Proceedings of the 2025 Winter Simulation Conference
E. Azar, A. Djanatliev, A. Harper, C. Kogler, V. Ramamohan, A. Anagnostou, and S. J. E. Taylor, eds.

MODELING SOCIAL INFLUENCE ON COVID-19 VACCINATION UPTAKE WITHIN AN AGENT-BASED MODEL

Sarah K. Mulutzie[1], Sebastian A. Rodriguez-Cartes[2], Osman Y. Özaltın[2]
Julie L. Swann[2] and Maria E. Mayorga[2]

[1]Operations Research Graduate Program, North Carolina State University, Raleigh, NC, USA

[2]Edward P. Fitts Department of Industrial and Systems, Engineering North Carolina State University, Raleigh, NC, USA

ABSTRACT

Vaccination is a critical intervention to mitigate the impact of infectious disease outbreaks. However, vaccination decision is complex and influenced by various factors such as individual beliefs, access to vaccines, trust in healthcare systems, and importantly, social norms within communities, the shared understandings and expectations about vaccination behavior. This paper analyzes the impact of social norms on vaccine uptake and subsequent disease transmission by explicitly incorporating these norms into an extension of the agent-based COVID-19 simulation model, COVASIM. We aim to analyze how social norms affect vaccination rates and disease spread. We demonstrate this by implementing community-specific vaccination norms that influence agents through the perceived vaccination behaviors of their social networks. Our simulated case study explored targeted communication about vaccination uptake through different age groups. Through this intervention, we examined the effectiveness of adjusting perceptions of community vaccine uptake to better align with its true value.

1 INTRODUCTION

Vaccination represents the most effective means of preventing infectious disease outbreaks (Betsch et al. 2017; Cascini et al. 2021). A high degree of vaccination coverage is paramount for protecting individual and public health, mitigating healthcare expenditures, and sustaining global health. While ensuring access to vaccines and building public trust are fundamental, the decision to get vaccinated is multifaceted. It is shaped by various factors beyond access and trust, including personal beliefs about health and disease, and the shared understandings and expectations within communities about vaccination.

These socially shared expectations, or *norms*, dictate what is considered acceptable or typical health behavior within a group or society, significantly influencing vaccination decisions (Rimal and Real 2003). Factors such as perceived social pressure to get vaccinated, observing peers and respected community figures receiving the vaccine, and prevailing attitudes toward vaccination within an individual's social network all contribute to shaping social norms. Moreover, these norms are not static—they can evolve in response to targeted interventions and shifts in public opinion (McDonald and Crandall 2015). As such, understanding and strategically influencing social norms presents a valuable opportunity to reduce vaccine hesitancy and achieve the high levels of coverage needed to ensure robust public health protection.

Agent-based models (ABMs) are increasingly employed to simulate the complex interplay between behavioral norms and social networks (Will, Groeneveld, Frank, and Müller 2020). These models define autonomous agents that interact within a network structure, allowing researchers to observe how norms emerge from local interactions, spread through the network via mechanisms such as conformity and imitation, and influence individual behavior. This individual-centric approach, with its capacity to model heterogeneous agents and dynamic network interactions, offers significant advantages over discrete event simulations (DES) for capturing complex social phenomena. Unlike DES, which primarily models the flow of entities through predefined processes, ABMs are uniquely suited to represent emergent properties, feedback loops, and the adaptive nature of social systems driven by individual decision-making and evolving relationships, all crucial for understanding social norm dynamics. ABMs can also incorporate agent heterogeneity and environmental influences to provide a nuanced understanding of how norms are established and maintained (Matthews et al. 2007; Bianchi and Squazzoni 2015). Beyond public health, ABMs explore the dynamics of social norms in contexts such as technological diffusion, environmental actions, consumer behavior, and political engagement. These models typically represent norms via adoption thresholds, social learning, or rules for updating opinions within networks, illustrating how collective behaviors and norms evolve from local interactions (Macy and Flache 2009; Martin-Lapoirie, d'Onofrio, McColl, and Raude 2023).

In this study, we incorporate social norms that affect vaccination behavior into COVASIM, an agent-based infectious disease simulation model designed to analyze the dynamics and control of the COVID-19 pandemic (Kerr et al. 2021). Network interactions influence the perceived social norms of agents, and we explore how these norms impact vaccine uptake and disease transmission. Integrating empirical data on perceived norms with the ABM capabilities of COVASIM, the proposed approach allows us to examine the effectiveness of various public health interventions aimed at correcting misperceptions about vaccination and promoting higher vaccination rates.

2 AGENT-BASED MODELING OF SOCIAL NORMS AND VACCINE ACCEPTANCE

In this section, we outline the methodologies employed to model the influence of social norms on individual vaccine acceptance in the context of a simulated COVID-19 pandemic. We begin by reviewing the existing literature that underscores the significant role of perceived social norms in shaping health behaviors, particularly the acceptance of vaccines. We then describe the survey data utilized to empirically assess the relationship between perceived social norms and vaccine acceptance. Finally, we detail the ABM, COVASIM, which serves as the platform for our simulations, and explain how we integrate the insights from our survey analysis to model the dynamic impact of social norms on vaccination behavior.

2.1 Social Norms and Vaccine Acceptance

The perceived social norms of friends and family strongly impact an individual's vaccination intentions, with this influence declining as social distance and group heterogeneity increase (Rabb et al. 2022; Lin et al. 2022). This is consistent with the findings that individuals are more likely to accept vaccines endorsed by their "in-group" (Cruwys et al. 2021) and when they believe that their friends and family support it (Tunçgenç et al. 2021). Therefore, effective interventions should leverage trust within small communities rather than rely solely on broad messaging.

Misconceptions about community vaccination acceptance, particularly the tendency to underestimate actual uptake, can negatively influence individual decisions and hinder overall vaccination efforts (Sinclair and Agerström 2023). Moehring et al. (2023) observed that providing normative information increased the fraction of people that respondents estimated would accept a vaccine, suggesting that these underestimations in vaccination acceptance can be corrected. Hence, interventions that accurately portray vaccine acceptance through trusted networks and public health messaging can effectively increase uptake (Tayloe 2021).

To further understand these dynamics, our study leverages data from a comprehensive survey conducted by the Massachusetts Institute of Technology (MIT) in collaboration with researchers from Johns Hopkins University, the World Health Organization, and the Global Outbreak Alert and Response Network (Collis et al. 2022). This rich dataset, which includes 66,045 responses from the United States collected between October 28, 2020, and March 29, 2021, offers insights into beliefs, behaviors, and social norms surrounding the COVID-19 pandemic. The survey specifically explored topics such as masking, vaccinations, individual willingness to get vaccinated, estimates of community acceptance, and factors that influence future vaccination decisions, including recommendations from friends, family, and health officials.

To identify key drivers of vaccine acceptance, in previous work, we analyzed a subset of these responses using a random forest model, examining the influence of demographic characteristics, individual risk perceptions, trust in news sources, past vaccination behavior, and perceived social norms among other factors (Mulutzie et al. 2025). Our analysis revealed that perceived social norms were a significant predictor of the acceptance of the COVID-19 vaccine, alongside other important factors such as recommendations from trusted health authorities. Three survey questions relevant to this study are presented in Table 1.

Table 1: Survey Questions and Response Options.

Survey Question	**Response Categories**
If a vaccine for COVID-19 becomes available, would you choose to get vaccinated?	yes, no, or don't know.
Out of 100 people in your community, how many do you think would take the COVID-19 vaccine if it were available?	integer values between 0 and 100
Would you be more or less likely to take a vaccine against COVID-19 infection if it were made available and recommended to you by each of the following: friends and family, local health workers, World Health Organization, government health officials, and politicians?	more likely, less likely, no impact

2.2 Agent-Based Simulation Model

COVASIM, is an open-source ABM developed in Python that simulates the complex dynamics of the COVID-19 pandemic and the impact of interventions (Kerr et al. 2021). Each agent in COVASIM has unique demographic and epidemiological characteristics, such as age and susceptibility to infection. Agents are embedded within a multi-layered contact network, representing their social interactions across different contexts. These layers include households (representing close and frequent contacts), schools (for relevant age groups), workplaces (for working adults), and communities (representing broader, casual interactions such as those in public spaces, transit, or social gatherings). The contacts within these layers define the potential pathways for disease transmission. By incorporating stochasticity in transmission and disease progression, COVASIM allows realistic simulations of epidemic spread and assessing key epidemiological outcomes. The strength of the model lies in its ability to simulate various interventions, such as social distancing, mask-wearing, testing, and vaccination campaigns, offering a valuable tool for public health research and response. COVASIM has been used to answer policy and research questions about COVID-19 in several countries, including the United Kingdom (Panovska-Griffiths et al. 2020), Vietnam (Pham et al. 2021), and Australia (Scott et al. 2021).

Within COVASIM, infections within the model occur daily. Here, t represents time, with each t equivalent to one day. When a susceptible agent comes into contact with an infectious agent, the probability of infection is calculated based on the transmission rate (β), modified by factors such as the agents' susceptibility and infectiousness levels. Every agent has a probability of being vaccinated. In our study, this vaccination probability is based on the agent's age and changes weekly based on historical vaccination uptake. When an agent is vaccinated, this reduces their susceptibility to infection. The logic governing these daily infection and vaccination processes within COVASIM is further detailed in Figure 2. This agent-based approach, modeling both infection and vaccination at the individual level, allows for studying how individual behaviors and interventions at the contact-layer level influence overall epidemic dynamics.

2.3 Quantifying the Social Norm for Vaccination

We quantified the social norm for vaccination, referred to as *vaccination norm*, for each agent j at time t based on the perceived vaccination status of neighboring agents $i = 1, ..., k_j$. These neighbors represent agents with whom agent j has direct social contact or interaction within the model's network structure. For agent j, its perception of agent $i's$ vaccination status was represented as a binary variable V_{it} (where $V_{it} = 1$ if agent j perceives agent i as vaccinated, and $V_{it} = 0$ otherwise). Note that the number of neighbors

k_j is different for each agent, but we drop the subscript j for simplicity, and use k to represent the number of neighbors of any agent.

At the start of the simulation, agents have an initial perception of vaccine uptake within their network, reflecting who they believed will get vaccinated when it becomes available. Building on existing research that highlights the underestimation of actual vaccination uptake (Moehring et al. 2023), we initialized each agent's overall perceived uptake percentage to be, on average, 10% lower than the historical vaccination uptake in North Carolina during the time period modeled (46%). Specifically, each agent's overall perceived uptake percentage was drawn from a normal distribution with a mean of 36% and a standard deviation of 4%, a value adopted to represent a plausible range of individual perceptions in the absence of more precise, literature-supported empirical data regarding this variability. For each agent j, its drawn overall perceived uptake percentage was then used to assign the binary perceived vaccination status (V_{it}) for a corresponding proportion of its neighbors i. For example, if agent $j's$ perceived uptake was 36%, then 36% of its neighbors were assigned $V_{it} = 1$, and the remaining 64% were assigned $V_{it} = 0$. This individualized initial perception then serves as the basis for calculating the agent's baseline value of the vaccination norm (VN_{jt}). Crucially, in the absence of other influences (such as the intervention described later), this baseline perception of vaccination status remains static for each agent's neighbors.

Recognizing that closer social ties exert a more substantial influence on perceived norms, we applied a weight γ_i^l where the weight depended on the contact layer l of the neighboring agent i; l can be household, work, school, or community. This weight determines the contribution of the vaccination status of agent i to the perceived norm of agent j. The baseline value for the vaccination norm, denoted by VN_{jt}, was calculated by the following equation:

$$VN_{jt} = \frac{\sum_{i=1}^{k} \gamma_i^l V_{it}}{\sum_{i=1}^{k} \gamma_i^l} \tag{1}$$

2.4 Modeling the Effect of Vaccination Norm

To model the relationship between the perceived vaccination norm and individual vaccine acceptance (a binary outcome where 1 indicates willingness to get vaccinated and 0 indicates non-willingness) we used a subset of the survey data described in section 2.1 (N = 38,987), which had complete responses for both relevant questions. The survey elicited the vaccination norm of respondents from responses to the second question in Table 1, showing an individual's estimate for the percentage of people in their community who would get vaccinated. We categorized these percentage responses into ten categories: 1 (0-9%), 2 (10-19%), 3 (20-29%), 4 (30-39%), 5 (40-49%), 6 (50-59%), 7 (60-69%), 8 (70-79%), 9 (80-89%), and 10 (90-100%). We employed a logistic regression model to predict individual vaccine acceptance (a binary outcome of 0 or 1) based on these perceived vaccination percentages. We selected the 50-59% category (category 6) as our reference group, as its midpoint closely approximated the survey's average perceived vaccination uptake percentage.

The resulting odds ratios of this regression model provided a measure of how much more likely an individual would be to accept vaccination for each category of the perceived community vaccination percentage, relative to the reference category. We focused on modeling the potential increase in vaccination probability due to higher perceived norms. Hence, in the simulation, we did not change the base vaccination probability of agents whose vaccination norm is less than or equal to the reference vaccination norm category. Table 2 provides the odds ratios for categories 6-10 with higher vaccination norm values.

Table 2: The odds ratio values from the logistic regression for each category of vaccination norms values used for the case study.

Vaccination Norm	50%-59%	60%-69%	70%-79%	80%-89%	90%-100%
Category (n)	6	7	8	9	10
Odds Ratio	1.0	2.5	3.3	4.0	5.3

Building directly on this survey derived relationship, we incorporated these norms into COVASIM to dynamically adjust the agents' vaccination probabilities (VP_{jt}). For each agent j at time t, we first converted their base vaccination probability ($baseVacProb_{jt}$) into odds using the formula:

$$odds_{base} = \frac{baseVacProb_{jt}}{1 - baseVacProb_{jt}} \tag{2}$$

Next, we multiplied these base odds by the odds ratio ($oddsRatio_n$) corresponding to the category n of the agent's perceived vaccination norm value VN_{jt} at time t to obtain the adjusted odds:

$$odds_{adj} = odds_{base} * oddsRatio_n \tag{3}$$

Finally, we converted these new odds back into a probability to determine the agent's adjusted vaccination probability (VP_{jt}):

$$VP_{jt} = \frac{odds_{adj}}{1 + odds_{adj}} \tag{4}$$

Thus, we have established a baseline mechanism within COVASIM where agents' perceived vaccination norms, derived from network interactions and subject to underestimation, influence their vaccination probability based on the survey-derived odds ratios.

3 CASE STUDY

Building upon this foundation, this section introduces a case study in which targeted interventions are applied to adjust these perceived norms, exploring their impact on vaccine uptake and disease spread. It details the framework and implementation of our ABM within COVASIM to explore the influence of vaccination norms on vaccine uptake. We outline the population characteristics, the integration of empirically derived norms effects based on survey data and logistic regression, and the design of an intervention to shift these norms.

3.1 Population

We simulated a population of 100,000 individuals between October 1, 2020 and July 1, 2021. We selected North Carolina as the state in COVASIM, which creates a synthetic population based on Census 2020 data. No vaccines were available to the general population in North Carolina before December 14, 2020. Vaccine availability expanded in phases, beginning with healthcare workers and long-term care residents in December 2020, and expanding to all adults by April 2021, and then to children aged 12-15 years old, as depicted in Figure 1.

To establish a baseline vaccination probability for agents within our model, we used weekly vaccine uptake data from the North Carolina Department of Health and Human Services (NC DHHS) from December 14, 2020, to July 1, 2021. We matched the timing of the peak vaccination rates in the simulation with the end of December 2020 to early January 2021, as observed in the historical data from North Carolina. By calculating the weekly proportion of the North Carolina population that received vaccinations during this period by age group, we derived an initial probability of vaccination for the agents by age group in each week within our model. This approach ensured that our simulation was grounded in actual vaccination behavior when vaccines were available, before introducing the influence of varying perceived vaccination norms.

3.2 Interventions to Change Vaccination Norms

Following the baseline implementation of vaccination norms, we introduced an intervention where vaccinated individuals, belonging to a targeted subset within the model, are encouraged by local health workers to

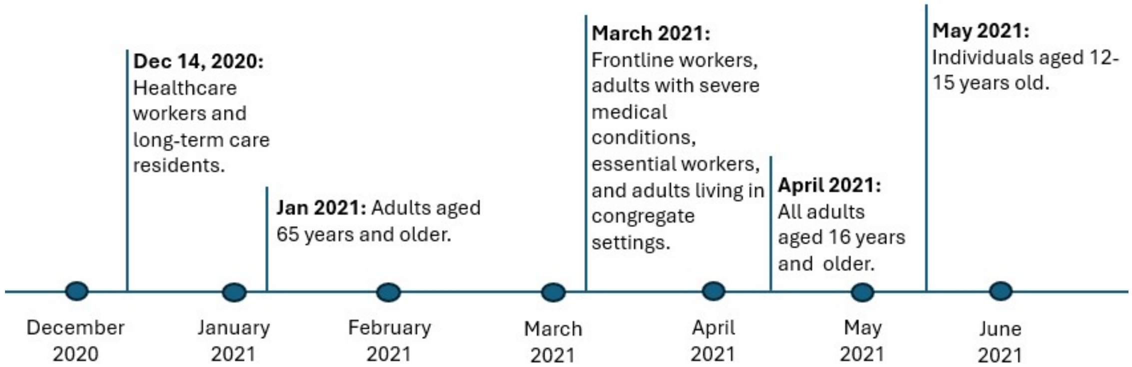

Figure 1: Timeline depicting the vaccination availability for different priority groups.

openly share their vaccination status within their household, work, school, and community networks. This intervention is designed to model the influence of trusted personal communication on vaccination norm perceptions, as individuals tend to find information more credible from family and friends (Thompson et al. 2024). We hypothesized that these interactions would correct underestimated vaccination norms, leading to an increase in perceived uptake, especially within closer social ties which exert a stronger influence on perceived norms (Rabb et al. 2022) (as modeled by our layer weights: $\gamma_i^{\text{household}} = 1$, $\gamma_i^{\text{workplace}} = \gamma_i^{\text{school}} = 0.005$, $\gamma_i^{\text{community}} = 0.001$), with these values representing the relative importance of influence from each respective social layer. This would then subsequently increase the vaccination probability of agents.

To quantify the effect of this intervention on perceived norms, we used an intervention indicator $I_{it} \in \{0,1\}$. If agent i openly shares its vaccination status at time t under the intervention, then $I_{it} = 1$, otherwise $I_{it} = 0$. The adjusted vaccination norm for agent j at time t is then calculated by:

$$VN_{jt} = \frac{\sum_{i=1}^{k} \gamma_i^l \max\{V_{it}, I_{it}\}}{\sum_{i=1}^{k} \gamma_i^l}. \tag{5}$$

In Equation (5), if agent i openly shares its vaccination status, then $I_{it} = 1$ and agent $i's$ perceived vacation of agent j is updated allowing the agents to perceive a more accurate vaccination norm. The intervention is applied to targeted population groups in our simulation. Based on age, these groups are: 18-24 year olds, 25-44 year olds, 45-64 year olds, and 65 years and older. Figure 2 depicts the general logic of our case study in COVASIM.

4 RESULTS

We established a baseline scenario to simulate vaccine uptake under the influence of social norms within COVASIM as described in section 3.1. Figure 3 illustrates weekly and cumulative vaccinations administered in our baseline scenario. This scenario was derived from COVASIM's simulation using a synthetic population constructed based on North Carolina census data. It achieves a final average vaccination coverage of 46.1% over all replications, consistent with the historical vaccine uptake of 46% observed in the North Carolina data during a comparable nine-month period. This alignment was achieved through the inclusion of vaccination norms within the model, without interventions applied. The subsequent section presents the results of our experiments. We examine whether the proposed intervention successfully corrects the underestimated vaccination norms and leads to the hypothesized increase in the perceived vaccination norm, vaccination probability, and correspondingly, vaccine uptake.

Figure 4 compares the simulated daily new infections in our vaccination norms baseline scenario (orange line at the top) with scaled historical daily new infections from North Carolina (blue line at the bottom),

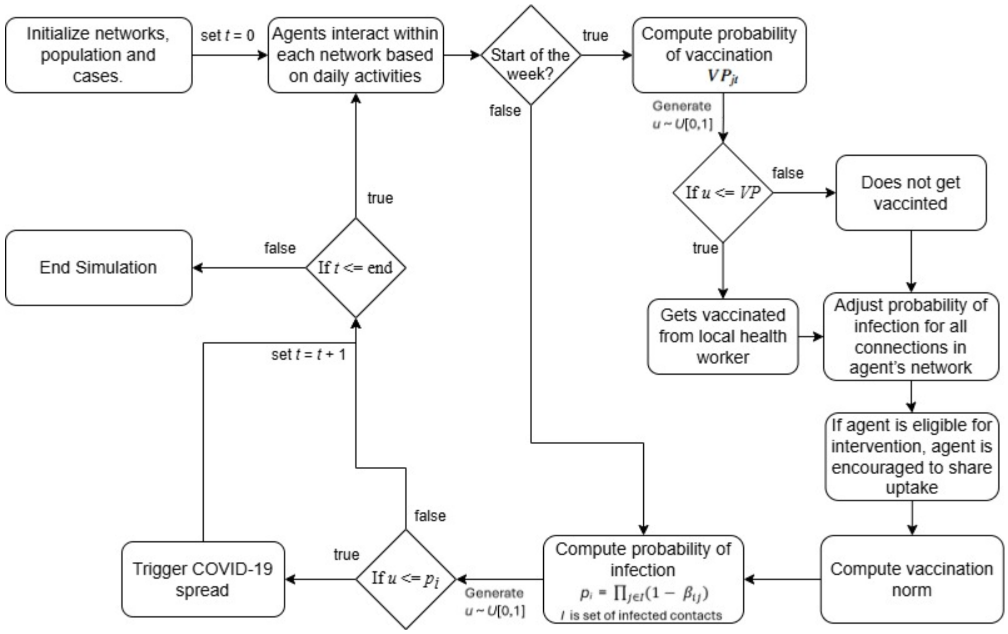

Figure 2: Flowchart depicting the case study logic in COVASIM.

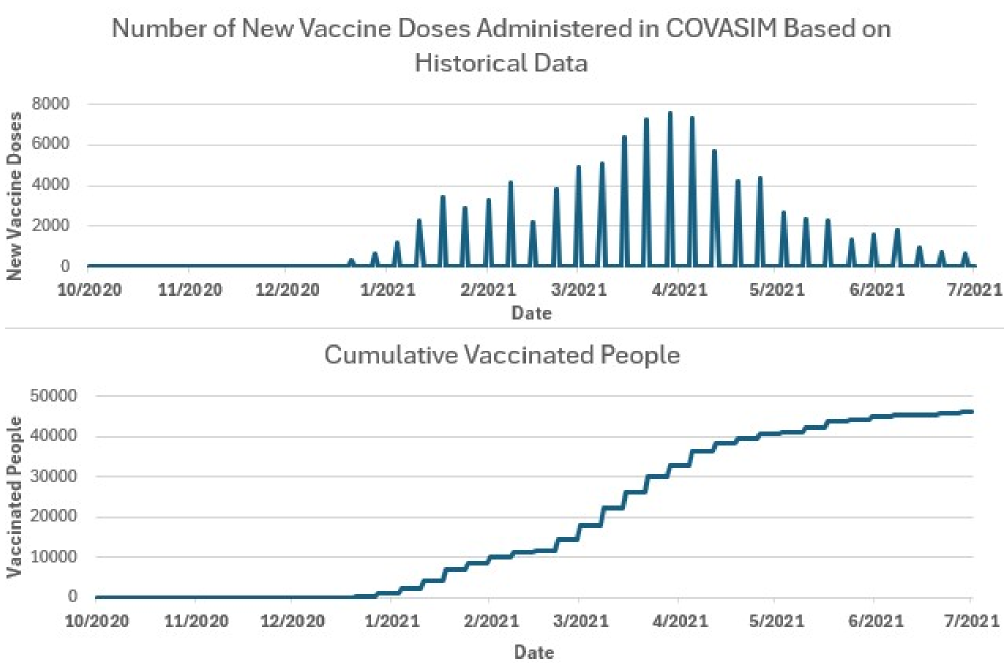

Figure 3: Vaccinations administered in baseline scenario. COVASIM vaccination norms baseline results illustrating the weekly number of new vaccination doses administered (top) and the corresponding cumulative number of people vaccinated (bottom), from October 1, 2020, to July 1, 2021 from a sample replication.

adjusted for our model's population size of 100,000 from October 2020 to March 2021. The observed difference, where simulated infections peak approximately 23.11% higher than reported data, likely reflects significant under-ascertainment in historical records due to incomplete testing and reporting, a phenomenon discussed in studies on both US (Pei, Yamana, Kandula, Galanti, and Shaman 2021) and UK data (Colman, Puspitarani, Enright, and Kao 2023). National studies in the US estimated ascertainment rates between 11.3% (Mar 2020) and 24.5% (Dec 2020) (Pei, Yamana, Kandula, Galanti, and Shaman 2021), highlighting the limitations of relying solely on reported case numbers for model validation.

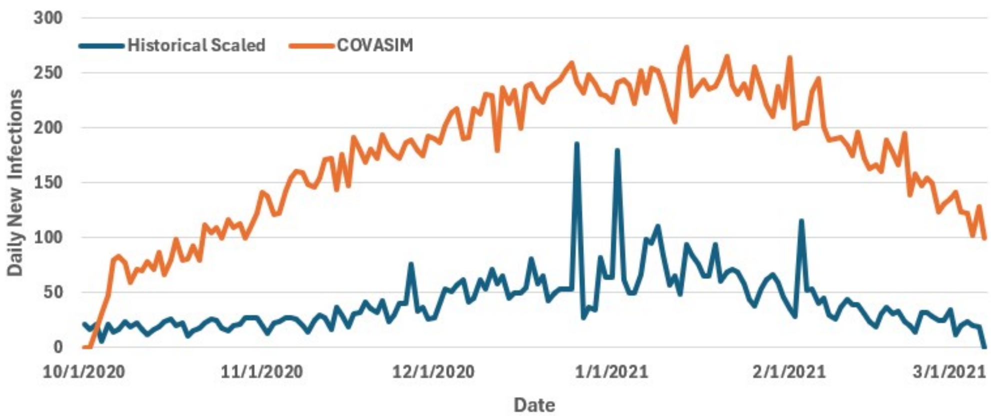

Figure 4: Simulated and scaled historical daily new infections. Daily new infections from the vaccination norms baseline simulation and scaled historical NC data (adjusted for population size) are shown for the period of October 1, 2020, to March 7, 2021.

To explore the impact of targeted interventions on vaccine uptake across the entire population and subsequent infection reduction, we implemented our intervention separately for four age groups: 18-25, 26-45, 46-64, and 65 and over. This focus on age was chosen because our random forest model indicated age as a significantly higher predictor of vaccine acceptance compared to other demographic factors such as educational level. Within our simulated population of 100,000 individuals, the sizes of these age groups varied across replications. On average, the groups comprised: 18-25 years (9,310 ± 75 adults), 26-45 years (24,480 ± 138 adults), 46-64 years (23,579 ± 139 adults), and 65 and over (16,273 ± 109 adults). Each targeted intervention scenario was run for 100 replications, with the network structure differing for each random seed. Each targeted age group is encouraged to tell the other agents of all ages in their networks of their vaccination uptake.

Figure 5(a) illustrates the percentage of the population vaccinated for the baseline (no intervention) and across the four target age group intervention. Figure 5(b) further details the cumulative additional vaccinations above baseline (no intervention) over time for with each of the four targeted group interventions. Notably, cumulative vaccination curves show an earlier increase when targeting the 65+ age group, followed by the 26-45 and 46-64 age groups.

These results consistently indicate that targeting each age group separately positively impacted vaccine uptake. The highest impact was observed when targeting the 46-64 age group, resulting in an average vaccination uptake of 48.5%. This higher impact was likely attributed to these individuals' earlier eligibility for vaccination and their proximity to the historical infection peak in the model This allows the older population to share their vaccination status earlier, potentially preventing infections. It is important to note that the effects of these interventions are not immediate. Real-world scenarios involve inherent delays in social network communication and individual vaccination scheduling. Similarly, within the model, daily agent interactions occur, but vaccination probabilities are adjusted weekly, contributing to a delayed effect on vaccine uptake from information sharing.

In terms of the impact on infections and deaths, Figure 6 presents the observed percent reductions across the different intervention groups. For total infections, our interventions led to observed reductions across all targeted age groups. However, these reductions were not statistically significant across the 100 replications (see Table 3 for detailed statistical results). Despite younger adults' significant role in social

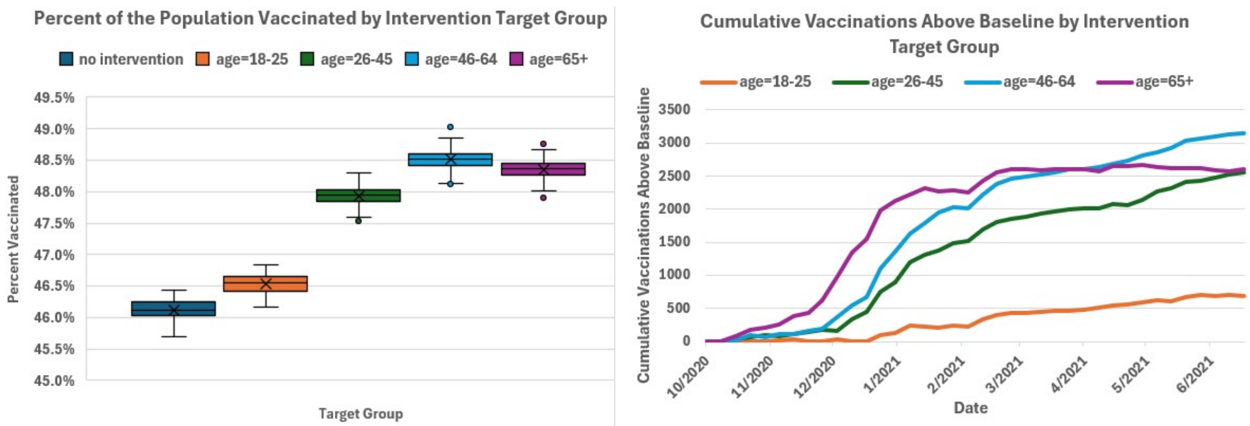

Figure 5: (a) The box plot on the left shows the distribution of final vaccination coverage for both the baseline without the intervention and the targeted vaccination interventions across multiple simulation runs. (b) The right figure illustrates the cumulative additional vaccinations above baseline runs with no intervention over time. Each line represents the results from a single model run where the intervention was applied to a different targeted group.

networks and disease transmission due to their connectivity and mobility, the timing of the intervention for these groups in our model did not lead to a substantial or significant reduction. For total deaths, our interventions resulted in observed reductions across all targeted age groups. However, none of these were statistically significant (see Table 3). Although a 0.85% mean reduction was observed for the 65+ age group, its variability and lack of statistical significance suggest a negligible impact on death reduction within our model.

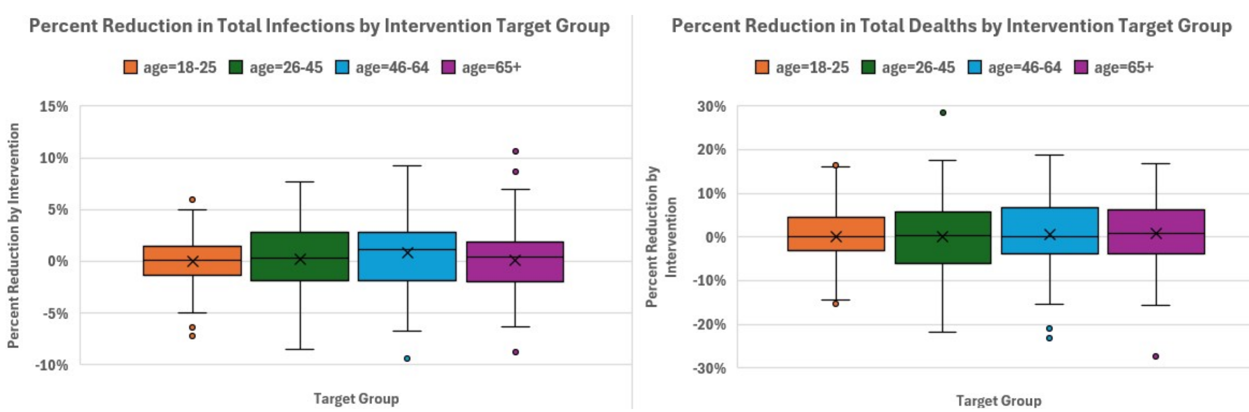

Figure 6: (a) The box plot on the left shows the impact of targeted vaccination interventions on total infections, with the largest average reduction percent observed when applying the intervention to adults aged 46-64. (b) The box plot on the right shows the impact of targeted vaccination interventions on total number of deaths, with the largest average reduction percent observed when applying the intervention to adults ages 46-64 and 65 and older.

Finally, we explored the effect of an earlier vaccination rollout, beginning nine weeks ahead of schedule on October 12th. This accelerated timeline provided a larger window to observe intervention effects during the initial infection surge, as illustrated in Figure 7. As anticipated, this earlier rollout resulted in a greater increase in vaccination compared to the baseline than the original schedule (compare Figure 7 (left) to Figure 5 (left)).

Table 3: Impact of Targeted Interventions on Total Infections and Deaths

	Targeted Age Group	Mean % Reduction	95% Confidence Interval
Infections	18-25 years	0.02%	±0.005
	26-45 years	0.18%	±0.007
	46-64 years	0.83%	±0.007
	65+ years	0.15%	±0.005
Deaths	18-25 years	0.15%	±0.013
	26-45 years	0.05%	±0.017
	46-64 years	0.47%	±0.016
	65+ years	0.85%	±0.015

In terms of the impact on infections with the earlier vaccination rollout, Figure 7 (right) presents the observed changes in total infections across the different target intervention groups. On average, the impact varied by age group: targeting 18–25 years showed an increase of 1.13% (95% CI: ±1.6%), while reductions were observed for 26–45 years (1.50%, 95% CI: ±2.1%), 46–64 years (0.16%, 95% CI: ±2.1%), and 65+ years (2.45%, 95% CI: ±1.6%). Of these observed changes, only the reduction for the 65+ age group was found to be statistically significant. This suggests that while an earlier start could improve overall vaccination coverage, its ability to substantially reduce the total infections might be limited without considering other factors or more targeted strategies.

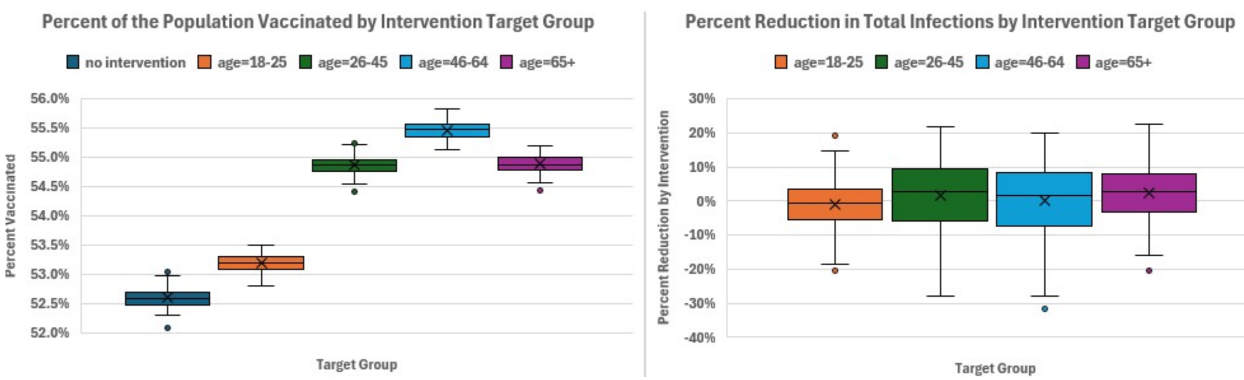

Figure 7: Impact of earlier vaccination rollout on vaccine uptake and total infections by targeted intervention groups. (a) The box plot on the left shows the distribution of final overall vaccination coverage for both the baseline without the intervention and for the targeted vaccination interventions across multiple simulation runs. The box plot shows results for all runs for each targeted intervention group. (b) The box plot on the left shows the impact of targeted vaccination interventions on total infections, with the largest average percent reduction observed when applying the intervention to adults aged 26-45 and adults aged 65 and up.

5 CONCLUSIONS

This study demonstrates the potential of targeted vaccination interventions to shape vaccination norms to influence vaccine uptake and reduce overall infections. Targeting older adults with our intervention to correct perceived vaccination norms resulted in the highest vaccination uptake, likely due to the earlier eligibility of older adults for vaccination and their higher baseline vaccination probability; that is, they are more likely to spread the news that they did get vaccinated earlier in the model horizon. The timing of vaccine availability relative to the infection peak also plays a crucial role in the effectiveness of vaccination interventions.

This study is not without limitations. First, our model primarily examines correcting underestimates of vaccination norms, not overestimates, and focuses on perceived rather than underlying social norms.

Real-world interventions may thus need to be more multifaceted. Second, the heterogeneous nature of the study population might make it difficult to assess the intervention's true effects on infection. A more homogeneous population would allow for clearer insights regarding its impact on vaccination uptake and total infections. Third, the true initial vaccination norms in real-world populations are unknown, potentially influencing baseline uptake and information response; future studies could explore their impact. Finally, the model does not account for variants or reinfections, which could alter the effectiveness of targeting younger age groups under different epidemiological conditions.

6 ACKNOWLEDGMENTS

This research was supported by the National Science Foundation under Grant 2327836. The researchers also received partial support from NC State University.

REFERENCES

Betsch, C., R. Böhm, L. Korn, C. Holtmann, and T. Rothmund. 2017. "On the Benefits of Explaining Herd Immunity in Vaccine Advocacy". *Nature Human Behaviour* 1(0056):1–6.

Bianchi, F., and F. Squazzoni. 2015. "Agent-Based Models in Sociology". *WIREs Computational Statistics* 7(126):284–306.

Cascini, F., A. Pantovich, Y. Al-Ajlouni, G. Failla, and W. Ricciardi. 2021. "Attitudes, Acceptance and Hesitancy Among the General Population Worldwide to Receive the COVID-19 Vaccines and their Contributing Factors: A Systematic Review". *EClinicalMedicine* 40:101113.

Collis, A., K. Garimella, A. Moehring, A. M. Rahimian, S. Babalola, N. H. Gobat, *et al*. 2022. "Global Survey on COVID-19 Beliefs, Behaviours and Norms". *Nature Human Behaviour* 6:1310–1317.

Colman, E., G. A. Puspitarani, J. Enright, and R. R. Kao. 2023. "Ascertainment Rate of SARS-CoV-2 Infections from Healthcare and Community Testing in the UK". *Journal of Theoretical Biology* 558:111333.

Cruwys, T., M. Stevens, J. L. Donaldson, D. Cárdenas, M. J. Platow, K. J. Reynolds *et al*. 2021. "Perceived COVID-19 Risk is Attenuated by Ingroup Trust: Evidence from Three Empirical Studies". *BMC Public Health* 21(869).

Kerr, C., R. M. Stuart, D. Mistry, R. G. Abeysuriya, G. Hart, K. Rosenfeld, *et al*. 2021. "Covasim: An Agent-Based Model of COVID-19 Dynamics and Interventions". *PLOS Computational Biology* 17(7):e1009149.

Lin, C., T. Parker, K. Pejavara, D. Smith, R. Tu, and P. Tu. 2022. ""I Would Never Push a Vaccine on You": A Qualitative Study of Social Norms and Pressure in Vaccine Behavior in the U.S.". *Vaccines* 10(1402).

Macy, M. W., and A. Flache. 2009. "Social Dynamics from the Bottom Up: Agent-Based Models of Social Interaction". In *The Oxford Handbook of Political Economy*, edited by B. Weingast and D. Wittman, 779–791. Oxford University Press.

Martin-Lapoirie, D., A. d'Onofrio, K. McColl, and J. Raude. 2023. "Testing a Simple and Frugal Model of Health Protective Behaviour in Epidemic Times". *Epidemics* 42:100658.

Matthews, R. B., N. G. Gilbert, A. Roach, J. G. Polhill, and N. M. Gotts. 2007. "Agent-Based Land-Use Models: A Review of Applications". *Landscape Ecology* 22:1447–1459.

McDonald, R., and C. S. Crandall. 2015. "Social Norms and Social Influence". *Current Opinion in Behavioral Sciences* 3:147–151.

Moehring, A., A. Collis, K. Garimella, M. A. Rahimian, S. Aral, and D. Eckles. 2023. "Providing Normative Information Increases Intentions to Accept a COVID-19 Vaccine". *Nature Communications* 14(126).

Mulutzie, S. K., S. A. Rodrguez-Cartes, M. E. Mayorga, O. Y. Özaltin, J. L. Swann, and P. J. Sheeran. 2025. "Local Health Care Workers Can Play a Crucial Role in Increasing Vaccine Acceptance During a Disease Outbreak". *PLoS ONE*. Unpublished manuscript, under review.

Panovska-Griffiths, J., P. H. Klepac, D. Kelleher, I. M. Blakey, H. Christensen, M. J. Keeling, *et al*. 2020. "Determining the Optimal Strategy for Reopening Schools, the Impact of Test and Trace Interventions, and The Risk of Occurrence of a Second COVID-19 Epidemic Wave in the UK: A Modelling Study". *The Lancet Child & Adolescent Health* 4:753–763.

Pei, S., T. K. Yamana, S. Kandula, M. Galanti, and J. Shaman. 2021. "Burden and Characteristics of COVID-19 in the United States During 2020". *Nature* 598:338–341.

Pham, Q. M., J. M. Trauer, J. McVernon, L. V. Phu, M. G. Lydeamore, R. Ragonnet, *et al*. 2021. "Estimating and Mitigating the Risk of COVID-19 Epidemic Rebound Associated with Reopening of International Borders in Vietnam: A Modelling Study". *The Lancet Global Health* 9:e916–e924.

Rabb, N., J. Bowers, D. Glick, K. H.Wilson, and D. Yokum. 2022. "The Influence of Social Norms Varies with "Others" Groups: Evidence from COVID-19 Vaccination Intentions". *PNAS* 119(29).

Rimal, R. N., and K. Real. 2003. "Understanding the Influence of Perceived Norms on Behaviors". *Communication Theory* 13:184–203.

Scott, N., A. Palmer, D. Delport, R. Abeysuriya, R. M. Stuart, C. C. Kerr, *et al*. 2021. "Modelling the Impact of Relaxing COVID-19 Control Measures During a Period of Low Viral Transmission in Victoria, Australia". *Medical Journal of Australia* 214:79–83.

Sinclair, S., and J. Agerström. 2023. "Do Social Norms Influence Young People's Willingness to Take the COVID-19 Vaccine?". *Health Communication* 38:152–159.

Tayloe, D. T. 2021. "Immunization Messaging, Communication, and Outreach Amidst the Growing Anti-vaccine Movement". *N C Med J* 82(2):118–121.

Thompson, O., M. Cristea, and M. Tamariz. 2024. "COVID-19 Vaccination Attitudes and Uptake: A Sociocultural Perspective Focusing on Parents and Peers". *PLoS ONE* 19(7):e0300771.

Tunçgenç, B., M. E. Zein, J. Sulik, M. Newson, Y. Zhao, G. Dezecache *et al*. 2021. "Social Influence Matters: We Follow Pandemic Guidelines Most When our Close Circle Does". *British Journal of Psychology* 112:763–780.

Will, M., J. Groeneveld, K. Frank, and B. Müller. 2020. "Combining Social Network Analysis and Agent-Based Modeling to Explore Dynamics of Human Interaction: A Review". *Socio-Environmental Systems Modelling* 2:16325.

AUTHOR BIOGRAPHIES

SARAH K. MULUTZIE is an Operations Research Ph.D. student at North Carolina State University. Her research interests focus on the development of mathematical models, including agent-based simulations, for infectious diseases, and the application of simulation-optimization techniques within healthcare systems. Her email address is skmulutz@ncsu.edu and her website is https://or.ncsu.edu/people/skmulutz/.

SEBASTIAN A. RODRIGUEZ-CARTES is an Industrial Engineering Ph.D. candidate at North Carolina State University. His research interests include the development of agent-based simulation models for infectious diseases and simulation-optimization applications in healthcare systems. His email address is sarodri4@ncsu.edu and his website is https://www.ise.ncsu.edu/people/sarodri4/.

MARIA E. MAYORGA is the Goodnight Distinguished Chair and Director of the Operations Research program at North Carolina State University, and a professor of personalized medicine in the Dept. of Industrial and Systems Engineering. Her research interests include predictive models in health care and health care operations management and humanitarian systems. Her email address is memayorg@ncsu.edu and her website is http://mayorga.wordpress.ncsu.edu.

OSMAN Y. ÖZALTIN is an associate professor in the Department of Industrial and Systems Engineering and a member of the Personalized Medicine Faculty Cluster at North Carolina State University. His research interests span theoretical, computational, and applied aspects of mathematical programming, focusing on optimization problems arising in public health, personalized medicine, and healthcare delivery. His methods include integer programming, stochastic programming and bi level programming. His email address is oyozalti@ncsu.edu and his website is https://www.ise.ncsu.edu/people/oyozalti/.

JULIE L. SWANN is A. Doug Allison Distinguished Professor and Department Head of the Edward P. Fitts Dept. of Industrial and Systems Engineering at North Carolina State University. Her research interests are using analytics and system approaches to enable health care and supply chains to become more efficient, effective, or equitable. Her email address is jlswann@ncsu.edu and her website is https://www.ise.ncsu.edu/people/jlswann/.