_{VAC}]$, and the costs in an unvaccinated state, $E[C_{UNV}]$. $E[C_{UNV}]$ is calculated as outlined in A.2.1 Rational Behavior (p. 76), using the parameters as described in Appendix B: Parameterization (p. 85), which are given for unvaccinated cases. This also generates $E[C_{VAC}]$.

In order to determine the probability of being infected, they select the higher of the background risk s_{BG} , and the global symptomatic rate in that time period, s_t . To generate the global symptomatic rate, an agent examines all agents in the world. They find the percent of the population that are currently

showing symptoms (either infective symptomatic or turned-the-corner). They then multiply this by a scalar value, I_{MULT} to account for the fact that disease spreads in an exponential, not linear, manner.

$$p_{UNV}(s_t) \times (T_{UNV} + D_{UNV} + M_{UNV}) \times (1 - kb) - P_V - A_V > \epsilon_i$$

Equation 4: Rational Agent Decision Rule

Once both $p_{UNV}(s_t)$ and $E[C_{UNV}]$ are found, the decision rule is simple. If Equation 4: Rational Agent Decision Rule (p. 80) evaluates as true, the agent immediately vaccinates.

In this model, all costs of vaccination occur upon the instant of vaccination.

A.3.2 Heuristic Vaccination Decision

Heuristic users instead use a simple threshold. As outlined in the previous section, heuristic users will *vaccinate* if $s_t > \bar{s}$.

A.3.3 Misinformed Vaccination Decision

Agents incorporating misinformation regarding ASD have an extra cost included in their analysis. For these agents, $A_V^* = A_V + y_i \times C_{ASD}$, where A_V^* is the agent's adjusted expected cost of a vaccine-adverse reaction, y is the agent heterogeneous believed probability of vaccine-induced ASD, and C_{ASD} is the true estimated cost of ASD.

To determine each agent's w , first agents are randomly assigned a misinformed status; each agent has a 16% chance to be misinformed. If an agent is not misinformed, w is set to 0, and thus the agent behaves rationally as in A.3.1 Rational Vaccination Decision (p. 79). Otherwise, the agent receives a random draw from a normal distribution and assign it to their belief, such that $y = N(\mu_{ASD}, \sigma_{ASD})$.

When calculating their expected costs as in Equation 4: Rational Agent Decision Rule, they use A_v^* for their expected costs from VARs. Note that the true probability of w is 0, so such costs are never actually incurred, though misinformed agents behave as if their belief were true.

A.4 Disease Spread

Agents start susceptible to infection. If infected, they go through a latent period, then begin infecting others while they themselves are asymptomatic. Then they begin to show symptoms, and after a period, cease infecting others. After a time, if they do not perish from the illness, they recover. Recovered individuals are presumed to be immune.

The length of these periods is heterogeneous to each agent. The length of these periods is drawn from triangular distributions, the parameters of which can be found in Appendix B: Parameterization, Subsection B.1 Measles (p. 85).

Having established the model that will be employed, now we must codify it. The model will run at 1 day time steps for 150 days³⁷. Agent position is spatially implicit to account for two factors: first, the disease only spreads within a certain range. This means the natural clustering of individuals simulates the hot-beds of measles outbreaks which we observe in the world. Second, it provides a basis for individuals to make their vaccination decisions in accordance with the above model.

Costs for being sick accrue the instant an agent enters the symptomatic state for the purposes of this model.

Agents continue to act as described in Section A.2 Theoretical Model (p. 75), picking the lowest cost of vaccinating or not vaccinating. The state transitions can be described by Figure 3: The Susceptible-

³⁷ This number is chosen because, after 100 days, the simulation has reached equilibrium, after which no changes are observed. 50 days are added for safety.

Exposed-Infected-Recovered (SEIR) Model (p. 20Error: Reference source not found). The Health State Chart describes how the agent's health changes over time.

A.5 Social Networks

Social networks have two components: neighbors, which represent any other agents an individual will have physical contact with, and thus can infect or be infected by; and friends, which represent individuals an agent keeps in contact with through nonphysical means, like social media. Both these groups are included in an agent's social network.

Social networks perform two functions. First, they can provide more relevant data on how a disease is spreading close to an agent, which I term "local information." Second, they provide a way for agents to modify their own belief that ASD is caused by vaccines (here termed "ASD belief"). If an agent is in a group with below average ASD belief, the agent's ASD belief will also go down. If, however, an agent is in a group with above average ASD belief, that individual's ASD belief goes increases. Such an increase is assumed to be asymmetrically larger than a decrease, to account for confirmation bias.

A.5.1 Social Network Specification

Neighbors within a social network are defined in-model as any individual within the distance across which disease can spread. That is, every neighbor is an individual who can have contact with and infect an agent, and every such individual is a neighbor.

Friends follow a scale-free network per Ma et al. (2013) and are randomly selected from the entire population. These agents can also be neighbors, but can come from anywhere in the simulation.

Note that if an agent is in both groups (i.e., both a neighbor and a friend) then they are counted twice in subsequent calculations. This is to reflect that known and proximate sources of information are taken to be more credible.

A.5.2 Disease Incidence in Social Networks

While social networks are active, there are two sources of information for current incidence. There remains the global incidence rate, which I denote $s_{G,t}$ and a local incidence rate, $s_{L,t}$, defined as the average symptomatic rate for all individuals within an agent's social network. An agent's adjusted estimate of risk accounts for both these factors, but may weigh one more heavily. They also will continue to use the background risk, if it is higher. As such, an agent's adjusted assessment of risk can be seen in Equation 5: Weighted Incidence in a Social Network (p. 83), where w is the relative weight an agent puts on local incidence, and s_t^* is an agent's adjusted assessment of risk.

$$s_t^* = \max \left\{ s_{BG}, I_{MULT} \cdot \frac{s_{G,t} + w s_{L,t}}{1 + w} \right\}$$

*Equation 5: Weighted Incidence
in a Social Network*

A.5.3 Misinformation in Social Networks

Agents in social network have their beliefs moderated by those in their networks. To do this, each agent determines the average level of ASD belief within their network (Y_i) including all agents who believe in no link. They then find how much their network's belief differs from the network average (\bar{y}) defining the difference as $y_{dif} = (\bar{y} - Y_i)/Y_i$. Finally, their own belief is modified proportional to this difference. However, the degree to which it changes depends on whether they have a more well-informed ($Y_i < \bar{y}$)

or more misinformed ($Y_i > \bar{w}$) network. For a well informed network, the adjusted ASD belief $y^* = (1 - y_{dif}) \times y$. For a misinformed network, individuals who already have a nonzero ASD belief will be more likely to take their peer's agreement as evidence of legitimacy. Thus, $y^* = (1 + y_{dif} \xi) \times w$, where ξ is a confirmation bias factor. Note that for well-informed agents, w^* remains zero.

When determining the costs in a vaccinated state, they use $A_v^* = A_v + y^* \times C_{ASD}$.

A.6. Spatial Distribution

The model is spatially implicit. Agents are distributed at random in 1000×1000 grid, with their location drawing from a uniform distribution. Agents can spread disease within a circular area centered around that agent. Neighbors are defined similarly. Agent positions are static.

A.7. Stochasticity

There are few sources of stochasticity in the model. All agent characteristics, such as location, agent-heterogeneous parameters, ASD belief, etc. are the same from simulation run to simulation run. Further, unless a parameter is changed for the purposes of analysis, the parameters are also unchanged across different scenarios. For instance, an agent will have the same ϵ_i term in the benchmark case as in the misinformation case. The agent will have a different ϵ_i only in the sensitivity analysis that changes the distribution from which ϵ_i is drawn.

The source of stochasticity comes from how the disease spreads. Because all parameters are the same across simulation runs, including the location of the initial infection, the only thing that changes is how the disease spreads. This comes from the fact that, each day, infected agents randomly select another agent to expose. It is that random selection that drives variance within simulation runs.

Appendix B: Parameterization

The parameters for the model are shown in Table 7: Measles Progression Parameters and Sources.

(p86) The exact process for determining these parameters, and of how they are use, are discussed in each subsection. Of particular note, however, are the parameters estimated by calibration. Although units (meters) are included, to maintain consistency with the model environment, these distances are abstractions, and not intend to be in any way representative of a physical fact.

B.1 Measles

Once infected, an individual begins to infect others. The mean daily infectivity of measles in 1973 was roughly 1 (London and Yorke, 1973) which, while it displays seasonal variation, is assumed to be constant for the purposes of this paper, though other potential data sources (Edmunds et al., 1994) remain. Each individual is able to expose any other individual within a certain radius to infection; that distance is determined by calibration. (See Appendix C: Calibration and Verification, p. 93) At a rate of once per day (ibid), an infected individual selects one of their neighbors, who may be of any infection and vaccination state, and exposes them to the virus.

This occurs only in the infective period. The disease begins with the host being exposed. 8-12 days later, they develop symptoms, which intensify and include a rash in 2-4 days (Perry and Halsey, 2004). 3-10 days after rash onset, the patient is recovered. The host is infectious starting 3-5 days before symptom onset, and ceases to be infectious up to 4 days after the appearance of a rash. (ibid).

Once symptomatic, some individuals will require a doctor's visit (Carabin et al., 2002). Among vaccinated individuals, this is 20%, while among unvaccinated, it is 50% (White et al., 1985). 1.5% of

them will require hospitalization, which will then require two follow-up visits. (ibid.) The costs for these have very wide error bars; hospitalization is estimated to cost \$4,032-46,060, while doctor visits are estimated to cost \$88-526 (Seither et al., 2014).

Table 7: Measles Progression Parameters and Sources

Parameter	Symbol	Value	Unit	Source
Daily infectivity	r	1	Infections per day	London & Yorke, 1973
Minimum latent noninfective period	$L_{N,MIN}$	3	days	Biellik & Clements, 1997; Perry & Halsey, 2004
Median latent noninfective period	$L_{N,MED}$	6	days	Biellik & Clements, 1997; Perry & Halsey, 2004
Maximum latent noninfective period	$L_{N,MAX}$	9	days	Biellik & Clements, 1997; Perry & Halsey, 2004
Minimum latent infective period	$L_{I,MIN}$	3	days	Biellik & Clements, 1997; Perry & Halsey, 2004
Median latent infective period	$L_{I,MED}$	4	days	Biellik & Clements, 1997; Perry & Halsey, 2004
Maximum latent infective period	$L_{I,MAX}$	5	days	Biellik & Clements, 1997; Perry & Halsey, 2004
Symptomatic infective period	S_I	4	days	Biellik & Clements, 1997; Perry & Halsey, 2004
Minimum symptomatic noninfective period	$S_{N,MIN}$	0	days	Biellik & Clements, 1997; Perry & Halsey, 2004
Median symptomatic noninfective period	$S_{N,MED}$	3	days	Biellik & Clements, 1997; Perry & Halsey, 2004
Maximum symptomatic noninfective period	$S_{N,MAX}$	6	days	Biellik & Clements, 1997; Perry & Halsey, 2004

While many cases of measles will be straightforward, 40% of all cases will have at least one complication (Bester, 2016). These include pneumonia (5%), diarrhea (8%), otitis media (7%), encephalitis (1-3%), panencephalitis (1%), subacute sclerosing panencephalitis (0.01%) and death (0.2%; ibid). All of these are significant enough to require hospitalization. Any fatalities from these complications is subsumed into the general 0.2% fatality rate. Encephalitis is carries with it special complications.

Roughly 25% of encephalitis cases result in neurologic damage (ibid). This can have a cost beyond the hospital care, which is estimated at \$877,150 in 1998 USD (Shiel et al., 1998). This amounts to roughly

1.4 M in 2021USD. Also worthy of elaboration is subacute sclerosing panencephalitis (SSPE) which is a fatal, long-term condition which can develop years after the initial infection. Its symptoms include personality change, cognitive decline, and motor dysfunction, culminating in death. This also is significant enough to warrant separate consideration.

Because SSPE is fatal, it ends in death. The average period from measles infection to developing SSPE is 6.84 years, and the period from diagnosis to death is 2.5 years (Farrington, 1991). Thus, a total average time between measles onset and death from SSPE is 9.34 years. Applying the discount factor (discussed in subsection B.4 Vaccination Decision, p. 91) to the value of statistical life (below) we can estimate the cost of SSPE to be \$528,000. Hospitalization costs for SSPE were not considered.³⁸

The cost of death is the value of statistical life, as taken from Kniesner & Viscusi, (2019). For this estimate, the least biased estimate reported by the authors of \$10 M is used.³⁹ While not a precise measure, its historical use in judging the merits of public programs makes it appropriate to use in this case.

Table 8: Measles Pathology Parameters and Sources

Parameter	Symbol	Value	Unit	Source
Infection spread distance	d_i	340	meters	Calibration
Encephalitis incidence	p_{enceph}	0.02	%	Buchanan and Bonthius 2012
Neurologic sequelae incidence	p_{NE}	25	%	Buchanan and Bonthius 2012
SSPE incidence	p_{SSPE}	0.01	%	White, Koplan, and Orenstein 1985
Hospitalization Incidence	p_{HOS}	1.5 ⁴⁰	%	White, Koplan, and Orenstein 1985
Outpatient visit incidence	q_{MD}	50	%	Carabin et al., 2003; White et al., 1985
Deafness incidence	p_{DEAF}	1.75	%	Cohen et al., 2014
Fatality rate	p_{DEATH}	0.2	%	Bester, 2016

38 A lack of reliable data on the hospitalization rate of SSPE patients made this value impossible to reliably calculate. Furthermore, as the cost of death is much higher than the cost of hospitalization, the cost of death is assumed to dominate.

39 This VSL is selected for three reasons. First, it uses revealed preferences by examining wage differentials, eliminating the possibility of stated preference bias. Second, it employs the Census of Fatal Occupational Injuries dataset, which removes much of the selection and publication bias which is found in other datasets. Finally, as a matter of practicality, the value lies firmly on the mean of other revealed-preference VSL estimates.

40 This probability greater than 1 reflects that most patients require more than one outpatient visit.

Two final outcomes of note are permanent deafness, which occurs at a rate between 0.1% and 3.4% of infections (Cohen et al., 2014). A mean value of 1.75% is used in the model. Lifetime cost of deafness is estimated to be between \$356,000 and \$609,000 (Welsh, 1991). A mean value of \$475,000 is used.

Because the exact health outcomes do not impact any model agents other than the patient, rather than modeling each outcome, when considering damages a simple expectation will be taken. While it is not the most realistic to have each mumps case pay for 5% of a meningitis infection (the expected costs from meningitis being the product of their price tag and their probability) it does greatly simplify the model and, as noted, the exact health incomes should have no impact on the spread.

Costs for these outcomes are calculated upon entering the noninfective symptomatic state and added to the total at that time.

Putting it all together, we can see that the expected cost of becoming sick with measles when unvaccinated is given by Equation 6: Evaluated Cost of Measles (p. 88), which evaluates to roughly $\$3.63 \times 10^5$.

$$\begin{aligned}
 E[C_m] &= C_H p(H) + C_{MD} p(MD) + C_{deaf} p(deaf) + C_{NS} p(NS) + C_{SSPE} p(SSPE) + C_{death} p(death) \\
 &= \$307 \cdot 0.5 + \$25,046 \cdot 0.015 + \$4.75 \cdot 10^5 \cdot 0.00005 + \$1.4 \cdot 10^6 \cdot 0.25 + \$5.28 \cdot 10^5 \cdot 0.0001 + \$1 \cdot 10^7 \cdot 0.00125 \\
 &= \$363,105.74
 \end{aligned}$$

Equation 6: Evaluated Cost of Measles

B.2 Vaccination

Preexposure vaccination is considered to have two benefits: First, it renders immune between 80% and 97% of those who are so vaccinated (*Publication on Cost-Benefits*, 2019). A value of 87% is used as a baseline (Herceg et al., 1994; Kim-Farley et al., 1985). Second, those who do become ill face a milder set of symptoms relative to those who were not vaccinated. (ibid.) For the purposes of this paper, two weeks are required between vaccination and full efficacy. (ibid.) If an individual is infected before that

allotted time has passed, they are presumed to become ill, but with the milder symptoms associated with having been prevaccinated.

Vaccination has been shown to reduce both the oral virus concentration and, presumably, infectivity by roughly 50% (Fanoy et al., 2011). While no quantitative figures were found on outcomes, it will be assumed that the probability of any negative outcome is reduced by a commensurate amount. This may be overstating the impact of mumps on vaccinated individuals; the virus often has to build to a critical level to cause harm, so a 50% reduction in viral population may result in a much larger reduction in negative health outcomes. However, due to the absence of harder figures, this neophyte will assume a linear relationship.

While a vaccine can also be administered therapeutically after exposure to prevent the case from experiencing clinical symptoms (Gershon, 2015), this possibility is dismissed for the purposes of this paper in the interest of expediency, and because the incubation period is sufficiently long that relatively few people will be made aware of their exposure within 3 days.

The value for background risk is chosen to produce an initial vaccination rate consistent with vaccine-hesitant communities. See Appendix C: Calibration and Verification, Section C.1.2 Background Risk (p. 94) for details. The value for the Vaccine Incidence Multiplier is chosen to account for how laypeople tend to estimate exponential values. Christandl (2008) shows that most individuals use a linear approximation plus a correction factor when performing such tasks. A value of 6 was chosen as an approximate representation of a reasonable projection for rapidly growing exponential function.⁴¹

The heuristic vaccination threshold is assumed, as no such data was found.

⁴¹ Respondents were asked to estimate the result of 25 years of exponential growth at a rate of 5%. The highest estimates were a 650% increase. Because measles spreads much faster than 5%, and the time frame for interactivity is in the neighborhood of 25 periods, a high value of 600% was selected.

Table 9: Vaccine Parameters and Sources

Parameter	Symbol	Value	Unit	Source
Vaccine-induced immunity rate	k	87	%	Herceg et al., 1994; Kim-Farley et al., 1985
Vaccine-induced symptom reduction	b	50	%	Fanoy et al., 2011
Vaccine Adverse Reaction Probability	V_{AR}	0.01	%	Hviid et al., 2008; Irani, 2008
Vaccination decision search radius	d_{infec}	314.5	meters	Calibration
Vaccination medial exemption rate	V_{EX}	0.2	%	Seither et al., 2014
Background Risk	S_{BG}	0	%	Calibration
Vaccine Incidence Multiplier	I_{MULT}	6	unitless	Christandl, 2008
Agent Error Standard Deviation	σ_E		\$	Calibration
Heuristic Vaccination Threshold	\bar{s}	7.5	%	Assumption
Mean ASD Expectation	μ_{ASD}	0.25	%	Assumption
ASD Standard Deviation	σ_{ASD}	0.1	%	Assumption
Confirmation Bias Factor	ξ	5	unitless	Assumption

The decision of whether or when agents choose to vaccinate is discussed in another section. The price for the vaccine (\$78.678) is taken from the current *CDC Vaccine Price List* (2021).

The characteristics for ASD Misinformation are selected in order to get the misinformed agents close to the 70% vaccination level described by Goldstein et al. (1996). They are subjected to sensitivity analyses in Appendix D: Sensitivity Analyses (p. 97).

B.3 Vaccination Adverse Reactions

Vaccination is presumed to have a low risk of complication, between 1 in 6,000 and in in 62,000 (median 1 in 11,000), as per Hviid et al. (2008) The main medically relevant complication is aseptic meningitis. This is assumed to occur over average of 15 days (Irani, 2008). At the median rate, the likelihood of any one patient developing an adverse reaction in the 15 day period is 9.09×10^{-5} . As shown in Equation 7: Daily probability of adverse vaccine reaction (p. 91) the daily probability of developing an adverse reaction is roughly 6×10^{-6} .

$$(1-p)^{15} = 1 - 9.09 \times 10^{-5} \rightarrow 1-p = (1 - 9.09 \times 10^{-5})^{\frac{1}{15}} \approx 1 - 6 \times 10^{-6}$$
$$p \approx 6 \times 10^{-6}$$

Equation 7: Daily probability of adverse vaccine reaction

The cost of a vaccine adverse reaction is given by the cost of hospitalization; it is uniformly nonfatal and leaves no lasting effects, so $E[C_{\text{VAR}}] = \$25,046$.

B.4 Vaccination Decision

The vaccination decision follows the process described in Chapter 2. The values for how big the “local area” is are obtained from calibration.

The constant for Betrayal Aversion comes from Aimone & Houser (2013). In this paper, participants play a variant of the Prisoner’s Dilemma against a computer for money, where they can either split a lesser amount of money, each getting \$5 with certainty, or else chose to trust the other player. If the other player similarly trusts, they each get \$15. If the other player betrays, the betrayer gets \$28 while the trusting player gets \$2.

They find that, when facing a computer that assigns betrayal at the same rate as human participants, 92% of participants extend trust, compared to 65% when they will be aware if they've been betrayed. This suggests that the value of a loss from betrayal is worth 7.5% more than a loss from random chance⁴². In other words, a \$1 loss from a betrayal is valued at \$1.075.

Table 10: Cost Parameters and Sources

Parameter	Symbol	Value	Unit	Source
Vaccine price	P_V	78.68	\$US	<i>CDC Vaccine Price List, 2021</i>
Hospitalization cost	C_{HOS}	25046	\$US	Seither et al., 2014
Outpatient visit cost	C_{MD}	307	\$US	Seither et al., 2014
Lifetime cost of Neurological Sequelae (NPV)	C_{NS}	1400000	\$US	Shiel et al., 1998
Lifetime cost of SSPE (NPV)	C_{SSPE}	528000	\$US	Farrington, 1991; Hotz & Miller, 1993
Lifetime cost of deafness (NPV)	C_{deaf}	475000	\$US	Welsh, 1991
Cost of death (NPV)	C_{DEATH}	10000000	\$US	Kniesner & Viscusi, 2019
Betrayal Aversion Scalar	B_{AV}	1.07	scalar	Aimone & Houser, 2012
Per-person education cost	C_{EDUC}	21	\$US	Karanth et al., 2017
Vaccine Ambassador Impact	VA_{imp}	0.35	%	Calibration

42 The expected utility of knowing one has been betrayed is: $EU(KNOW) = 0.346 \cdot \$5 + 0.654 \cdot [0.692 \cdot \$2 + 0.308 \cdot \$15] = \5.66 . The expected utility from not knowing is: $EU(DON'T KNOW) = 0.080 \cdot \$5 + 0.920 \cdot [0.679 \cdot \$2 + 0.321 \cdot \$15] = \6.06 . Thus, $[EU(DON'T KNOW) - EU(KNOW)]/EU(DON'T KNOW) = 0.0695$. So a \$1 gain in a random experiment is worth $1 - 0.0695 = \$0.9305$ in an experiment with betrayal. Thus, a \$1 loss from non-betrayal is worth $\$1/0.9305 = \1.07 in the presence of betrayal.

Special thanks to Aziz Saglam for his help formulating this.

Appendix C: Calibration and Verification

C.1 Calibration

C.1.1 Infection Spread Distance

For calibration, we need to determine the appropriate the radius within which measles can spread. In order to be sure we are matching the spread correctly, I select from literature data where both the mean incidence rate was known and where the vaccination rate was known. For this, the 2013-2014 measles outbreak in the Netherlands is taken to be the model.

In this outbreak, there were two stages of infection. First, the initial infection began on May 27, 2013 with two cases, and peaked by July 14, 2013 with 180 reported cases. On June 17, a national outbreak management team (OMT) began acting to stop the spread of Measles. That marks the end of the initial infection. Subsequently, parents were allowed and encouraged to vaccinate their children.

The precise vaccination rate for the originating community was not reported, but a range of 60-90% vaccine coverage was stated by the authors. A middle value of 75% is assumed. Furthermore, Woudenberg et al. (2017) notes that only 7-9% of Measles cases are reported. Correcting for this, we would expect 2,250 individuals to have become infected after 19 days. This is out of a population of 25,000 as reported by Lisowski et al. (2019). This works out to an attack rate of 9%.

Having established that we expect a 9% attack rate at 75% vaccination at 19 days after the first symptoms, I ran an optimization in Anylogic to minimize the squared difference between my simulation and the actual results. To do this I count the total number of infected and recovered individuals at time=29 in my model.⁴³

⁴³ This time was chosen to allow individuals to develop symptoms in the model. With a median of 10 days between being infected and developing symptoms, I expect the first agents to show symptoms at $t=10$, allowing the 19 days after the

The results of the calibration, however, have experienced substantial variation. Out of 20 calibration attempts, the Sum of Squared Errors was equal to 0.007 or 0.008. For each attempt, a median of the lowest SSE was computed. These values ranged from 102 m (minimum) to 425 m (maximum), with a median of 297.5 m, and standard error of 111 m.

This implies that the model is insensitive to infection radius. This is explored in greater detail in Appendix D: Sensitivity Analyses (p. 97) but for the baseline model the median result (297.5 m) is used.

C.1.2 Background Risk

Similar to the above method, the value for background risk (s_{BG}) was selected by matching the model with empirical data. Since I am focusing my analysis on vulnerable communities, I calibrate my background risk to produce an initial vaccination level of 68%⁴⁴.

I set the initial infection to 0 to ensure I capture only background risk. I also ensure social networks, misinformation, and heuristics are disabled. I then vary the background risk level between 0 and 0.16% at 0.001% intervals. By minimizing the squared difference between the resulting vaccination rate and 68%, I find that a background risk level of 0.025%.

A wide range of parameter values produce similar results. This implies that the model is relatively insensitive to the value of background risk.

C.1.3 Agent Error

When determining the agent's error term (ϵ_i), I need to determine both the mean and the standard deviation of the normal distribution it is drawn from. The mean I assume is 0; that is, agents have no

first infection in the model to mirror what was observed in the Netherlands.

44 This value is selected to be consistent with Dayan and Rubin (2008).

bias towards or against vaccination on average. The standard deviation (σ_E) I determine in a similar process to the above. I set the initial infection rate to 0, but this time permit both social networks and misinformation to be active. I allow the standard deviation σ_E to vary from \$1 to \$100,000.

The resulting values range extremely widely. This is expected; because the distribution is still centered at 0, on average the dispersion does not change the overall vaccination rate. Ultimately, I select a value of \$5,000 because it produces a reasonably smooth curve of vaccination over time. The model is highly insensitive to this value, so my selection does not seem to be of importance. Further analysis is discussed in Appendix D: Sensitivity Analyses (p. 97).

C.1.4. Vaccine Ambassador Impact

To calibrate the effectiveness of the vaccine ambassadors, I seek to recreate the 2.7 percentage point increase in vaccine uptake in Attwell and Freeman (2015). To do this, I set the initial infection to 0 and enable misinformation and social networks. I also require an appropriate number of vaccine ambassadors. In the Atwell and Freeman paper, the 2.7 pp increase in vaccination came from around 12,000 “views” of a Facebook poster, out of a population of 29,000. I approximate this to, on average, 40% of individuals having been exposed to the vaccine ambassador intervention.

I vary the vaccine ambassador impact variable, α , between 0 and 50 in 0.1 unit increments, to minimize the squared difference between my final vaccination rate, and a 2.7 pp increase from my baseline, resulting in a final vaccination rate of 47.5%. This resulted in a value of $\alpha=16.2$. The model was highly insensitive to this value; above roughly 16, the results were nearly identical, indicating that there may be a threshold effect.

C.2 Verification

In this outbreak, there were two stages of infection. First, the initial infection began on May 27, 2013 with two cases, and peaked by July 14, 2013 with 180 reported cases. On June 17, a national outbreak management team (OMT) began acting to stop the spread of Measles. That marks the end of the initial infection. Subsequently, parents were allowed and encouraged to vaccinate their children.

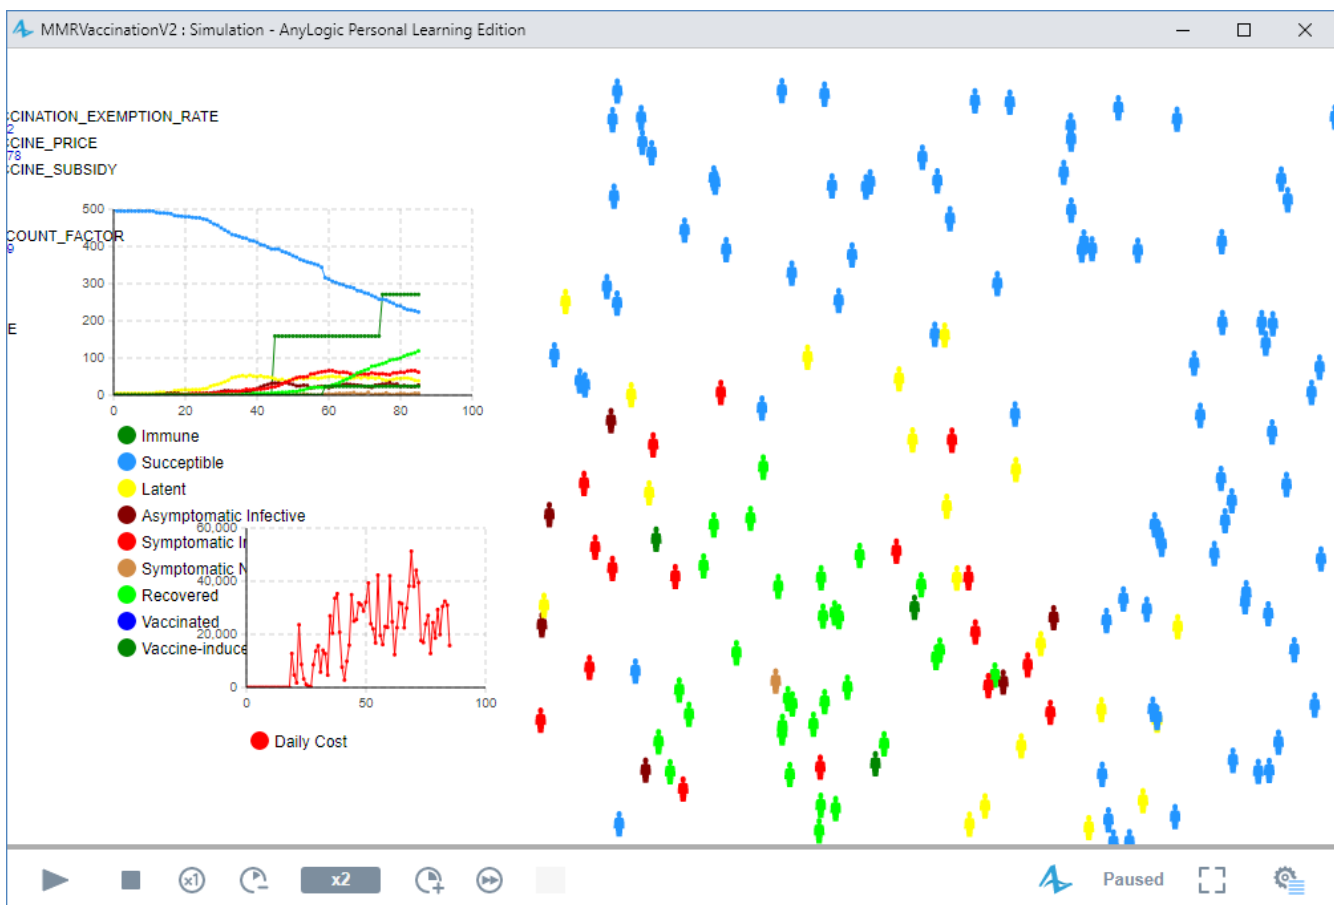


Figure 22: Model Verification

As expected, the disease starts at the point of infection, and spreads outward. As a larger portion of agents begin to show symptoms, we start to see healthy agents choose to vaccinate. This shows that our model is behaving as intended.

Appendix D: Sensitivity Analyses

This appendix contains the sensitivity analyses run on the base model as presented in Chapter 2. Each sensitivity analysis modifies one of the base parameters, and runs 1,000 simulations with each new parameter value. Unless otherwise noted, all simulations are run with heuristics, misinformation, and social networks present in the model. The metric used to evaluate performance is total societal cost per capita, measured in USD. Error bars indicate 95% confidence intervals.

D.1 Rational Agent Parameters


This section varies parameters governing all rational agents.

D.1.1 Incidence Multiplier (I_{MULT})

Our selection of I_{MULT} was initially arbitrary. Because the agents are using a linear measure of present incidence to approximate the future exponential spread of disease, any value will necessarily be an approximation, and one driven by individual psychology. Thus, I vary the incidence multiplier I_{MULT} from 0.5 to 10 in 0.5 increments. Figure 23: Impact of Incidence Perception Multiplier on Total Societal Cost (p. 98) shows that as I_{MULT} gets higher, the total societal cost decreases.

This is expected; as individuals consider each observed case to be riskier (as codified by being multiplied by I_{MULT} to determine the agent's risk) they are more likely to vaccinate sooner, leading to better societal outcomes. Furthermore, for small levels of I_{MULT} , the model is highly sensitive to the exact value chosen. However as I approach high values, the marginal impact of each unit increase of I_{MULT} becomes negligible.

Figure 23: Impact of Incidence Perception Multiplier on Total Societal Cost




This suggests that the way in individuals judge risk matters. Direct proportionality (that is, $p(\text{infection}) = I_{MULT} \times s_i$) was chosen for simplicity. However, if further evidence comes to light which suggests a more accurate mental model by which individuals gauge risk, this could be incorporated to further refine the model.

D.1.2 Agent Decision Error (ϵ_i)

To determine the impact of the agent’s heterogeneous error term ϵ_i on outcome, I vary the distribution from which that term is drawn. Recall that $\epsilon_i \sim N(0, \sigma_E^2)$. I vary σ_E from \$0 to \$10,000 in \$1,000 increments. Figure 24: Impact of Agent Error Term on Total Societal Cost (p. 99) shows that as the error term increases, so too does total societal cost. Small changes from my selected value of \$5,000 don’t change the magnitudes of the results appreciably, and the general ordering of results among the different models is unaffected.

This marginal sensitivity is expected. $\sigma_E = 0$ implies that agents perform perfect cost-benefit analyses, subject to any heuristics or misinformation present. As σ_E increases, there will be larger groups of agents that vaccinate early or late. Because the normal distribution is symmetrical, I expect the number of early vaccine recipients to be approximately equal to the number of late vaccine recipients. However, because herd immunity requires a high threshold to be effective, raising the number of late vaccine recipients has a much larger cost (in terms of propagated disease) than the benefits from the early vaccine recipients.

Figure 24: Impact of Agent Error Term on Total Societal Cost



D.1.3 Social Network Incidence Weighting (w)


I vary the amount of weight individuals put on incidence within their social network (w) from 0.5 to 10.0 in 0.5 increments. Figure 25: Impact of Increasing Social Network Weight on Total Societal Cost

(p. 100) shows that as individuals weight observations among their social network more heavily than the global average, total societal costs decrease, but at a decreasing rate.

This at first glance seems to suggest that the most effective decision making metric is based only on our local area. However, this may be an artifact of my small simulation size. It is conceivable that, for a very large population for which one individual is unlikely to know any agents in an outbreak location, using global data may yield better results. This is a limitation of the model and computing power, and so I decline to speculate on the interpretation of this relationship.

Instead, I note only that the model is sensitive to the exact value of w , and that the value it takes would help determine whether adding a social network results in an improvement or loss to society in the context of vaccination decisions.

Figure 25: Impact of Increasing Social Network Weight on *Total Societal Cost*



D.2 Misinformation Parameters

This section governs parameters that only affect misinformed agents who mistakenly believe in a link between vaccines and ASD.

D.2.1 ASD Expectation Mean (μ_{ASD})


I vary the misinformed agent belief that vaccines cause ASD by changing the distribution from which that value is drawn. Recall that the perceived probability of vaccines causing ASD is $y_i \sim N(\mu_{ASD}, \sigma_{ASD}^2)$.

I vary μ_{ASD} from 0% to 1% in 0.1% increments.

Figure 26: Impact of Mean ASD Expectation on Total Societal Cost (p. 102) shows that, as agents on average believe in a stronger link between vaccines and ASD, societal costs increase. This increase tapers off for very large values, but around the chosen ASD mean of 0.25%, small changes can lead to fairly large differences in outcome.

This is because of a threshold effect caused by the binary choice between vaccinating and not vaccinating. For very high values of μ_{ASD} , vaccine hesitant individuals consider the cost of vaccinating to be much higher than the cost of not vaccinating. As such, any small deviations caused by random chance, or by marginal changes in μ_{ASD} itself, aren't enough to change the balance point. Put another way, if an agent perceives vaccinating to cost \$300 and not vaccinating to cost \$100, even a 50% change in the cost of vaccinating will not change the outcome.

Figure 26: Impact of Mean ASD Expectation on Total Societal Cost

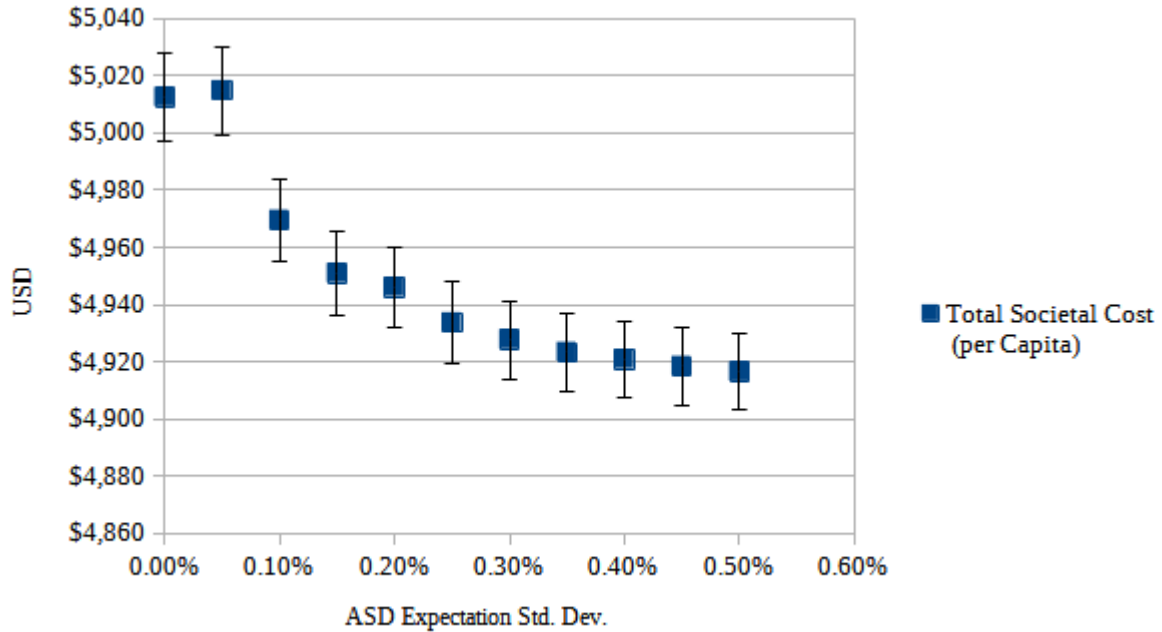


This implies that, while the model is somewhat sensitive to the mean belief in how likely vaccines are to cause ASD, it will impact the magnitudes of the results, the order of the results (namely, that misinformation and heuristics unambiguously decrease welfare) is unlikely to change.

D.2.2 ASD Expectation Standard Deviation (σ_{ASD})

Similar to above, I vary σ_{ASD} from 0 to 0.5% in 0.05% increments. Figure 27: Impact of the Distribution of ASD Expectation on Total Societal Cost (p. 103) shows that, as agents experience more variation in their belief of how likely vaccines are to cause ASD (as codified by σ_{ASD}), costs to society go down.

Figure 27: Impact of the Distribution of ASD Expectation on Total Societal Cost




This may seem counter intuitive at first. Given the behavior of the mean (see D.2.1 ASD Expectation Mean (μ_{ASD}), p. 101), I would expect the threshold effect to mean that higher dispersion would lead to an increase in social costs. By increasing the number of very-high ASD belief agents, I would expect more hold-outs, and thus a higher social cost.

Instead, for very low values of σ_{ASD} , social cost is quite high, but as σ_{ASD} rises, social costs fall, though the magnitude of the change is very small. This decrease is driven by the social aspect. Agents' socially adjusted ASD belief is changed asymmetrically by their social network's average ASD belief. Thus, when agents' individual ASD expectations vary more greatly, the smaller (more accurate) expectations weigh more heavily in all agents' socially adjusted ASD expectation.

This can be seen more clearly when I run an analysis with social networks disabled. Figure 28: Impact of ASD Expectation Distribution Without Social Networks (p. 104) shows such simulations, in which increasing dispersion actually sees an increase in social costs, as I would expect. This confirms that it's

the social network which means variation in ASD belief and improve outcomes. In a situation for which social networks are disabled, increasing the variation in ASD belief unambiguously and significantly increases social costs.

Figure 28: Impact of ASD Expectation Distribution Without Social Networks



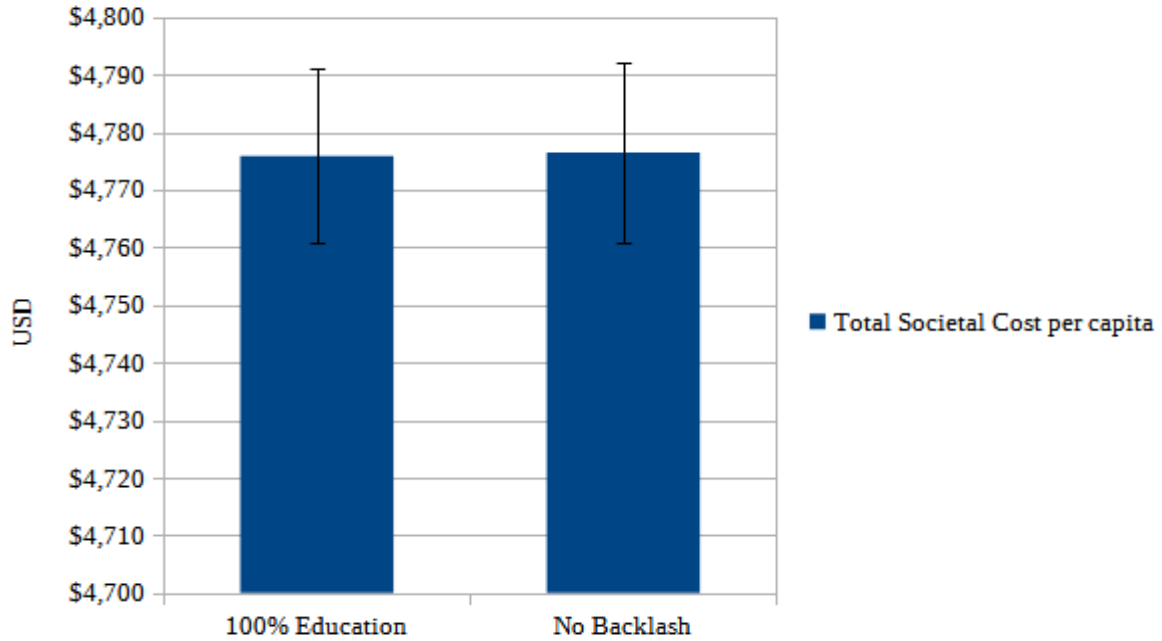
D.3 Education Backlash

The education intervention includes a 23% chance that an educated agent will be less willing to vaccinate, contrary to the intent of the intervention. To test the impact, we run a series of 1,000 simulations where, instead of resulting in a higher belief that vaccines cause ASD for 23% of agents, it instead has no effect. That is, their ASD expectation remains unchanged by education. The remaining 77% of agents have improved ASD expectations as before.

Figure 29: Impact of Education Backlash (p. 105) shows that the model is highly insensitive to this specification. Not only is there no statistically significant difference between the two cases, but the

baseline values are nearly identical. This suggests that it is not a backlash effect that drives the inefficacy of the education intervention, but rather how the education intervention itself works in the first place.

Figure 29: Impact of Education Backlash



Appendix E: Supplementary Tables

Table 11: Results of Single Interventions

	Total Societal Cost Per capita	Vaccination Rate	Attack Rate	Mortality (per 10k pop)
Baseline	\$4,764 (\$15)	81.60% (0.18%)	52.40% (0.18%)	2.16 (0.01)
Mass Vaccination	\$95 (\$2)	99.80% (0.01%)	0.20% (0.03%)	0.04 (0)
Ring Vaccination	\$1,628 (\$21)	99.70% (0.04%)	23.70% (0.32%)	0.72 (0.01)
Full Price Subsidy	\$4,760 (\$15)	81.60% (0.18%)	52.40% (0.18%)	2.12 (0.01)
\$100 Rebate	\$4,711 (\$15)	81.60% (0.17%)	51.10% (0.18%)	2.12 (0.01)
\$500 Rebate	\$4,602 (\$16)	81.40% (0.17%)	48.30% (0.19%)	2.08 (0.01)
\$1000 Rebate	\$4,473 (\$16)	81.10% (0.16%)	44.90% (0.19%)	2 (0.01)

	Total Societal Cost Per capita	Vaccination Rate	Attack Rate	Mortality (per 10k pop)
5% Vaccine Ambassador	\$1,228 (\$17)	99.00% (0.06%)	17.20% (0.26%)	0.56 (0.01)
25% Vaccine Ambassador	\$1,193 (\$17)	99.30% (0.05%)	16.70% (0.26%)	0.52 (0.01)
50% Vaccine Ambassador	\$1,212 (\$17)	99.30% (0.05%)	17.00% (0.26%)	0.56 (0.01)
100% Vaccine Ambassador	\$1,214 (\$17)	99.40% (0.05%)	17.10% (0.25%)	0.56 (0.01)
50% Education	\$4,768 (\$15)	81.80% (0.17%)	52.40% (0.19%)	2.16 (0.01)
100% Education	\$4,776 (\$15)	81.70% (0.18%)	52.40% (0.18%)	2.16 (0.01)
No Backlash	\$4,776 (\$16)	81.80% (0.18%)	52.40% (0.18%)	2.16 (0.01)

Table 12: Results of Combined Interventions

	Total Societal Cost Per capita	Vaccination Rate	Attack Rate	Mortality (per 10k pop)
Baseline	\$4,764 (\$15)	0.816 (0.18%)	0.524 (0.18%)	2.16 (0.01)
Mass Vaccination	\$95 (\$2)	0.998 (0.01%)	0.002 (0.03%)	0.04 (0)
Ring Vaccination	\$1,628 (\$21)	0.997 (0.04%)	0.237 (0.32%)	0.72 (0.01)
\$500 Rebate	\$4,602 (\$16)	0.814 (0.17%)	0.483 (0.19%)	2.08 (0.01)
5% Vaccine Ambassador	\$1,228	0.99	0.172	0.56

	Total Societal Cost Per capita	Vaccination Rate	Attack Rate	Mortality (per 10k pop)
	(\$17)	(0.06%)	(0.26%)	(0.01)
25% Vaccine Ambassador	\$1,193	0.993	0.167	0.52
	(\$17)	(0.05%)	(0.26%)	(0.01)
100% Education	\$4,776	0.817	0.524	2.16
	(\$15)	(0.18%)	(0.18%)	(0.01)
25% Vac. Amb. + Ring	\$1,079	0.997	0.152	0.48
	(\$16)	(0.03%)	(0.25%)	(0.01)
Rebate + Vac. Amb.	\$1,272	0.992	0.152	0.56
	(\$16)	(0.06%)	(0.24%)	(0.01)
Education + \$500 Rebate	\$4,630	0.814	0.484	2.08
	(\$16)	(0.18%)	(0.18%)	(0.01)
Ring + \$500 Rebate	\$1,619	0.996	0.208	0.72
	(\$20)	(0.06%)	(0.31%)	(0.01)
Educ. + Vac. Amb. + Ring	\$1,124	0.997	0.156	0.52
	(\$16)	(0.02%)	(0.25%)	(0.01)
Rebate + Educ. + Vac. Amb.	\$1,279	0.991	0.149	0.56
	(\$17)	(0.07%)	(0.25%)	(0.01)
Rebate + Educ. + Ring	\$1,656	0.996	0.21	0.76
	(\$20)	(0.05%)	(0.31%)	(0.01)
Rebate + Vac. Amb. + Ring	\$1,170	0.997	0.138	0.52
	(\$16)	(0.03%)	(0.24%)	(0.01)
All Interventions	\$1,188	0.996	0.138	0.52
	(\$15)	(0.03%)	(0.24%)	(0.01)