medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

Allocation of COVID-19 Vaccines Under Limited Supply

Xin Chen[1], Menglong Li[2], David Simchi-Levi[3], and Tiancheng Zhao[4]

[1]Department of Industrial Enterprise and Systems Engineering, University of Illinois at Urbana-Champaign, Urbana, IL 61801. xinchen@illinois.edu

[2]Department of Industrial Enterprise and Systems Engineering, University of Illinois at Urbana-Champaign, Urbana, IL 61801. ml10@illinois.edu

[3]Institute for Data, Systems, and Society, Massachusetts Institute of Technology, Cambridge, MA 02139. dslevi@mit.edu

[4]Gies College of Business, University of Illinois at Urbana-Champaign, Urbana, IL 61801. tz14@illinois.edu

Abstract

This paper considers how to allocate Covid-19 vaccines to different age groups when limited vaccines are available over time. The disease dynamics is specified by an age-structured SAPHIRE model whose parameters are estimated by the standard least square method using the epidemic data from New York City. We derive optimal static allocation policies with different objectives under different amounts of daily available vaccines, and examine several dynamic allocation heuristics including old-first policy, infection-first policy, myopic policy, death-weighted myopic policy, and two-day myopic policy. For static policies, our numerical study shows that to minimize the total deaths, it is optimal to allocate limited vaccines to the oldest group first and then the younger group if there are capacities remaining. In contrast, to minimize the total confirmed cases, the optimal static policy allocates a considerable portion of vaccines to younger groups even if the daily available vaccines are very limited. The optimal static policies achieve a much smaller number of confirmed cases and deaths compared to two benchmark policies: a uniform allocation policy that allocates available vaccines equally to each age group, and a proportional allocation policy that allocates available vaccines proportionally to the population of each age group. For dynamic allocation policies, the myopic policy and the two-day myopic policy have similar performance and significantly outperforms the other dynamic heuristics and the static policies in terms of the confirmed cases and deaths.

NOTE: This preprint reports new research that has not been certified by peer review and should not be used to guide clinical practice.

medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

1 Introduction

The pandemic Covid-19 has caused tremendous loss to the global economy since December 2019. To contain this pandemic, over 165 vaccines are being developing, among which 37 vaccines are in human trials (Corum et al. 2020). Most experts believe that a vaccine is likely to become available by mid-2021, about 12-18 months after Covid-19 first emerged (Gallagher 2020). However, the early Covid-19 vaccine supplies are far from enough for everyone, even for people at high risk such as healthcare workers (Loftus 2020). In order to have an early access to Covid-19 vaccines, the US government has made an agreement with Pfizer Inc. and BioNTech SE to provide 100 million doses of Covid-19 vaccines once proved to work safely (Wack 2020). An important question is how to allocate the limited vaccines to different groups of people when it becomes available over time. It is commonly agreed that highest-risk medical, national security and other essential workers should get the vaccine first (Sun 2020). A challenging problem is to how to allocate vaccines for the remaining large population.

In this paper, we focus on vaccine allocation policies for ordinary people grouped by ages. We follow the Centers for Disease Control and Prevention (CDC) to classify people into 5 age groups: 0-17, 18-44, 45-64, 65-74, 75+. The epidemiological model we use is an age-structured SAPHIRE model whose single population version is proposed by Hao et al. (2020). See Zhou et al. (2019) and the reference therein for more on age-structure in epidemic models. Our model has seven compartments: susceptible compartment, exposed compartment, presymptomatic infectious compartment, unascertained infectious compartment, ascertained infectious compartment, isolated compartment, and removed compartment. Each compartment is further divided into five age groups. We use the Covid-19 data from New York city during March 17 to June 8 to estimate parameters in our model. The parameter estimation method we use is the standard least squares method, i.e., to minimize the square errors of the prediction made by the epidemic model. Since the disease dynamics is nonlinear, the accumulated squared error is a highly nonlinear and nonconvex function of the parameters, which leads to many local optimal solutions. To get a reasonable optimal solution, we further impose conditions on the mixing pattern of age groups and solve the optimization problem many times with randomly choosing initial parameters.

Based on the estimated parameters, we solve the optimal static allocation policies under different objectives and different amount of daily available vaccines and conduct sensitivity

medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

analysis. The objectives we consider include total number of confirmed cases, total number of death cases, and their weighted sum. Our numerical study shows that to minimize the total deaths, it is optimal to allocate limited vaccines to the oldest group first and then the younger group if there are capacities remaining. When the daily supplies increase, the older group still gets the majority of the vaccines while the younger groups gradually get more. The optimal static policy that minimizes the total confirmed cases follows a similar pattern but allocates doses to younger groups even if the daily available vaccines are limited. When the daily supplies increase, the younger groups get the majority of the vaccines. We also compare these two optimal static polices with two benchmark polices: a uniform allocation policy that allocates available vaccines equally to each age group, and a proportionally allocation policy that allocates available vaccines proportionally to the population of each age group. Our numerical study shows that the optimal static policies have a significant decrease of total confirmed cases and deaths compared to the two benchmarks.

We evaluate several dynamic allocation policies under the setting of constant daily supply including an old-first policy which allocates daily supplies in a decreasing age order, an infection-first policy which allocates daily supplies proportionally to the current infection ratio of each group, a myopic policy which determines today's allocation by minimizing tomorrow's total new infections of all groups, a death-weighted myopic policy, and a two-day myopic policy which determine the allocation for today and tomorrow by minimizing the total new infections of the day after tomorrow. Our numerical study shows that the myopic policy and the two-day myopic policy have similar outcomes and significantly outperforms the other dynamic allocation heuristics and optimal static policies in terms of both confirmed cases and deaths. The myopic policy allocates daily available vaccines to the group with the largest marginal effect of a unit vaccine, which leads to a single group allocation. The two-day myopic policy has a mixing allocation over groups, but still follows a similar pattern with the myopic policy in that the group that gets the majority of the daily supply is roughly the same with the group that gets all under the myopic policy. This suggests the decision maker to take the marginal effect of a unit vaccine of each group into consideration when making allocation.

The organization of the paper is as follows. In Section 2, we formally introduce the age-structured SAPHIRE model. Section 3 provides the description of the data set and details of parameter estimation. Section 4 contains comprehensive numerical experiment of allocation policies. We conclude the paper in Section 5.

medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

2 Model

In this section, we present an age-structured SAPHIRE model. In this model, the total population is divided into seven compartments including susceptible compartment, exposed compartment, presymptomatic infectious compartment, unascertained infectious compartment, ascertained infectious compartment, isolated compartment, and removed compartment. Each compartment is further divided into five age group (labeled 1,...,5 in ascending order): 0-17, 18-44, 45-64, 65-74, 75+. The dynamics of the age-structured SAPHIRE model is given by the following ODEs (1)-(7). Firgure 1 illustrates the status transition of age group 1 and how susceptible individuals in age group 1 are infected. We should mention that for simplicity, this figure does not specify all the transitions.

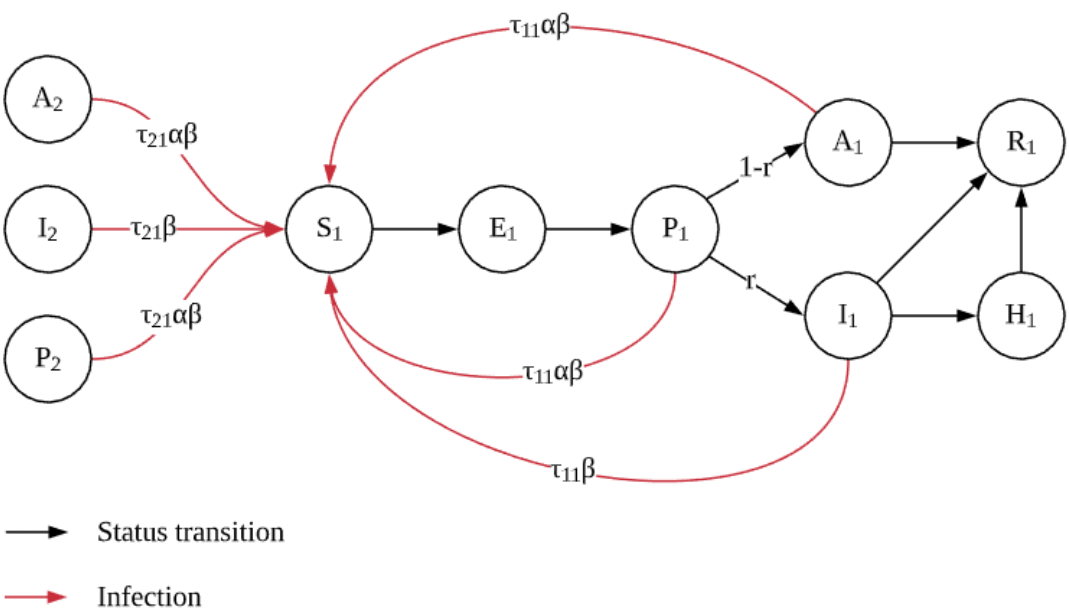

Figure 1: Disease dynamics for the SAPHIRE model

medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

$$\frac{dS_i}{dt} = -a(t)(S_i - v_i(t))\beta \sum_j \frac{\tau_{ji}(I_j + \alpha(P_j + A_j))}{N_j} - v_i(t) \tag{1}$$

$$\frac{dE_i}{dt} = a(t)(S_i - v_i(t))\beta \sum_j \frac{\tau_{ji}(I_j + \alpha(P_j + A_j))}{N_j} - \frac{E_i}{D_e} \tag{2}$$

$$\frac{dP_i}{dt} = \frac{E_i}{D_e} - \frac{P_i}{D_p} \tag{3}$$

$$\frac{dA_i}{dt} = \frac{(1 - r_i)P_i}{D_p} - \frac{A_i}{D_r} \tag{4}$$

$$\frac{dI_i}{dt} = \frac{r_i P_i}{D_p} - \frac{I_i}{D_r} - \frac{I_i}{D_q} \tag{5}$$

$$\frac{dH_i}{dt} = \frac{I_i}{D_q} - \frac{H_i}{D_h} \tag{6}$$

$$\frac{dR_i}{dt} = \frac{A_i + I_i}{D_r} + \frac{H_i}{D_h} + v_i(t), \tag{7}$$

where the meaning of the notations in (1)-(7) are listed below.

S_i: the number of susceptible individuals in age group i;

E_i: the number of exposed individuals in age group i;

P_i: the number of presymptomatic infectious individuals in age group i;

A_i: the number of unascertained infectious individuals in age group i;

I_i: the number of ascertained infectious individuals in age group i;

H_i: the number of isolated individuals in age group i;

R_i: the number of removed individuals in age group i;

N_i: the population size of age group i;

β: the transmission rate due to the contact between an infectious individual and a susceptible individual;

τ_{ij}: the contact rate of an individual in age group i with an individual in age group j;

α: the discount factor of the transmission rate due to the contact between an unascertained infectious individual and a susceptible individual;

D_e: the average time from exposed to infectious;

r_i: the fraction of ascertainment in age group i;

D_p: the average time from presymptomatic infectious to symptomatic infectious;

D_r: the average time from symptomatic infectious to recovered;

D_q: the average time from ascertained infectious to isolation;

D_h: the average time from isolation to recovered;

medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

$a(t) \in [0, 1]$: the level of permitted economic activities at time t;

$v_i(t)$: the amount of vaccination allocated to age group i at time t.;

The dynamics (1)-(7) is based on the following assumptions. The rate of new infections follows a mass action law, i.e., the number of new infections is proportional to the number of susceptible individuals and the number of infectious individuals (commonly assumed in the epidemiological literature, see Hamer 1906). We assume that the rate of new infections is discounted by the level of permitted economic activities (see Birge et al. 2020). The transmission of unascertained infectious is discounted by a factor α (Li et al. 2020). We assume that the birth rate and the death rate is zero in our model as the data we use ranges around 3 months and the birth and the death from natural causes can be ignored. We use $(\tau_{ji})_{i,j}$ to model the mixing rate of different age groups. See Fumanelli et al. (2012), Liu et al. (2020) for more on age-specific social contact characterizations.

3 Data Description and Parameter Estimation

In this section, we describe the data set and how we use it to estimate parameters in our model. The epidemic data we use is disclosed from the NYC Department of Health and Mental Hygiene (NYC Health 2020), which covers the epidemic trajectory of the New York City, including daily confirmed cases and death cases, etc. Reported cases are divided by 5 age groups, 0-17, 18-44,45-64, 65-74, and 75+ respectively. The age group population information is drawn from a 2017 census (Baruch College 2017). Some summary statistics of different age groups is provided in Table 1. The table includes the total number of confirmed cases and deaths in NYC from March 17 to June 8. It also covers the proportion of population that belongs to the group, where the total population in NYC (Manhattan area) is around 1.58 million.

age	total cases	total deaths	population ratio	population
0-17	6590	0	0.228	360240
18-44	74828	678	0.384	606720
45-64	77460	4019	0.245	387100
65-74	25307	4309	0.080	126400
75+	24607	8583	0.064	101120

Table 1: Age groups

A brief timeline of the epidemic progression and the city's response is outlined as the follows

medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

(Wikipedia contributors 2020). On March 1, the first case of COVID-19 was confirmed in New York State. On March 3, the first recorded person-to-person spread cases was confirmed. The epidemic then went through an exponential growth, with number of confirmed cases of 17,800 by March 25. On the other hand, on March 14, all New York public libraries were shut down. On March 17, facilities including theaters, concert venues, and nightclubs are closed. And on March 22, a stay at home order was put into effect where the majority of businesses are paused. As the situation started to turn better in May and June, the city resumed operations according to a four-phase reopening plan, which began on June 8, and on July 20, the final phase is being executed.

We use the epidemic data of New York City, from March 17 to June 8, a total of 84 days, to estimate our model parameters. The reasons for this selection can be summarized as the follows. First, New York City is the first major US city struck by the epidemic, and it went through both the increasing phase and declining phase of the epidemic, which provides us rich data about different stages of the epidemic. Second, the time period of our study is approximately during the execution of the stay-at-home order, where the level of social and economical activity stays roughly constant. This simplifies the model as we do not need to incorporate different level of economic activity in the model, which is not our focus in this paper. Also, since lock down is a common action taken when facing COVID-19, the estimated model parameters and insights drawn from the model can be more useful in places out of New York City at different times. Third, since we are focusing on the disease spread between different age groups rather than geographical regions, the New York city is relatively the best fit because it has high population density in a rather small area, compared to places like Florida or Texas. But we may expect that the transition parameters can be quite different at different locations, due to factors including local age structure, level of urbanization, etc.

The way we estimate the parameters is described as follows. Some parameters related to the Covid-19 disease are set according to Hao et al. (2020): $D_r = 2.9$, $D_h = 30$, $D_e = 2.9$, $D_p = 2.3$, $D_q = 6$, $\beta = 1.4$ and $\alpha = 0.55$. Although these parameters are estimated from the epidemic data of Wuhan, we would expect disease-specific parameters to be similar in NYC. See Birge et al. (2020) for the same treatment of using disease-specific parameters from Wuhan for epidemic models of NYC. But we caution that these parameters can be quite different at different locations. The parameters remain to be estimated are the contact rate matrix $(\tau_{ij})_{1\leq i,j\leq 5}$ among age groups and the ascertainment rate $r_1, ..., r_5$ for the five age groups.

medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

Given a set of all the parameters of the model and the initial number of individuals in each age group and compartment pair, we are able to compute the number of individuals in each age group and compartment pair in all future dates including the daily number of confirmed cases and death cases of each age group. Our objective function (a function of the parameters) is the accumulated squared error of the daily number of confirmed cases predicted by the model to the actual data. We optimize the objective function with python's built-in minimization function (SLSQP algorithm). In fact, since the ODEs that define our model are non convex, there exist many local optimal solutions to the objective function. A detailed procedure of how we deal with the local optimal solutions can be found in Appendix 6. The parameter estimation results are in Table 2 and Table 3.

τ	0-17	18-44	45-64	65-74	75+
0-17	0.0657	0.0255	0.0000	0.0000	0.0000
18-44	0.0151	0.0551	0.0053	0.0053	0.0053
45-64	0.0000	0.0084	0.0832	0.0131	0.0130
65-74	0.0000	0.0256	0.0400	0.0397	0.0397
75+	0.0000	0.0320	0.0496	0.0496	0.0878

Table 2: Estimation of contact rates among age groups

	0-17	18-44	45-64	65-74	75+
r	0.014	0.105	0.142	0.142	0.183

Table 3: Estimation of ascertainment rates

This set of parameters provides a fairly good prediction given the initial states. The root mean squared error (RMSE) for confirmed cases and deaths in NYC are 1183 and 83, respectively. The plot for the predicted confirmed cases and death against actual data are shown in Figure 2 and Figure 3 respectively. From the figure we can see that our prediction fits pretty well with the actual data, both in confirmed cases and deaths. Figure 4 represents the predicted daily confirmed cases from each age group. From this figure, we can see that the peak of each age group arrives approximately at the same time, which happened in reality. The relative ratio between each groups agrees with the data with an error less than 5%.

Some insights can be drawn from the estimated parameters. First, from the scales of τ we can see that the contact rates are quite heterogeneous. The highest frequency of contact is within the 45-64 group and the 75+ group, which we argue could be because the 45-64 aged people are still carrying out some economical activities during lock down that are necessary to keep the city running, while for younger people, there are more proportion of office workers who

medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint **(which was not certified by peer review)** is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

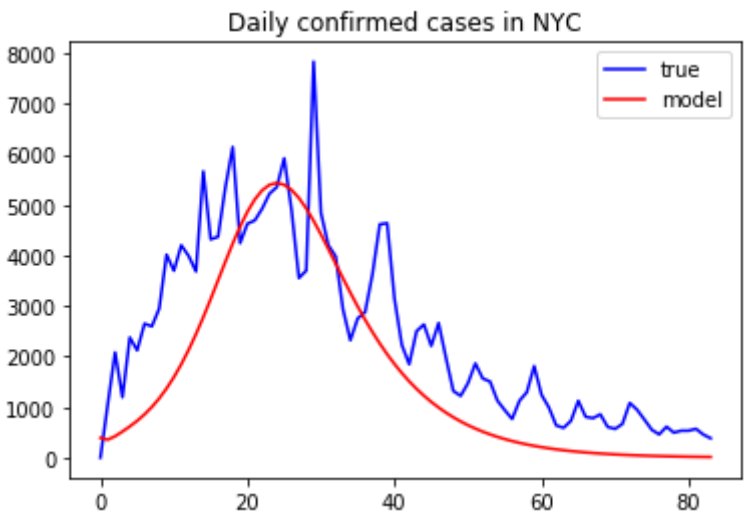

Figure 2: Predicted confirmed cases and actual data

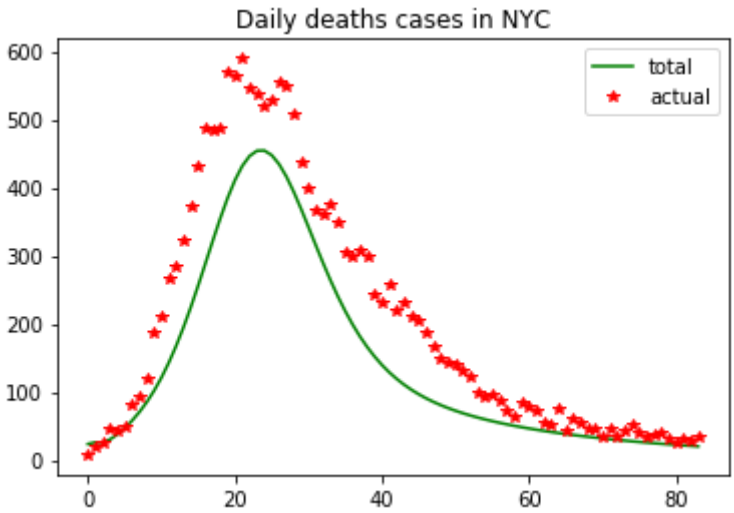

Figure 3: Predicted death cases and actual data

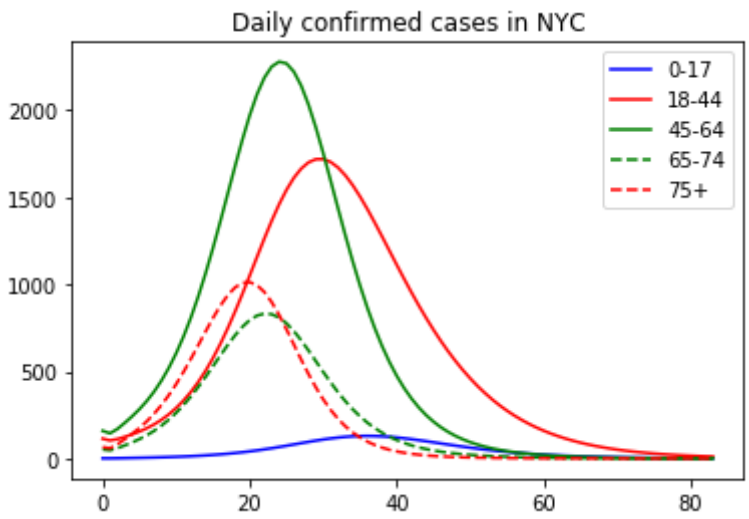

Figure 4: Predicted confirmed cases in each age group

medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

could possibly work remotely and have smaller frequency of contact with others. For the older group, the 65-74 aged may have a larger proportion of living alone, while the 75+ group may live in nursing homes and contact more often with others. Second, the ascertainment rate for different age groups also varies significantly, with a overall ascertainment rate of around 11%. For the kids, it is actually very small, which partially agrees with the recent news (see, e.g., CBS News 2020) that many children are tested positive on the virus, and would be otherwise not found infected if not tested. This raises concern that children may also be impacted by the virus. For other groups, it generally follows the order of age, where the older, the more likely to be ascertained. This observation is also consistent with our knowledge of the coronavirus, that it is more dangerous to older people.

4 Vaccine Allocation Policy

In this section, we evaluate several vaccine allocation policies using our age-structured SAPHIRE model and the parameters estimated. For simplicity, we assume that the effect of one dose of vaccine moves one susceptible individual to the removed compartment. We also assume that a test is performed prior to vaccination, so we don't give vaccine to infected but unascertained people.

4.1 Static Allocation

We consider static allocation polices, i.e., the amount of vaccines allocated for each group does not change over time. We assume that if all the susceptible individuals in a group are vaccinated and there are doses allocated to this group in the future, then the vaccines cannot be transferred to other groups and are wasted. In the following, we derive optimal static allocation policies with respect to different objectives (total confirmed cases, the total deaths, and their weighted summation) under different daily available doses.

We first consider minimizing the total number of deaths across all groups. We assume the daily number of available doses take values from 2,500 to 15,000. The total population of our interest is around 1.58 million, so if we give 15,000 doses of vaccine per day for 84 days, we should be able to cover around 80% of the total population. In this case, the vaccine allocation is shown in Table 4. From this table, we can see that when the daily available vaccines are very limited (e.g., 2500 doses), the optimal static allocation policy only focuses on the 75+ age

medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

group since they are the most vulnerable to the virus and the most likely to endure serious consequences. As the available vaccines increase, the policy allocates some doses for younger groups as well. Although these groups are more resistant to severe disease outcomes, providing vaccination to them can help protect the most vulnerable group as well, as they have a large contact rate with the oldest group. Meanwhile, for people aged between 0-44, the policy suggests not to provide vaccination even if the daily available doses increase to 15000. This is because the fatality rate among these groups are extremely low, and their contact level with the most vulnerable groups are relatively small, so in order to achieve the minimum deaths, we prefer not to give vaccination to them.

doses	0-17	18-44	45-64	65-74	75+
2500	0	0	0	0	2500
5000	0	0	0	1184	3816
7500	0	0	0	2507	4993
10000	0	0	620	3322	6057
12500	0	0	2041	3779	6680
15000	0	0	3506	4102	7392

Table 4: Daily vaccine allocation for minimizing deaths

We also look at the allocation amount of each age group scaled by their population. Table 5 shows the daily number of vaccine per 10,000 people, allocated to each age group. From this table we can observe that even though the absolute number of vaccines per capita is increased, the scale of increase for younger groups are not that significant as indicated by the previous table. The oldest group still has the largest vaccines per capita when the daily supply of vaccines increases.

doses	0-17	18-44	45-64	65-74	75+
2500	0	0	0	0	246
5000	0	0	0	93	376
7500	0	0	0	198	492
10000	0	0	16	262	597
12500	0	0	53	298	658
15000	0	0	90	323	728

Table 5: Daily vaccine allocation for minimizing deaths, per 10,000 people

Figure 5 shows the estimated number of deaths as a function of total available doses. The horizontal axis represents the daily available doses, the vertical axis represents the minimal deaths by implementing the optimal static allocation policy. We can see from this figure that the decrease of deaths becomes less significant when the daily available vaccines becomes large.

medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

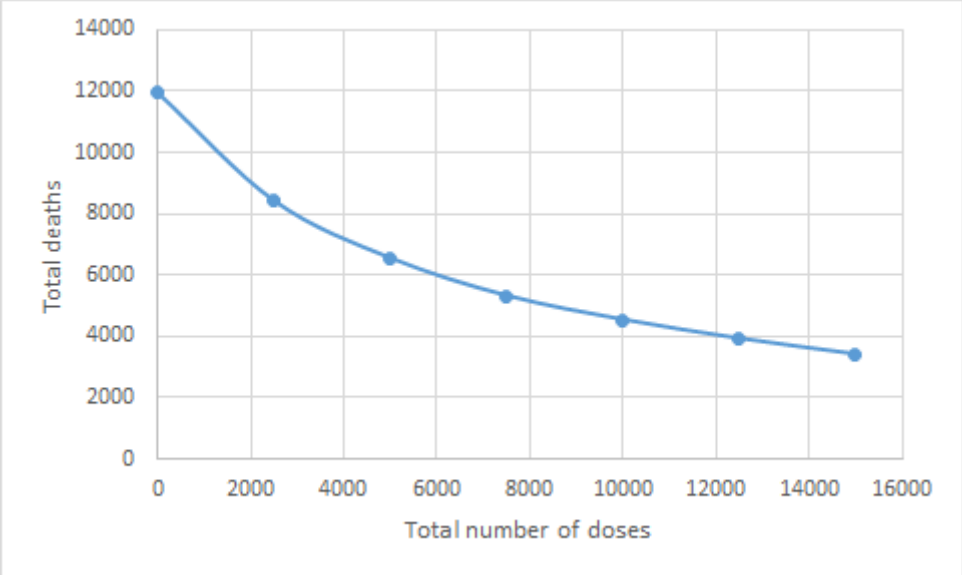

Figure 5: Estimated number of deaths vs. available doses

We now consider the optimal static policy by minimizing the total confirmed cases from all age groups. Table 6 shows the optimal vaccine allocation to each age group for different daily available doses. We observe from this table that the pattern is similar but Table 4 allocates more to older groups. Besides, the optimal static policy allocates vaccines to younger groups even when the supply is limited. This illustrates that when supply is limited, allocates all the available vaccines to the oldest group has a smaller marginal effect.

doses	0-17	18-44	45-64	65-74	75+
2500	0	0	384	619	1498
5000	0	0	1880	1037	2083
7500	0	695	2892	1435	2479
10000	0	1758	3616	1641	2986
12500	0	2838	4344	1936	3383
15000	0	3909	5143	2177	3771

Table 6: Daily vaccine allocation for minimizing confirmed cases

Again, an allocation per capita and the effect of total doses on total confirmed cases are provided in Table 7 and Figure 6 respectively. We can see that to better prevent the virus from spread in the population, the oldest group is given less vaccine compared to the previous allocation policy, while other groups are given more to decrease the spread within them. The effect of available doses on the total number of confirmed cases are almost linear, meaning that the incremental doses has an equal margin. This indicates that in order to prevent spread of the virus, it is better to provide as many doses of vaccine as possible, as the increased doses always have a nearly constant marginal effect.

medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

doses	0-17	18-44	45-64	65-74	75+
2500	0	0	10	49	148
5000	0	0	48	82	205
7500	0	11	74	113	244
10000	0	29	93	129	294
12500	0	47	112	153	333
15000	0	64	132	172	372

Table 7: Daily vaccine allocation for minimizing confirmed cases, per 10,000 people

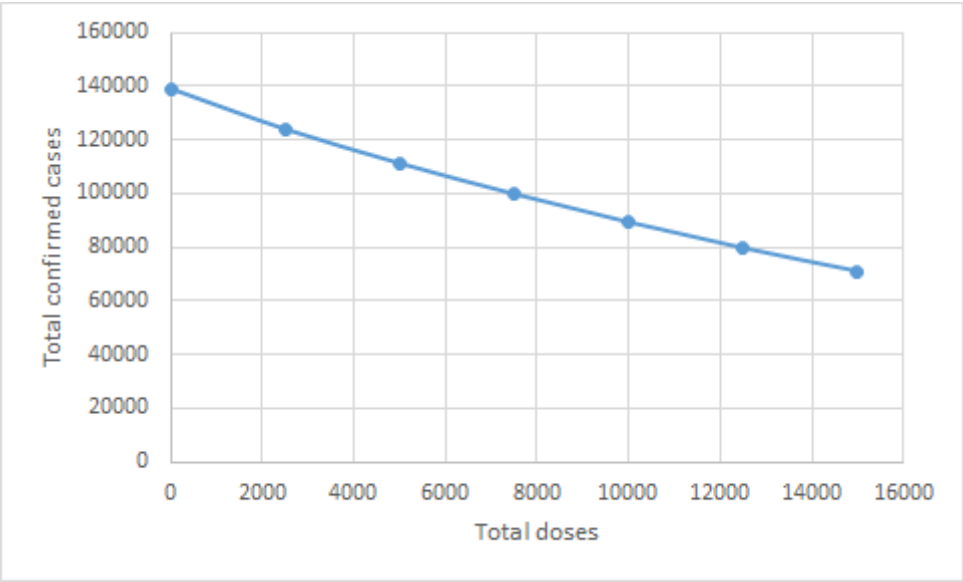

Figure 6: Estimated number of confirmed cases vs. available doses

medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

The decision maker may not simply want to minimize the total number of confirmed cases or the total deaths. Instead, it is more likely to set an objective which combines these two measures. So, here we consider an objective which is defined as:

$$weight\ of\ death \times deaths + confirmed\ cases$$

In this case, we fix the total number of daily available doses to be 10,000. Then we increase the weight of death from 1 to 40 to see how the policy is changed. Table 8 shows the optimal allocation under different weight of deaths. We can observe that when death gets more and more weighted, the policy converges slowly to the allocation which minimize only death and allocate more and more vaccine to the old people. We can see that as we increase the weight of death in our objective, the predicted confirmed cases increase almost linearly, but the decrease of total death becomes very flat quickly. So when making allocation decisions, the policy maker should properly balance the two objectives, and sometimes the over emphasise of deaths will significantly increase the total confirmed cases, which is not desired as we aim to eventually stop the virus from spreading.

death weight	0-17	18-44	45-64	65-74	75+	total confirmed	total death
1	0	1445	3382	1786	3386	89339	5619
5	0	0	3407	2451	4142	90993	4915
10	0	0	2757	2687	4555	92027	4756
15	0	0	2277	2719	5004	93196	4659
20	0	0	2111	2884	5006	93523	4641
25	0	0	1947	2893	5159	94019	4620
30	0	0	1711	2907	5383	94779	4592
35	0	0	1586	2914	5500	95198	4579
40	0	0	1495	3004	5501	95431	4573

Table 8: Daily vaccine allocation for minimizing combined objective and the outcomes

To illustrate the performance of our static policy, we also run two heuristic policies as benchmarks. In the benchmark test, we set the daily available doses to be 10,000 as the results would not change qualitatively if the amount of daily available doses is different. The first benchmark is a uniform allocation policy which gives 2,000 doses for each group per day. Note that the two oldest groups actually have less population, so this benchmark is indeed in favor of the older groups. The total number of confirmed cases under this policy is 95,171, while the total deaths is 7051, which is far inferior than what achieved by the previous policies. The optimal static policy for death minimization has an estimated 98,678 confirmed cases but only 4,541

medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

deaths. The one for confirmed cases minimization achieves 89,183 confirmed cases and 5,942 deaths. Another benchmark is a policy that allocates vaccines to each age group proportionally to its population. This policy has an estimated confirmed case of 100,215 and 9,432 deaths, which is even worse.

In the following, we relax the assumption that the amount of daily available vaccines does not change over time. It is likely that the supply is increasing at the beginning and after some time periods does not change to much. To model this, we assume the amount of daily available vaccines C_t has the following form:

$$C_t = \begin{cases} \gamma t, & t \leq C/\gamma \\ C, & t > C/\gamma. \end{cases}$$

In the simulation study, we set $C = 10,000$ and γ take value between 100 and 600. We consider static policies that allocate a constant percentage of available doses each day to each age group. Under different values of γ, the allocation to minimize death cases is shown in Table 9, while the allocation to minimize confirmed cases is in Table 10. We can observe from Table 9 that when the supply is limited at the beginning, i.e., γ is small, the majority of supplies is allocated to the oldest group. When the supply at the beginning increase, more vaccines are allocated to the group 64-74, but no vaccine is allocated to 0-64 aged individuals. This is almost consistent with the pattern of the optimal static policy when the daily supply is constant. When the objective is minimizing confirmed cases, we observe from Table 10 that groups 18-44 and 45-64 get more percentages of the supply, and their shares does not change too much when γ varies. This illustrates that when the steady supply C is moderately limited (i.e., 10000 in our case), the optimal static policy allocates more to younger groups regardless the beginning limited supplies.

γ	0-17	18-44	45-64	65-74	75+	death	confirmed
100	0.00%	0.00%	0.00%	6.17%	93.83%	10101	131321
200	0.00%	0.00%	0.00%	17.42%	82.58%	8997	125761
300	0.00%	0.00%	0.00%	19.31%	80.69%	8173	121689
400	0.00%	0.00%	0.00%	26.39%	73.61%	7536	117932
500	0.00%	0.00%	0.00%	28.93%	71.07%	7039	115186
600	0.00%	0.00%	0.00%	29.61%	70.39%	6648	113128

Table 9: Percentage allocation of minimizing deaths

We mention that under the the assumption that superfluous vaccines allocated to a group

medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

γ	0-17	18-44	45-64	65-74	75+	death	confirmed
100	0.00%	31.53%	27.50%	15.36%	25.62%	10946	129580
200	0.00%	37.00%	29.80%	12.42%	20.79%	10390	121859
300	0.00%	36.21%	32.63%	12.07%	19.11%	9876	115438
400	0.00%	31.87%	33.50%	13.56%	21.12%	9238	110503
500	0.00%	28.74%	34.49%	13.95%	22.87%	8712	106860
600	0.00%	27.24%	35.58%	15.24%	21.99%	8402	104181

Table 10: Percentage allocation of minimizing total confirmed cases

is wasted, the optimal static policy will cause a huge waste even when the supply is limited. For example, when the daily supply is a constant 2500, the total vaccines allocated to group 75+ is $84 \times 2500 = 210000$, which is more than twice of its population 101120. This suggests considering dynamic allocation policies which are examined in the next section.

4.2 Dynamic Allocation

We consider dynamic allocation policies in this section. Since our age-structured SAPHIRE model has 35 age-compartment pairs and the population in each pair may range to several thousands (some ranges to hundreds of thousands, e.g., susceptible compartment), it is challenging to compute the optimal dynamic allocation using dynamic programming. One may want to use approximate dynamic programming to obtain good heuristics. However, as illustrated in Long et al. (2018) with a space-structured epidemic model for the 2014 Ebola outbreak, the heuristic obtained from approximate dynamic programming performs worse than simple heuristics such as myopic policy (to be discussed below for our model). Hence, we only provide evaluations of several dynamic allocation heuristics.

- (**Old-First Policy**) This policy allocates available daily vaccines to age group 5 first. If there are any remaining vaccines, it allocates to age group 4, and so on.
- (**Infection-First Policy**) This policy allocates available daily vaccines proportionally to the infection ratio of each group (i.e., $\frac{A_i+E_i+P_i+I_i}{N_i}$).
- (**Myopic Policy**) In each time t, the myopic policy determines the amount of vaccines v_i allocated to age group i by minimizing the total new infections of all groups in time $t+1$,

medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

i.e.,

$$\begin{aligned} \min \sum_{i=1}^{5} (S_i - v_i)\beta \sum_{j=1}^{5} \frac{\tau_{ji}(I_j + \alpha(P_j + A_j))}{N_j} \\ \text{s.t.} \ \sum_{i=1}^{5} v_i \leq C_t, \\ 0 \leq v_i \leq S_i, \ i = 1, ..., 5. \end{aligned}$$

Since this is a linear program with a capacity constraint and a box constraint, the myopic policy will first allocate vaccines to the group with the largest coefficient. If the number of susceptible individuals in this group is less than C_t, then the myopic policy allocates the remaining doses to the group with the second largest coefficient, and so on. The coefficient of v_i in the objective can be regarded as the marginal effect of a unit vaccine allocated to group i.

- (**Death-Weighted Myopic Policy**) In each time t, the death-weighted myopic policy determines the amount of vaccines v_i allocated to age group i by minimizing the total weighted infections of all groups in time $t+1$, i.e.,

$$\begin{aligned} \min \sum_{i=1}^{5} w_i(S_i - v_i)\beta \sum_{j=1}^{5} \frac{\tau_{ji}(I_j + \alpha(P_j + A_j))}{N_j} \\ \text{s.t.} \ \sum_{i=1}^{5} v_i \leq C_t, \\ 0 \leq v_i \leq S_i, \ i = 1, ..., 5, \end{aligned}$$

where w_i is the death rate of group i.

- (**Two-Day Myopic Policy**) In each time t, this policy determines $v_i(t), v_i(t+1)$, $i =$

medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

$1, ..., 5$ to minimize the new infections of all groups in time $t+2$, i.e.,

$$\begin{aligned}
&\min \sum_{i=1}^{5} (E_i(t+2) - E_i(t) + \frac{E_i(t+1)}{D_e}) \\
&\text{s.t. } (1) - (7), \\
&\quad \sum_{i=1}^{5} v_i(t) \leq C_t, \\
&\quad \sum_{i=1}^{5} v_i(t+1) \leq C_{t+1}, \\
&\quad 0 \leq v_i(t) \leq S_i(t), \\
&\quad 0 \leq v_i(t+1) \leq S_i(t+1), \ i = 1, ..., 5.
\end{aligned}$$

In the numerical study, we set the daily supply C_t to be a constant 10,000. For the two-day myopic policy, as we mentioned in Section 3, the disease dynamics (1)-(7) is nonlinear. Hence, its optimization problem is a nonconvex problem, which potentially has local optimal solutions. To obtain a good local optimal solution, we solve the optimization problem 100 times at the beginning of every two days and pick the one that gives the minimal objective value. The outcomes of the above heuristics are listed in Table 11. Here, we do not put the outcome of the

	confirmed						death					
	0-17	18-44	45-64	65-74	75+	total	0-17	18-44	45-64	65-74	75+	total
old-first	3638	42787	35937	5709	1898	89969	0	359	1792	954	676	3781
myopic	2119	35435	21487	6985	887	66913	0	295	1069	1134	334	2832
infection-first	2285	31254	29341	5576	4315	72770	0	263	1450	923	1482	4118
two-day myopic	2121	35382	21735	6696	883	66818	0	295	1081	1090	332	2798

Table 11: Outcomes of different dynamic heuristics

death-weighted myopic policy as it yields the same allocation with the old-first policy in this numerical study. Table 11 shows that the two-day myopic policy performs the best in terms of both the total confirmed cases and the total deaths. The myopic policy has a similar outcome. The old-first policy has the largest total confirmed cases and the infection-first policy has the largest total death cases. This is because the old-first policy postpones younger groups that have the most infections, and the infection-first policy only consider the infection ratio $\frac{A_i+E_i+P_i+I_i}{N_i}$. Although the old-first policy itself seems to be beneficial to decrease death cases, it actually causes more deaths compared to the myopic policy the and two-day myopic policy which do not

medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

explicitly minimize deaths in their optimization problems. This illustrates that the decreasing of confirmed cases leads to fewer deaths as well and a good allocation can decrease both deaths and confirmed cases.

In the following ,we compare the allocation of different dynamic policies. Since the old-first and the myopic policy only allocate the daily supply to one group (except when the remaining susceptible individuals of the group cannot consume all the supply), we can visualize them as in Figure 7. In the figure, the different color blocks represent time periods of the allocation to different groups, and the vertical axis represents the time. The first row corresponds to the old-first policy, and the second row corresponds to the myopic policy. We do not specify the mixing of allocations on the boundary of different color blocks. The coefficients of v_i for different age groups under the myopic policy is shown in Figure 8. Note that even the group 75+ has the largest coefficient from day 10 (starts to allocate to group 65-74) to 35, the vaccines will not be given to group 75+ after day 10 as all its susceptible individuals are vaccinated. The allocation of the two-day myopic policy is shown in Figure 9.

We can observe that the execution of the myopic policy follows a similar pattern as the old-first policy at the beginning in that they both start from the oldest age group and gradually move to younger groups. There are some differences between these two policies that explain superiority of the myopic policy. First, the myopic policy spends less time in the 65-74 and 45-64 age group. We can observe that the myopic policy do have slightly more confirmed and deaths for the 65-74 group. Although the myopic policy spends less time in these groups, it moves to other groups (0-17, 18-44) with larger vaccine marginal effect (see Figure 8), and helps controlling the disease transmission in all groups. Second, the old-first policy still allocates vaccines to group 45-64 after day 34 even though the vaccine marginal effect of this group is smaller than group 0-17, while the myopic policy skips the 18-44 age group at day 34 and goes directly to the 0-17 age group. At the end of the horizon, the myopic policy fluctuates among the groups 65-74, 18-44 and 45-64 as their vaccine marginal effects are nearly the same and a small periods of allocation will change the vaccine marginal effect. At this stage, the allocation is insensitive to different groups and a change of the myopic policy will not hurt the performance.

The two-day myopic policy has a mixing allocation as shown in Figure 9. The majority of the supply is first given to the group 75+ and moves to the group 65-74 roughly on day 10. Around day 15 the majority of the supply is given to 45-64 and then moves to the group 0-17 around day 34. Then the policy gives most of the vaccines to the group 18-44 after day

medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

52. Comparing with the allocation of the myopic policy (Figure 7), the two-day myopic policy roughly follows the pattern of the myopic policy but has more mixing on the exchange time boundary and the end of the horizon, which may be the reason for its slightly better outcomes.

5 Conclusion

This paper considers vaccine allocation policies of Covid-19 to different age groups under limited supply. We use an age-structured SAPHIRE model and estimate relevant parameters with the epidemic data from NYC. Base on this model and the estimated parameters, we evaluate the performance of the optimal static policies under different daily vaccine supply and different objectives, and those for several dynamic allocation heuristics. Our numerical study shows that generally the optimal static policy allocates most of the vaccines to older groups when the objective is minimizing deaths, and if the objective is minimizing confirmed cases, then younger groups will get more. This suggests the decision maker to balance very carefully between different objectives which leads to a policy performs much better than ad-hoc allocation policies. The dynamic allocation heuristics in general perform better than the static ones. This is not surprising since we allow different allocation each day. Among the dynamic policies, the best are myopic policies and two-day myopic policies, and they performs much better than the other heuristics.

Our paper though has the following limitations. One is on the assumption placed on static policies, i.e., once the people in one group have all been vaccinated, the doses allocated to that group cannot be transferred to other groups and are wasted. This is not the case in reality, and it would be interesting to see how the static policy performs when the transfer of vaccines to other groups is allowed. Alternatively, a more complicated policy, for instance, piecewise static policies can be used to reduce the waste of vaccines under this assumption. Another limitation is that we do not solve the optimal dynamic policy. This is due to the curse of dimensionality and the nonlinearity of the disease dynamics. Finally, although our numerical study only uses data from NYC, we believe the insights of allocation policies draw from it could provide reference to decision makers on the allocation of the upcoming Covid-19 vaccines.

medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint **(which was not certified by peer review)** is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

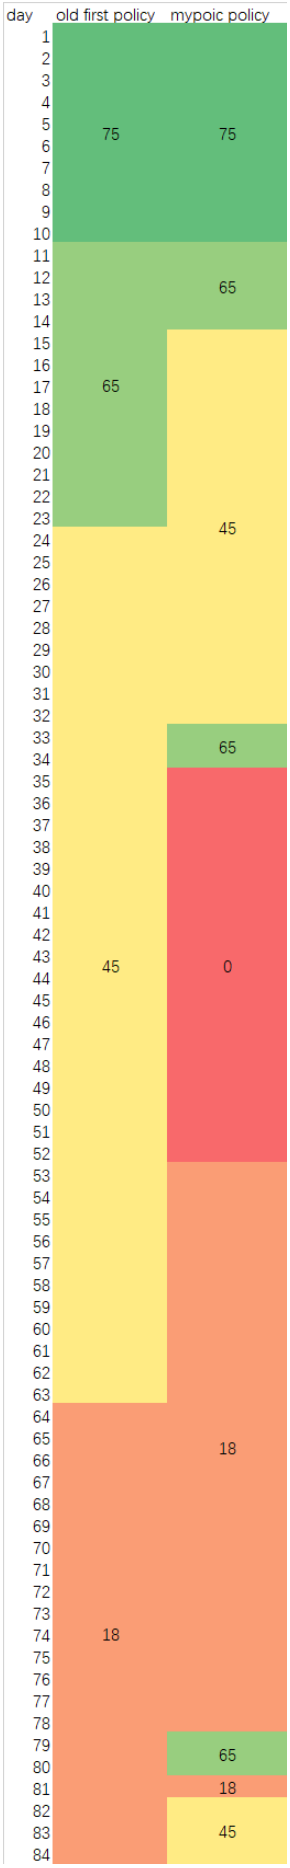

Figure 7: Allocation for old-first policy and myopic policy

medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

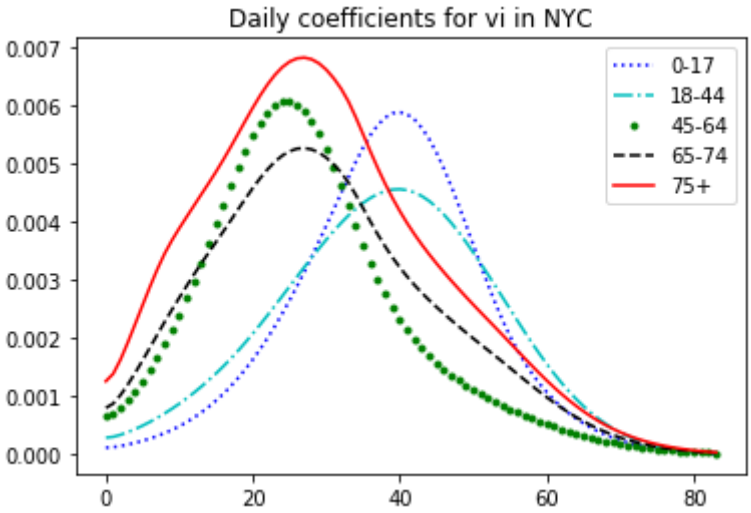

Figure 8: Coefficients of v_i for different age groups under myopic policy

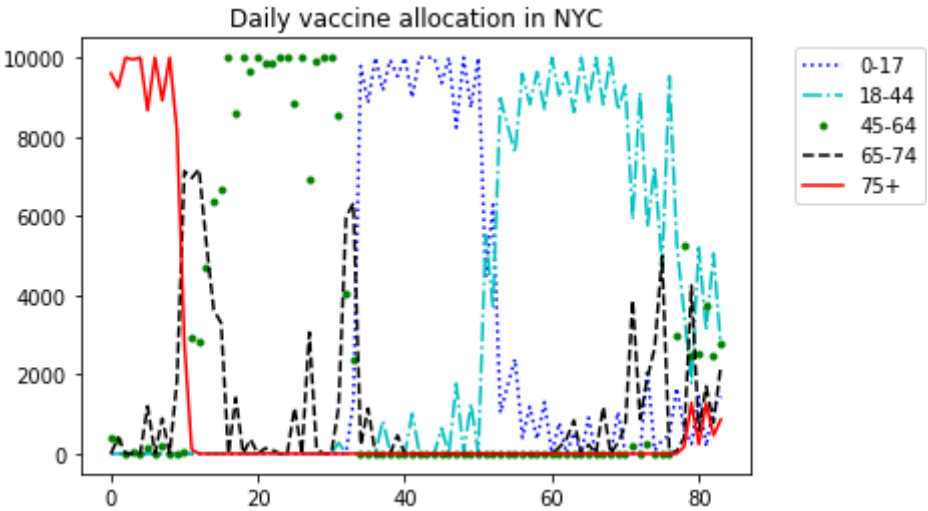

Figure 9: Daily allocation under the two-day myopic policy

medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

6 Appendix: Details of Parameter Estimation

In this appendix, we provide details of the parameter estimation method.

The input to our SAPHIRE model is simply the parameters and the observed initial conditions. Then the model can compute the daily number of each groups from previous days' data with the ODEs defined in Section 2. The initial condition comes from the observation on Mar 17, 2020. We set the initial susceptible population to be the total population, compute the initial number of ascertained people from each group by the total infected number times the overall ratio of different age groups, and set the initial removed population to be 0. Note that the number of some disease group is not observable, for example, the unascertained infected people. However, we argue that as early as Mar 17, there is not too much people in this group, so a slight error in this number will not cause a huge difference in our estimation. With our final set of parameters, if we multiply or divide the initial number in these unobserved compartments by 2, the change in our estimated total number of infected cases would be around 1.5%. To further deal with this, we set the first ten days in our model to be the burn-in period, and we will only take predictions of our model from Mar 27 to June 8.

As we have mentioned before, to estimate the parameters, our loss function to be minimized is the squared error of the prediction of confirmed cases and deaths in each groups. Formally, let $x(t)$ be the state vector at time t satisfying ODEs $\frac{dx(t)}{dt} = f_t(x(t), \theta)$ and $x^o(t)$ be the observable state subvector. Let $y(t)$ be the observable data in time t. The loss function is written in (8). This is a common measure used in the literature, see, for example, Cantó et al. (2017). Due to the highly non-convex nature of the ODEs that define the epidemic dynamics, the loss function possess a great number of local optimal solutions, many of which have similar level of mean squared error in the prediction. These local optimal solutions can be obtained from applying optimization algorithm to the loss function from different starting points.

$$\min_{\theta \in \Theta} \sum_{t=0}^{T} ||y(t) - x^o(t)||^2 \tag{8}$$

$$\text{s.t. } \frac{dx(t)}{dt} = f_t(x(t), \theta), \; x(0) = x, \tag{9}$$

A critical part of our estimation is to select the local optimal solution that best characterizes the nature of transmission of the virus. When selecting among the multiple estimation of contact

medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

matrix τ between age groups, and the ascertainment rate r of different age groups, we consider the following constraints on the estimation.

First, the contact between age groups should be balanced and bi-directional. In other words, when two person makes contact, it is possible to transfer virus from the first person to the second person, or the reverse. In terms of contact between age groups, the following equation must hold for all elements in τ

$$\tau_{ij} \times \text{population in group } i = \tau_{ji} \times \text{population in group } j$$

Second, there are some studies in finding the contact rate between people using socioeconomic data, for example, Fumanelli et al. (2012) uses the data of school, workplace and community to estimate the contact rate between age groups in the UK. Refer to Figure 2 of the paper for a heatmap of the estimated contact level between age groups. Though as a metropolitan area, people in New York City during the lock down period may not have exactly the same contact level as shown in this paper, but some insights should be similar. For example, the diagonal elements of the matrix τ should be the largest in a row, and the more off-diagonal an elements is in a row, the smaller it is. Formally, we incorporate the following constraint, with a small tolerance.

$$\tau_{i1} \leq \tau_{i2} \leq \cdots \leq \tau_{ii} \geq \cdots \geq \tau_{i5}$$

Third, as data have shown, the hospitalization rate and fatality rate of the older group is significantly higher than the younger group, we expect the ascertainment rate r for the groups follow the same fashion. In this sense, we incorporate the constraint that the ascertainment rate should be increasing in age. Formally,

$$r_1 \leq r_2 \leq r_3 \leq r_4 \leq r_5.$$

Based on the above criteria, we find the local minimum solution that satisfy the above constraints, that yields the smallest objective value. The results are shown in the main text.

medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

References

Baruch College (2017). New york city age and sex distribution. https://www.baruch.cuny.edu/nycdata/population-geography/age_distribution.htm. Accessed 20-August-2020.

Birge, J. R., Candogan, O., and Feng, Y. (2020). Controlling epidemic spread: Reducing economic losses with targeted closures. Available at SSRN https://papers.ssrn.com/abstract=3590621.

Cantó, B., Coll, C., and Sánchez, E. (2017). Estimation of parameters in a structured sir model. *Advances in Difference Equations*, 2017(1):33.

CBS News (2020). 97,000 children reportedly test positive for coronavirus in two weeks as schools gear up for instruction. https://www.cbsnews.com/news/covid-19-kids-school-children-positive-tests-coronavirus-reopening/. Accessed 20-August-2020.

Corum, J., Grady, D., Wee, S.-L., and Zimmer, C. (2020). Coronavirus Vaccine Tracker. *The New York Times*. https://www.nytimes.com/interactive/2020/science/coronavirus-vaccine-tracker.html. Accessed 20-August-2020.

Fumanelli, L., Ajelli, M., Manfredi, P., Vespignani, A., and Merler, S. (2012). Inferring the structure of social contacts from demographic data in the analysis of infectious diseases spread. *PLoS Comput Biol*, 8(9):e1002673.

Gallagher, J. (2020). How close to developing a vaccine are we? *BBC News*. https://www.bbc.com/news/health-51665497. Accessed 20-August-2020.

Hamer, W. H. (1906). Epidemic disease in England — the evidence of variability and of persistence. *The Lancet*, 167:733–738.

Hao, X., Cheng, S., Wu, D., Wu, T., Lin, X., and Wang, C. (2020). Full-spectrum dynamics of the coronavirus disease outbreak in Wuhan, China: A modeling study of 32,583 laboratory-confirmed cases. medRxiv 2020.04.27.20078436.

Li, R., Pei, S., Chen, B., Song, Y., Zhang, T., Yang, W., and Shaman, J. (2020). Substantial undocumented infection facilitates the rapid dissemination of novel coronavirus (SARS-CoV-2). *Science*, 368(6490):489–493.

medRxiv preprint doi: https://doi.org/10.1101/2020.08.23.20179820; this version posted August 25, 2020. The copyright holder for this preprint **(which was not certified by peer review)** is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC-ND 4.0 International license .

Liu, Y., Gu, Z., Xia, S., Shi, B., Zhou, X.-N., Shi, Y., and Liu, J. (2020). What are the underlying transmission patterns of COVID-19 outbreak? An age-specific social contact characterization. *EClinicalMedicine*, 22:100354.

Loftus, P. (2020). Early coronavirus vaccine supplies likely won't be enough for everyone at high risk. *Wall Street Journal*. https://www.wsj.com/articles/early-coronavirus-vaccine-supplies-likely-wont-be-enough-for-everyone-at-high-risk-11596706202. Accessed 20-August-2020.

Long, E. F., Nohdurft, E., and Spinler, S. (2018). Spatial resource allocation for emerging epidemics: A comparison of greedy, myopic, and dynamic policies. *Manufacturing & Service Operations Management*, 20(2):181–198.

NYC Health (2020). COVID-19: Data Main - NYC Health. https://www1.nyc.gov/site/doh/covid/covid-19-data.page. Accessed 20-August-2020.

Sun, L. H. (2020). Who should get a coronavirus vaccine first? *Washington Post*. https://www.washingtonpost.com/health/2020/07/29/covid-vaccine-essential-workers-high-risk-populations/. Accessed 20-August-2020.

Wack, J. S. H. a. C. (2020). Pfizer, BioNTech Get $1.95 Billion Covid-19 Vaccine Order From U.S. Government. *Wall Street Journal*. https://www.wsj.com/articles/pfizer-biontech-get-1-95-billion-covid-19-vaccine-order-from-u-s-government-11595418221. Accessed 20-August-2020.

Wikipedia contributors (2020). Covid-19 pandemic in new york city — Wikipedia, the free encyclopedia. https://en.wikipedia.org/w/index.php?title=COVID-19_pandemic_in_New_York_City&oldid=973893024. Accessed 20-August-2020.

Zhou, L., Wang, Y., Xiao, Y., and Li, M. Y. (2019). Global dynamics of a discrete age-structured SIR epidemic model with applications to measles vaccination strategies. *Mathematical Biosciences*, 308:27–37.