

IMPACT OF VACCINATION POLICIES FOR COVID-19 USING HYBRID SIMULATION

Felisa J. Vázquez-Abad

Daniel Dufresne

Computer Science, Graduate Center CUNY
365 Fifth Ave, New York, NY 10016

Mathematics and Statistics, Concordia University
1455 Maisonneuve Boul W, Montreal H3G 1M8

Gi-Beom Park

Computer Science, Hunter College CUNY
695 Park Ave, New York, NY 10065

ABSTRACT

A stochastic model for individual immune response is developed. This model is then incorporated in a larger simulation model for the spread of COVID-19 in a population. The simulator allows random transitions between being susceptible, exposed, having mild or severe symptoms, as well as random non-exponential sojourn times in those states. The model is more efficient than others based on geographical location, where the virus spreads according to actual distance between individuals. We are able to simulate much larger populations and vary parameters such as time between vaccinations, probability of infection, and so on. We present an application to study the effects on healthcare as a function of vaccination policies.

1 INTRODUCTION

We propose a model that describes the spread of COVID-19 in a population with varying levels of vaccination. Our model is a hybrid one, in the sense that (1) transmission of the virus occurs randomly but for the whole sub-population of susceptible people at once, not based on geographical mobility, and (2) the evolution of the infection and/or protection (from vaccines or natural immunity) is process-oriented (individuals have their own trajectory). As far as we know, this is the first model with these features.

Vaccines have been available since early 2021. In the US, the FDA gave approval for emergency use to three vaccines, with the goal of starting the vaccination process early. Even though many vaccines were produced and used around the globe (including the Chinese and Russian vaccines) the production rate and delivery logistics were insufficient in many countries. Three of the best known vaccines require people to have two doses of the vaccine in order to build protection. Our model has states representing either being unaffected by the virus or being in one of the possible stages of infection (see Section 2 for details); transitions between states are random, and sojourn times in a state is also random. When simulating the model it is straightforward to set the times when people are vaccinated, and also to decide what proportion of the population will be vaccinated. Suppose that a community does not have enough supplies to fully vaccinate 80% or more of the population. The next delivery of vaccines is expected at some time in the future. How should vaccination be administered? For instance, in early 2021 several governments decided to extend the interval between doses of the vaccine beyond what was recommended by the makers of the vaccines. That way, a greater proportion of the population was vaccinated early. Some claim that there is scientific support for this decision (Sherwin 2021). But the protection level from a single dose may be dramatically lower than that of full vaccination. Was that decision a good one? How does this affect the effects of the pandemic? We consider this problem in Section 5.

In a letter to the editor of the New England Journal of Medicine, Canadian researchers Skowronski and De Serres (2021) state that the high protective level of the first dose of the BioNTech-Pfizer vaccine justifies deferring administration of the second dose of the vaccine. In particular, they claim: “*Given the current vaccine shortage, postponement of the second dose is a matter of national security that, if ignored, will certainly result in thousands of Covid-19 related hospitalizations and deaths this winter in the United States—hospitalizations and deaths that would have been prevented with a first dose of vaccine.*” However, their letter only shows confidence intervals for the efficacy of one dose of vaccine after one and two weeks. The recommended interval by the manufacturer (BioNTech-Pfizer) is 3-4 weeks, and at the time that letter was written there was no supporting data for the effectiveness of single doses for prolonged periods, as stated in the subsequent letters to the same editor in the same journal. A UK study of the Oxford-AstraZeneca vaccine claims that longer intervals between the first and second doses yield higher antibody titers (concentration levels) after the second dose (Flaxman et al. 2021). For this vaccine, McQuade and Breskin (2021) observe: “*However, the total public health impact of the extended prime-boost interval is unclear given the trade-off between a longer period at the lower level of protection afforded by a single dose and the higher level of protection obtained after a delayed second dose. [... Antibody] titers waned over that period, suggesting that the risk of infection might increase between doses as the interval extends.*” They conclude that “*Evidence to inform decisions about COVID-19 vaccine timing and dosing is urgently needed, despite imperfect data.*”

Those citations show how important it is to study the consequences of the length of the interval between the two doses of the vaccine, particularly on (1) the proportion of people who get ill, or who are protected, and (2) on the pressure on the health system. We address this goal in this paper. The available data is scarce, because not enough tests have been performed about antibody titers in the blood as a function of the time between doses, and the available data may follow different formats for different countries, population characteristics, or vaccines. We believe that in order to study this question, a mathematical model for the simulation of the mixed immunization level within a community provides the best tool to clarify the validity of such policies. To do so, we propose in Section 2 a mathematical model for the immune response. Section 3 presents the process-oriented model for the infection, based on what is currently known about the evolution of the disease. In Section 4 we describe one way to use this model to perform simulations of the pandemic for mixed levels of vaccination in the community. Section 5 presents the results of our experiments using our simulation model with a cost function measuring the pressure on healthcare capacity in a community. While our simulations are model-based rather than data-driven, we have used real data to calibrate the hyper-parameters of the simulation.

2 MODEL FOR THE SHIELD : BUILDING PROTECTION

Our mathematical model for the protection due to vaccination, which we call *the shield*, is based on the immune response mechanisms, described in detail by Abbas et al. (2021). For completeness, Section 2.1 summarizes these mechanisms. Readers may choose to go directly to Section 2.2 if they wish to omit the biological description.

2.1 Immune Response and Vaccine Mechanisms.

When antigens enter the body and are identified by dendritic cells as “danger signals”, the dendritic cells process these antigens and carry the “T-cell epitope” (which is the *key* danger signal that can be seen by T-cells) through the lymphoid vessels towards the lymph nodes. There, they present the key danger signal to the naive T-cells. The latter differentiate into CD4 T-cells via presentation of the key danger signals and a second signal called co-stimulation. This double stimulation, produces interleukin 4 (IL4) which stimulates the proliferation of differentiated CD4 T-cells called Th2 (carrying the receptor specific for the key danger signal of the antigen). Th2 cells communicate with the B-cells, which are in the B-cell zone at the edges of the lymph nodes, activating them. Activated B-cells then differentiate into plasma cells.

All activated lymphocytes then travel through the blood towards the tissues in our bodies. These activated B-cells and plasma cells are differentiated to bond the specific antigen (the danger). When the danger is identified, the short-lived plasma cells produce the specific **antibodies** (IgG or IgM) that will bind to the identified danger. Figure 1 (adapted from Vázquez-Abad (2022)) illustrates this mechanism.




Figure 1: Immune response: triggering the production of Th2 cells with the coded ID of the danger cells.

When a vaccine is injected in a person, it is designed to be identified as a danger signal, triggering the specific immune response for that pathogen (danger). We summarize here the *primary* and *secondary* response mechanisms that explain why some of the COVID-19 vaccines require two doses.

It takes time for the body to go through the process just described, that is, dendritic cells that carry the T-cell epitope (key) of the danger signal, traveling to lymph nodes, co-stimulating T-Cells that activate the B-cells, which in turn differentiate into the short-lived plasma cells that travel through the blood to the site of danger where they synthesize antibodies.

Thus, it takes a certain amount of time for the antibodies to be produced; however, once the mechanisms are in place the antibody IgM titers increase up to a peak level; the peak level reached may or may not be protective. It normally takes between five and fifteen days from the time antigens enter the body until the peak level of antibodies is reached. (In our probabilistic model this time period is a random variable). After this primary response to antigens the antibody titers go down, but not as low as their original level, because there are long-lived plasma cells that produce antibodies. The lower level of antibodies can be protective in some cases, depending on the danger signal; for instance, humans seem to retain memory against the measles virus for life. On the other hand, for influenza viruses these lower levels of specific antibodies fall below the *protective* level. Memory is stored in some cells (called Memory B-cells) ready to respond in case of further exposure to the same danger. If Memory B-cells are still present, then when the body encounters the same danger signal a second time the production of antibody (now the more efficient IgG) occurs at a much higher rate. Not only is the time until the next peak much shorter (it can be as fast as three to five days), but the actual level of antibodies is much higher than what was achieved with the primary response. Figure 2, taken from Vázquez-Abad (2021), illustrates the process using a logarithmic vertical scale. In the figure, exposure to danger cells is referred to as “challenge”.

Vaccine development is a very complicated and long process that normally takes between five and fifteen years. In an unprecedented way, scientists and pharmaceutical companies around the world joined efforts and were able to produce vaccines within a year of first detection of COVID-19 in the Chinese province of Wuhan. Of these, we will focus here on the three best known vaccines that require an interval between the first 2 doses. Two are mRNA based : *Comirnaty*[®] (BioNTech-Pfizer) and *Spikevax*[®] (Moderna), and the last one is based on adenovirus: *Vaxzevria*[®] (Oxford-AstraZeneca).

Vaccines are meant to promote an antibody response that will be protective against the virus, while using danger signals that are actually safe for the organism. Some vaccines use a modified version of the virus, and others use a sub-unit of the virus. For the particular case of SARS-CoV-2, it was determined early in 2020 that most COVID-19 patients who recovered had naturally developed antibodies to the *spike*




Figure 2: Typical evolution of antibody levels (titers) as a function of days from exposure.

protein (S1) of the virus, while antibodies present in patients with severe and serious disease (including deaths) were not specific to this protein (minute 11:18 of Vázquez-Abad (2020b)). Thus the S1 protein of the virus was chosen as a good candidate for vaccine production (Voss et al. 2021; Malik et al. 2021). Particularly, the mRNA vaccines include the *receptor binding domain* (RBD) protein inside the S1 protein. This is the site that binds to the human cells in the alveoli eventually leading to the disease known as COVID-19. This protein itself is not dangerous to humans, but in mounting an immune attack to this as a danger antigen the virus will be recognized and neutralized, not being able to bind to the human cells. Spikevax and Comirnaty use an engineered messenger RNA (mRNA) that encodes for the surface antigen (the S1 protein that includes the RBD); Vaxzevria uses adenovirus with engineered DNA material that encodes the S1 protein that includes the RBD (see Vázquez-Abad (2020a)). After being injected with engineered genetic material that encodes the S1 protein of the virus, our body produces the aforementioned antigen; this in turn triggers the immune response to the antigen, that produces antibodies that will identify and hopefully neutralize the virus.

When developing the vaccines, clinical trials helped determine the timing and levels of the antibodies, shown schematically in Figure 2 (see Sahin et al. (2021) for details on the Comirnaty development, and Flaxman et al. (2021) for Vaxzevria). For the three vaccines that we have in mind, the primary response does not provide full protection against the COVID-19. To achieve a level of antibody that can successfully neutralize the virus, it was determined that a second dose was required. After experimentation and careful study, BioNTech-Pfizer set the interval to three weeks between doses, Moderna to four weeks and Oxford-AstraZeneca to three months. This is when the levels of IgG are low, but Memory B-cells are present so that a second exposure triggers a much higher response. Within three to five days after the second dose the levels are high enough to protect the patient. The studies report average levels from a sample population of thousands of subjects. While these results can help as a guide, individual immune responses vary depending on unknown factors, including nutrition, genetic predispositions, compromised immune systems due to other health conditions, gender, age, etc. (Levin et al. 2021).

REMARK. Attentive readers may think that after only one dose, exposure to the virus itself will trigger the much higher antibody titers. This is, unfortunately, not true. First, when administering a second dose, the amount of viral material is controlled and contains the same selected S1 protein as the first vaccine. In contrast, when the SARS-CoV-2 enters the organism it begins auto replicating at very high speed, and this does not usually allow enough time for the immune system to neutralize the virus. Second, when analyzing people who contracted COVID-19 Voss et al. (2021) observed that most of the severe cases had developed antibodies to different parts of the virus than the targeted spike protein of the current vaccines. It is as though vaccination and infection proceed independently, the two do not add up.

2.2 Mathematical Model for the Simulation

We now propose a novel stochastic model for antibody titers that captures individual variability of immune responses. Research studies identified a minimum level of antibody concentration L that *neutralizes* the virus (Sahin et al. 2021). There are only two possibilities: either a person’s antibody level is below L , in which case the person’s immune system cannot neutralize the virus, or the antibody level is above L , and the virus is neutralized (there is no partial immunity). In practice, it is impossible to measure the antibody titers every day (it would require drawing blood constantly and carrying out expensive analyses). The plots and statistical estimations for the titers found in the literature report values at the chosen timepoints when they were measured for the subjects in the trials, providing the “vertical variability” of the titers across a population (see, for example Flaxman et al. (2021) and Shrotri et al. (2021)). For our purposes we need to focus on the evolution over time of antibody titers for a single individual, that is, the *horizontal variability*.

On average (for Comirnaty), individuals who develop the shield after the first dose are protected between 14 and 21 days after vaccination; those who develop the shield after the second dose are protected from 3-4 days after the second dose to 6 months later. In Figure 3 the red dotted line is an example when level L is not reached after only the first dose. For a given individual, we denote $T(d)$ the day when vaccine dose $d \in \{1, 2\}$ is administered and $A(k)$ the level of antibody for that individual on day $k \in \mathbb{N}$. As explained above, following each dose a person may (or may not) develop the shield for some time interval.

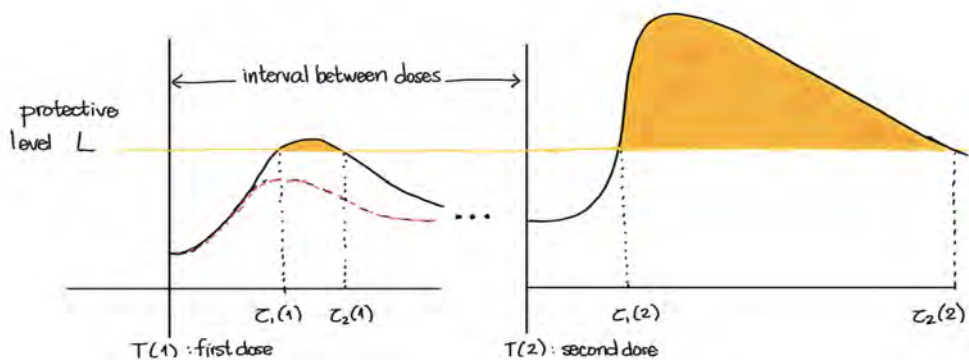


Figure 3: Evolution of antibody level over time, for one individual.

We use the following notation (refer to Figure 3) for the duration of the shield, if it occurs.

$$\tau_1(d) = \min(k : T(d) \leq k < T(d+1) : A(k) \geq L) \quad (1a)$$

$$\tau_2(d) = \min(k > \tau_1(d) : A(k) \leq L), \quad (1b)$$

where $T(d+1) = \infty$ if the person does not have a flooring dose. We say that this person has the shield at day t if $\tau_1(d) \leq t \leq \tau_2(d)$ (regions highlighted in yellow in Figure 3). The variables $\tau_i(d)$ exist only if antibody level L is reached, otherwise we define them as infinite.

Even when the person is not vaccinated, it is possible (though unlikely) that an individual’s autoimmune system will produce neutralizing antibodies very fast after exposure to the virus, so the person would not be infected. We denote π_0 the probability that such an event occurs for one non-vaccinated infected individual. The probability that someone develops the shield after just one dose is $\pi_1 = \mathbb{P}(\tau_1(1) < T(2))$. This parameter will influence how effective the shield may be for people that have had only one dose of the vaccine. As far as we are aware this important quantity (π_1) has not been estimated: it is, of course impossible to know from reported data because people who successfully combat the virus may never even realize that they were exposed, thus the numbers are not reported. Skowronski and De Serres (2021) provide some estimates for the Comirnaty vaccine around 0.45, but the confidence intervals are very wide. On the other hand the probability $\pi_2 = \mathbb{P}(\tau_1(2) < \infty)$ is the effectiveness of the vaccine in question and has

been calculated through the various clinical trials and extensive data analysis. Katella (2022) reports that Comirnaty has $\pi_2 = 0.95$, Spikevax has $\pi_2 = 0.90$ and for Vaxzevria $\pi_2 = 0.76$. Naturally this parameter will also depend on the variant of the virus in question; changes can be easily implemented in our simulator.

While the start of the protection is reported with very little variability (if it occurs), the duration of the shield has non negligible variance among individuals. We use a three parameter Beta Binomial distribution that can be used for statistical fitting to better model the immune response. It has finite support, which agrees with what we know about COVID-19 immunity, and it is also relatively efficient to generate in simulations.

3 MODEL FOR THE COVID-19 INFECTION

This section presents a *process-oriented* model that describes how individuals can change from one health state to another, once infected. The states represent the stages in the infection process and the transitions represent the evolution of the infection for a given individual. We assume that once a person is infected (meaning that their immune response did not successfully produce the protective level of antibody titers) the evolution of the disease (COVID-19) is independent of whether this person was vaccinated or not. It may, however, depend on the variant of the virus. Transitions have two parameters: a transition probability and a delay distribution (in days).

In this paper we use a terminology that differs somewhat from the one used in epidemiology. An *exposed* person carries the virus and is contagious, but there are no symptoms. An *infected* person is one who has the COVID-19 disease: the virus has successfully bound to the cells in the organism, causing the illness. We remark that the term ‘infected’ is used elsewhere to denote either those who are exposed or those who have the disease. Similarly, the word ‘exposed’ is used by the CDC and others to refer to a person who has been in close contact with other persons that have tested positive for COVID-19 (CDC 2022), but they do not necessarily carry the virus.

| Name | Symbol | Description |
|-------------------------------|--------|---|
| Susceptible | S | Individual who has no SARS-CoV-2 in their system and no COVID-19. |
| Protected | P | People who have the shield. |
| States for Infection Dynamics | | |
| Exposed | E | Individual who has SARS-CoV-2 but no symptoms (not infected). |
| Mild | IM | Infected with COVID-19 disease, mild symptoms. |
| Severe | IS | Infected with COVID-19 disease, severe symptoms. These individuals require medical care and are either in hospital or in self-quarantine. |
| Dead | D | Died from COVID-19. |

Table 1: Description of the states of each individual.

A significant difference between our model and others in the epidemiological literature is the inclusion of a state “Protected” and the omission of a state “Recovered”. Evidence shows that people who have recovered from COVID-19 may have high protective levels of antibodies (thus they are Protected). Kojima and Klausner (2022) report that the risk of re-infection can be between 1% and 4.3% (during a period of six months after recovery). Because recovered people will necessarily be either Susceptible or Protected, our model does not require explicit inclusion of the Recovered state, simplifying the simulation.

Figure 4 presents the complete model that we use for the simulation. Transitions for the infection process are represented by black arrows. The transition from Susceptible to Exposed depends on the contagion from contacts in the community, described in Section 4. Some of the transitions to and from the Protected state depend on the immune response from vaccination, as described in Section 2. The required informa-

tion for these transitions is incorporated in the data structure of the simulation as *attributes* for each individual.




Figure 4: States of the Model for COVID-19 Infection.

In the absence of vaccinations, from state E there are five transition sequences (called “infection paths”) that lead to S, and there is one transition sequence that leads to D. These are shown in (2), including the range of the time required for each transition and the corresponding probability. The symbol (—) means instantaneous transition.

| Sequence (“future”) | Probability |
|---|-------------|
| $E \xrightarrow{(-)} P \xrightarrow{(14-30)} S$ | 0.010000 |
| $E \xrightarrow{(2-14)} IM \xrightarrow{(7-21)} S$ | 0.049500 |
| $E \xrightarrow{(2-14)} IM \xrightarrow{(7-21)} P \xrightarrow{(14-30)} S$ | 0.881100 |
| $E \xrightarrow{(2-14)} IM \xrightarrow{(4-12)} IS \xrightarrow{(1-7)} S$ | 0.023760 |
| $E \xrightarrow{(2-14)} IM \xrightarrow{(4-12)} IS \xrightarrow{(30-42)} P \xrightarrow{(14-30)} S$ | 0.047520 |
| $E \xrightarrow{(2-14)} IM \xrightarrow{(4-12)} IS \xrightarrow{(4-8)} D$ | 0.009504 |

The numbers in Figure 4 and in equation (2) were calculated using a number of sources (Grant et al. (2022); DOHMH (2022); Worldometer (2021); WebMd (2022); Chung, E. (2020); Iuliano et al. (2022); Kojima and Klausner (2022)). Available data does not always refer to the same concepts we use in our model, so we had to estimate our parameters from the reported data (Park 2022). We found no statistical basis for the probability of spontaneous immune response π_0 , we chose 0.01 in our simulation. Notably, in our model the shield (protection) is an explicit attribute of individuals, but no studies have been made that correlate the various states with the antibody level of the subjects. Chen et al. (2020) and Zheng et al. (2020) use a Markov model to describe the “compartments” in the population that show the different stages, from exposure to SARS-CoV-2 to various classes of COVID-19 symptoms. Our approach is different predominantly because we use the model of the shield in order to incorporate the immune response of individuals from the vaccination. In addition, we use a Semi-Markov process to describe the evolution of the infection paths, where the distribution of the occupation time at each state is approximated using a Beta Binomial distribution, which reflects more accurately the true evolution of the disease. In Section 4 we model transmission of the virus via community contagion and incorporate this to the process oriented description of the individual infection.

REMARK. An individual is Protected when his antibody levels are high enough to neutralize the virus. That is, the viral particles in the person’s body, if any, are not capable of binding. However, if present these viral particles may stay in the person’s body for some time before degrading. Under this situation it is theoretically possible that a person with the shield who has been exposed to the virus may remain contagious for some days. In this paper we do not consider this scenario, but our model can be adapted accordingly.

4 SIMULATION MODEL

In this paper the word *transmission* means that SARS-CoV-2 virus passes from one person to another. The receiving person may or may not be able to neutralize the virus. Because of the physical way transmission is known to occur, we postulate that the frequency of transmission depends on how prevalent the virus is in the community. In reality, different people are in contact with different numbers of people over a day, and the proportion of people who are contagious varies geographically. To simplify matters, we let the probability of transmission for each susceptible person depend only on the fraction of people in the community who are contagious, among those not in confinement. In the simulations, this probability is

$$P(S \rightarrow E) = c \times \frac{E_n + IM_n}{P_n + S_n + E_n + IM_n}, \tag{3}$$

where a_n is the number of people on day n that are in state ‘ a ’. This makes sense if we believe that each day every susceptible person meets a group of people of equal number and who are exactly as contagious as the whole population. In the simulations we use $c = 0.01$ for the daily probability of contagion for each susceptible individual. This is based on heuristics, and in any case c would vary according to the particular variants present and other factors. Notice that those with severe symptoms of COVID are not counted in the contagious (see Figure 4), as they are supposed to be in quarantine or in a hospital.

Our model differs from others that have focused on physical distance between people, for instance Washburn et al. (2021). In that paper a continuous-time Markov chain model (CTMC) is used for the simulation of the transmission in public transportation, based on time of exposure to sources of contagion. Another example is Perez and Dragicevic (2009), who used an agent-based mobility model (ABM), where transmission is determined by the proximity to viral sources. A disadvantage of models of physical distance or mobility of individuals is that only a relatively small number of people can be simulated. Our model differs from SIR-type models (e.g. Liu et al. (2021)) in that ours is probabilistic. This allows the simulations to measure the variability of global outcomes around their mean. However, our model differs from the other probabilistic models of COVID we are aware of, in particular our model is not Markovian, and its structure allows the consideration of relatively large populations over periods of months or years.

The simulation begins when vaccination starts in the community. The initial population contains only E, S, IM, IS , because nobody has been vaccinated yet. The vaccination dates are generated in batches at time 0 for the susceptible and exposed (vaccines are not administered for people with symptoms (IM, IS)). The number of days between the two doses is a control variable θ (the same for everyone) and the days on which the doses are administered are staggered, assuming that v people are vaccinated each day until all those eligible have received their doses. The periods, if any, during which a person should have the shield are generated at time 0 as well, according to (1). This is the *vaccination future* that each person will have, unless the person becomes infected before the shield holds. For all those not in S, we initially generate their *infection paths* according to (2). An important function for the simulator is what we name the *consolidation* of futures. It follows some 20 rules according to various scenarios, using the following assumptions. The function is called with a vaccination future for a given person, which is assumed to just have entered the state E. Then we use the rules to determine if the illness will prevent the shields, if it will prevent vaccinations, or if the illness will finish before shields from the vaccine are created, in which case we re-write in the persons consolidated future the consecutive days of next changes and corresponding

states. The maximal known future contains 16 elements. We add also the dates of vaccination schedules and other information. When illness occurs to a person who on the day of a scheduled vaccination has symptoms (or is dead), the corresponding vaccines are freed and can be used for other eligible people in the community. All of this information about the future is processed in the consolidation function which returns for every person their particular “known future” that consolidates the infection paths and the vaccination shields. The probability of transiting from S to E depends on the current proportion of contagious in the population, so those transitions must be put into effect daily (they cannot be generated at time 0).

Once initialized, the simulation executes the following tasks for every day:

1. Allocate freed vaccines for THISDAY.
2. Find all the S that become E by generating a Binomial random variable with parameters $\|S\|, P(S \rightarrow E)$ and assign their infection paths according to (2). Finally, consolidate these infection futures with the corresponding vaccination futures.
3. Update the rest of the population: if the next scheduled change is THISDAY then execute the change of state by shifting the vector.

The way the simulation model is constructed avoids the problem with the order in which people are infected, which arises in agent-based models (ABMs). In our model, the probability of contagion (3) for a day is determined once for that whole day. It is as though that proportion was calculated every day at midnight, and didn’t change for the rest of the day. This is similar to what is done in some discrete-time actuarial models, and is the major reason why this model can be run for relatively large populations (hundreds of thousands of people over a year requires about 5 sec of CPU time on an Intel Core i7-6950x CPU with 10 cores, using the Julia programming language). By contrast, in ABM models at each time step the program must test each of the ‘agents’ (individuals) in order to determine their new position and health state and therefore contagion is not commutative, so the order in which the program passes through the list of individuals is important. Our model avoids this problem: at every time step the aggregate changes are calculated in bulk for the set of all Susceptible individuals, generating their future infection paths using a Multinomial random variable. After that, the program loops around every other individual in the population. Because the evolution of the infection path is not dependent on the rest of the population, the individual changes in state can be performed in any order on any given day. This is the process-oriented part of the simulation model, which works in parallel for all those individuals in a state within an infection path.

5 EXPERIMENTAL RESULTS: CASE STUDY

To illustrate the use of our simulation model we focus on the following scenario. The simulation starts when vaccination is commenced in a community of N people, with an initial distribution of people in various states. There are a total of $D \ll 2N$ doses of a vaccine available, and the next shipment is not expected before time T . We perform the simulation using the data consistent with the Comirnaty vaccine, but this can easily be modified to adjust to other vaccines, or to a mixture of vaccines. The decision variable θ is the *number of days between the doses* administered. The healthcare system has a vaccination capacity for ν people per day and we assume that exactly ν doses are administered per day. Assume that $D/\nu < T$, so the total number of doses is depleted before the next shipment of vaccines arrives at the community.

The cost function focuses on the overburdening of the healthcare system due to the COVID-19 pandemic. Specifically, let X_n denote the number of people in state IM on day n , and Y_n the number on state IS. The community has a healthcare capacity H_n at day n . In our model Y_n patients are taxing the system, and for the community to keep a reasonable response to the multiple healthcare demands it would be desirable to keep the COVID-19 related cases below the capacity H_n , but we have seen worldwide that this has unfortunately not been possible. We define the cost as the quantity:

$$J(\theta) = \mathbb{E} \left[\sum_{n=1}^T (Y_n - H_n)_+ \right],$$

which reflects the total amount of person/days *above* the community’s capacity. In order to model H_n we now reason that it is the essential and medical personnel the one that places themselves most at risk in a pandemic. We postulate that the expected number of healthcare personnel that will not be available due to being themselves infected, is a proportion of the total number of infected people, that is $H_n = \max(0, H_0 - \alpha(X_n + Y_n))$, where H_0 is the total capacity when all personnel is available.

Figure 5 shows the average number of people per day on states P, IM and IS. Simulating a community of $N = 100,000$ people over $T =$ six months using 100 replications took 2.5 minutes of CPU time.

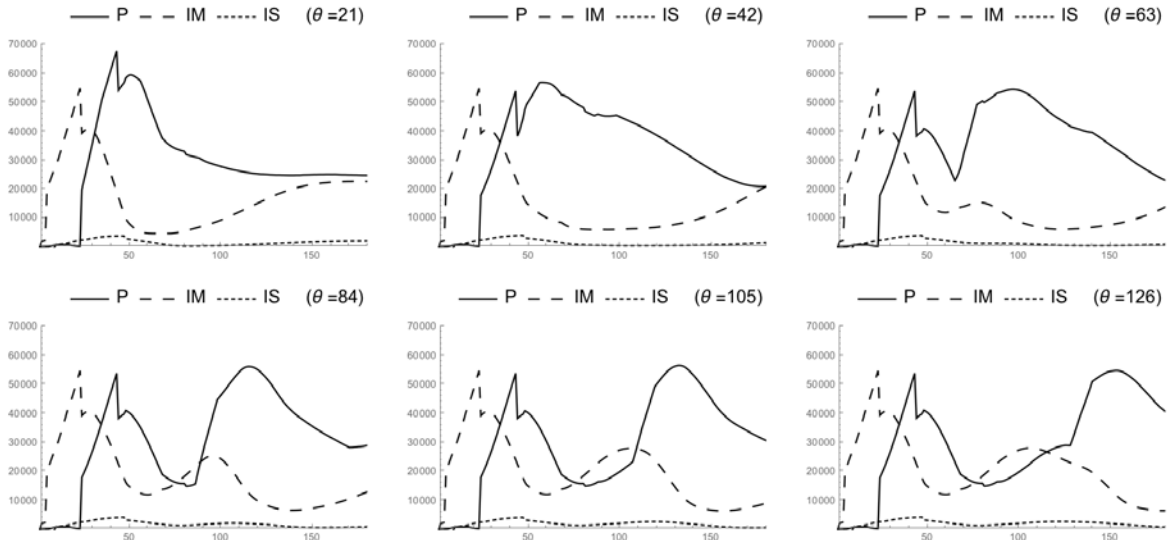


Figure 5: Average trajectories (100 replications) of the number of people in each of P, IM and IS states for $\theta \in \{21, 42, 63\}$ (top, left to right) and $\theta \in \{84, 105, 126\}$ (bottom, left to right).

It is interesting to see how much of the protection can be lost when θ is too large. Not surprisingly, because protection from vaccination is short lived for the first vaccine, long intervals between vaccines leave many people in the community without the shield, which leads to more infections early on. Figure 6 shows the results for estimating $J(\theta)$ at each of the values of θ (shown in weeks). Overlapping confidence intervals prevent us from asserting the actual value of the optimal θ^* , although it is apparent that either vaccinating too soon or waiting too much for the second dose may cause significant overburdening of the health care system. Further investigation will be performed using our simulator.

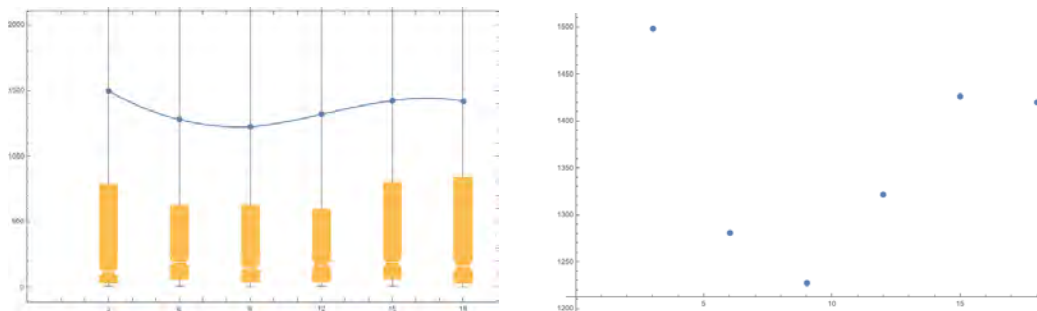


Figure 6: Box plot (left) for the estimated values of $J(\theta)$ (right) from 1000 replications of the simulation for each value of θ . The x -axis shows the control variable in number of weeks.

6 CONCLUDING REMARKS

We developed a simulation model for a mixed population integrating a process-oriented model for the immune response and the evolution of the disease with an aggregated probabilistic model for transmission of the virus in a community. Our simulator is versatile; different vaccines and variants of the virus can easily be included by adjusting the hyper parameters of the model. To illustrate the potential use of our simulator we considered a situation motivated by real life during the current pandemic. Public health officials need to put in place vaccination strategies in order to achieve herd immunity and control the load of health care. Simulation models for pandemics such as ABM models cannot handle large populations, due to the many intricate rules required to model a mixed population in different stages of protection from vaccination. On the other hand, traditional SIR models that have been designed to describe very large populations using ordinary differential equations cannot deal with issues like staggered vaccinations. Our model achieves this goal using aggregate transitions for the transmission events, while generating the infection paths as in a process oriented simulation. The simulator can be used to perform simulation optimization in a more efficient way. The model may be used for sensitivity analysis. For example, estimating the sensitivity of the effects of the pandemic with respect to π_1 may provide important insight about public health policies and vaccine administration strategies.

ACKNOWLEDGMENTS

This research was partially supported by the CUNY Institute for Computer Simulation, Stochastic Models and Optimization. Immune response pictures are the courtesy of Dr. María Dolores (Loli) Vázquez-Abad.

REFERENCES

- Abbas, A., A. Lichtman, and S. Pillai. 2021. *Cellular and Molecular Immunology*. 10th ed. ELSEVIER.
- CDC 2022. “Quarantine & Isolation”. <https://www.cdc.gov/coronavirus/2019-ncov/your-health/quarantine-isolation.html>. Accessed 28 March 2022.
- Chen, J., M. C. Fu, W. Zhang, and J. Zheng. 2020. “Supporting Real-time COVID-19 Medical Management Decisions: The Transition Matrix Model Approach”. *arXiv preprint arXiv:2007.01201*.
- Chung, E. 2020. “How COVID-19 Symptoms Develop, and How Long They Last”. <https://www.cbc.ca/news/health/typical-covid-19-progression-1.5546949>. Accessed 17 February 2022.
- DOHMH 2022. “COVID-19 Daily Counts of Cases, Hospitalizations, and Deaths”. <https://data.cityofnewyork.us/Health/COVID-19-Daily-Counts-of-Cases-Hospitalizations-and-Deaths-rc75-m7u3>. Accessed 21 March 2022.
- Flaxman, A., N. G. Marchevsky, D. Jenkin, J. Aboagye, P. K. Aley, B. Angus, S. Belij-Rammerstorfer, S. Bibi, M. Bittaye, F. Cappuccini et al. 2021. “Reactogenicity and Immunogenicity after a Late Second Dose or a Third Dose of ChAdOx1 nCoV-19 in the UK: A Substudy of Two Randomised Controlled Trials (COV001 and COV002)”. *The Lancet* 398(10304):981–990.
- Grant, R., T. Charmet, L. Schaeffer, S. Galmiche, Y. Madec, C. Von Platen, O. Chény, F. Omar, C. David, A. Rogoff et al. 2022. “Impact of SARS-CoV-2 Delta Variant on Incubation, Transmission Settings and Vaccine Effectiveness: Results from a Nationwide Case-control Study in France”. *The Lancet Regional Health-Europe* 13:100278.
- Iuliano, A. D., J. M. Brunkard, T. K. Boehmer, E. Peterson, S. Adjei, A. M. Binder, S. Cobb, P. Graff, P. Hidalgo, M. J. Panaggio et al. 2022. “Trends in Disease Severity and Health Care Utilization during the Early Omicron Variant Period Compared with Previous SARS-CoV-2 high Transmission Periods—United States, December 2020–January 2022”.
- Katella, K. 2022. “Comparing the COVID-19 Vaccines: How Are They Different?”. <https://www.yalemedicine.org/news/covid-19-vaccine-comparison>. Accessed 16 February 2022.
- Kojima, N., and J. D. Klausner. 2022. “Protective Immunity after Recovery from SARS-CoV-2 Infection”. *The Lancet Infectious Diseases* 22(1):12–14.
- Levin, E. G., Y. Lustig, C. Cohen, R. Fluss, V. Indenbaum, S. Amit, R. Doolman, K. Asraf, E. Mendelson, A. Ziv et al. 2021. “Waning Immune Humoral Response to BNT162b2 Covid-19 Vaccine over 6 Months”. *New England Journal of Medicine* 385(24):e84.
- Liu, Y., S. Srivastava, Z. Huang, and F. J. Vázquez-Abad. 2021. “Pandemic Model with Data-driven Phase Detection, a Study Using COVID-19 Data”. *Journal of the Operational Research Society*:1–15.
- Malik, Y. S., P. Kumar, M. I. Ansari, M. G. Hemida, M. E. El Zowalaty, A. S. Abdel-Moneim, B. Ganesh, S. Salajegheh, S. Natesan, S. Sircar et al. 2021. “SARS-CoV-2 Spike Protein Extrapolation for COVID Diagnosis and Vaccine Development”. *Frontiers in Molecular Biosciences* 8.

- McQuade, E. T. R., and A. Breskin. 2021. “Longer Intervals and Extra Doses of ChAdOx1 nCoV-19 Vaccine”. *The Lancet* 398(10304):933–935.
- Park, G. 2022. “ABM Simulation Model of a Pandemic for Optimizing Vaccination Strategy”. Master’s thesis, Hunter College of CUNY.
- Perez, L., and S. Dragicevic. 2009. “An Agent-based Approach for Modeling Dynamics of Contagious Disease Spread”. *International Journal of Health Geographics* 8(1):1–17.
- Sahin, U., A. Muik, I. Vogler, E. Derhovanessian, L. M. Kranz, M. Vormehr, J. Quandt, N. Bidmon, A. Ulges, A. Baum et al. 2021. “BNT162b2 Vaccine Induces Neutralizing Antibodies and Poly-specific T Cells in Humans”. *Nature* 595(7868):572–577.
- Sherwin, C. 2021. “An Immunologist Makes the Case for Delaying Quebec’s Second Dose of the COVID-19 Vaccine”. <https://montreal.ctvnews.ca/an-immunologist-makes-the-case-for-delaying-quebec-s-second-dose-of-the-covid-19-vaccine-1.5261221>. Accessed 23 February 2022.
- Shrotri, M., A. M. Navaratnam, V. Nguyen, T. Byrne, C. Geismar, E. Fragaszy, S. Beale, W. L. E. Fong, P. Patel, J. Kovar et al. 2021. “Spike-antibody Waning after Second Dose of BNT162b2 or ChAdOx1”. *The Lancet* 398(10298):385–387.
- Skowronski, D. M., and G. De Serres. 2021. “Safety and Efficacy of the BNT162b2 mRNA Covid-19 Vaccine”. *New England Journal of Medicine* 384(16):1576–1577.
- Vázquez-Abad, M. D. 2020a. “Evolution: The Making of Vaccines”. <https://youtu.be/JWK9rHWuik8>. Accessed 24 March 2022.
- Vázquez-Abad, M. D. 2020b. “Hope: Developing a COVID-19 Vaccine”. <https://www.youtube.com/watch?v=gTkBVIX0-vQ&t=10s>. Accessed 24 March 2022.
- Vázquez-Abad, M. D. 2021. “In 10 Minutes: Primary and Secondary Antibody Response”. <https://youtu.be/K54HI8K6JF8>. Accessed 24 March 2022.
- Vázquez-Abad, M. D. 2022. “Antibody Response Th2 Cells and B Cells”. <https://www.youtube.com/watch?v=xzKIYU3UwUg>. Accessed 24 March 2022.
- Voss, W. N., Y. J. Hou, N. V. Johnson, G. Delidakis, J. E. Kim, K. Javanmardi, A. P. Horton, F. Bartzoka, C. J. Paresi, Y. Tanno et al. 2021. “Prevalent, Protective, and Convergent IgG Recognition of SARS-CoV-2 non-RBD Spike Epitopes”. *Science* 372(6546):1108–1112.
- Washburn, A., Y. Paing, P. Lin, and F. Vázquez-Abad. 2021. “Simulating SARS-CoV-2 Transmission in the New York Subway”. Proceedings of the 2021 Winter Simulation Conference.
- WebMd 2022. “Coronavirus Recovery”. <https://www.webmd.com/lung/covid-recovery-overview#1>. Accessed 30 March 2022.
- Worldometer 2021. “Coronavirus Symptoms (COVID-19)”. <https://www.worldometers.info/coronavirus/coronavirus-symptoms/>. Accessed 14 March 2022.
- Zheng, Z., K. Wu, Z. Yao, X. Zheng, J. Zheng, and J. Chen. 2020. “The Prediction for Development of COVID-19 in Global Major Epidemic Areas through Empirical Trends in China by Utilizing State Transition Matrix Model”. *BMC Infectious Diseases* 20(1):1–12.

AUTHOR BIOGRAPHIES

FELISA J. VÁZQUEZ-ABAD is Professor of Computer Science (CUNY) and Principal Investigator at the School of Computing and Information Systems (University of Melbourne). She is interested in stochastic optimization and computer simulation of complex systems under uncertainty. She obtained a B.Sc. in Physics (1983) and a M.Sc. in Statistics and Operations Research (1984) from UNAM (Mexico). In 1989 she obtained a Ph.D. in Applied Mathematics from Brown University. She was postdoctoral researcher at the INRS-Telecom in Montreal, Canada from 1990 to 1993. She was a professor at the University of Montreal (1993-2004), and then a professor at the University of Melbourne, until 2009 that she moved to New York to join the CUNY Faculty. In 2000, she was a recipient of the Jacob Wolfowitz award for advances in the mathematical and management sciences. Her team was a finalist in the MEDSTART competition (2014) and she lead the award-winning Hunter Hawks team in the CUNY-IBM Watson Competition (2017). She has participated in Grant Selection Committees and has been Associate Editor for IEEE Transactions on Automatic Control, Management Science, and Operations Research Letters, Area Editor of the ACM Transactions on Computer Modeling and Simulation, and web editor of the INFORMS College on Simulation. She actively participates in events and programs to encourage women and minority students to succeed in Computer Science. Her email addresses are felisav@hunter.cuny.edu, felisav@unimelb.edu.au.

DANIEL DUFRESNE is Affiliate Professor at the Department on Mathematics, Statistics and Actuarial Science at Concordia University in Montreal. Email address: daniel.dufresne@concordia.ca

GI-BEOM PARK is currently finishing his master’s degree in computer science at Hunter College, CUNY. Research interests are stochastic modeling and simulation, stochastic optimization, formal methods and applied logic. Email address: gibeom.park04@myhunter.cuny.edu.