2022 IEEE 61st Conference on Decision and Control (CDC)
December 6-9, 2022. Cancún, Mexico

Optimal Capacity-Constrained COVID-19 Vaccination for Heterogeneous Populations

Raghu Arghal, Shirin Saeedi Bidokhti, and Saswati Sarkar

2022 IEEE 61st Conference on Decision and Control (CDC) | 978-1-6654-6761-2/22/$31.00 ©2022 IEEE | DOI: 10.1109/CDC51059.2022.9992682

Abstract— **COVID-19 and the ensuing vaccine capacity constraints have emphasized the importance of proper prioritization during vaccine rollout. This problem is complicated by heterogeneity in risk levels, contact rates, and network topology which can dramatically and unintuitively change the efficacy of vaccination and must be taken into account when allocating resources. This paper proposes a general model to capture a wide array of network heterogeneity while maintaining computational tractability and formulates vaccine prioritization as an optimal control problem. Pontryagin's Maximum Principle is used to derive properties of optimal, potentially highly dynamic, allocation policies, providing significant reductions in the set of candidate policies. Extensive numerical simulations of COVID-19 vaccination are used to corroborate these findings and further illicit optimal policy characteristics and the effects of various system, disease, and population parameters.**

I. INTRODUCTION

Since its beginning in December 2019, the COVID-19 pandemic has resulted in nearly 500 million infections and over 6 million deaths as of March 2022 [1]. Vaccines have proven to be the most effective countermeasure to the pandemic by limiting further transmission and protecting especially vulnerable populations [2]. During its early stages, the vaccination drive was heavily capacity constrained with demand far outstripping supply and administration capability – a challenge that continues to plague Low- and Middle-Income Countries (LMICs) [3]. This is bound to be the case for vaccines developed for every infectious disease. Under such constraints, governments and public health organization must make the critical choice of whom to vaccinate first: 1) those who are likely to transmit the disease most, 2) those who are at risk for developing a serious form of the disease due to age or comorbidity, or 3) a combination of the first and second set. For COVID-19 most public health bodies opted for the second category first, but was it the optimal choice even if we consider the limited objective of minimizing say only the fatality count?

To appreciate the complications in resolving this decision process consider the example scenario of a retirement community, which comprises of two categories of individuals: 1) residents and 2) employees who serve the residents (eg. essential service providers). The residents have an increased risk of developing a serious form of the disease due to age, while employees who are younger usually suffer from mild symptoms even when infected, but transmit the disease to a large number of individuals due to their contacts with large and dynamic sets. The residents usually come in contact with the employees regularly and with other residents infrequently. Thus, the residents have a high contact rate with the employees and a low contact rate among each other. Suppose our goal is to minimize the fatalities and only the residents are at risk of succumbing to the disease.

The question now is whether the optimal vaccination policy first vaccinates the residents, the employees, or resorts to a potentially complex combination of the two extremes. The answer is far from clear even under the above simplifying assumptions. For example, it is entirely conceivable that the optimal strategy will first vaccinate the residents or that it may vaccinate the employees first, particularly when 1) the number of employees is much smaller than the number of residents, 2) contact rates between the employees and residents is much higher than that between residents, or 3) vaccination capacity is low. In this scenario, it will take a long time to vaccinate a substantial number of residents if the decision is to vaccinate the residents first, meanwhile if some employees imbibe the disease they can spread the disease to a large number of residents who are yet to be vaccinated leading to a large death count. In contrast, the small number of employees can be vaccinated in a short time, thus the disease can now spread among the residents only through direct contacts between them which happen infrequently. Thus, the disease spreads slowly, allowing enough time for the residents to be vaccinated before a substantial fraction among them incurs the disease. The optimal policy may also in principle be a complex temporal combination of the two extremes. We therefore need a systematic methodology to determine the optimal strategy which is the focus of this paper.

The challenge in determining the optimal vaccination strategy is multi-fold. First, the populations are naturally heterogeneous with different individuals exhibiting different social contact patterns and risk factors. Yet, considering networks where each individual is a separate entity usually leads us to analytical intractability and numerical infeasibility because of the curse of dimensionality. In contrast, homogeneous abstractions, though both numerically and analytically tractable, lose the essence of the system. We resolve this dilemma by grouping different sections of the populace in accordance with their risk factors and contact rates. Each group is referred to as a cluster, and the individuals in the same cluster are assumed to be statistically identical in terms of risk factors and contact rates among each other and across clusters. The contact rates between clusters depend on pairs in question. This clustered modeling provides a tunable tradeoff between retaining analytical and numerical tractabil-

The authors are with the Department of Electrical and Systems Engineering at the University of Pennsylvania, Philadelphia PA 19104.

ity and capturing the inherent heterogeneity. Vaccination in different clusters however remains coupled due to capacity constraints on overall rate of vaccination across all clusters together.

We subsequently formulate the determination of the optimal vaccination strategy that attains a global objective, eg, minimizing the fatality count, as an optimal control problem. The formulation accommodates arbitrary time-varying combinations of vaccination strategies that potentially simultaneously vaccinate individuals from different clusters subject to the capacity constraints at rates that can constitute complex functions of time. Despite the rich decision space, we prove that there exist optimal vaccination strategies which devote the entire vaccination capacity to one group at a time, except when the available capacity exceeds the number of individuals who are yet to be vaccinated in some groups, in which case, the spare capacity can be directed to other groups. These optimal vaccination strategies are *mixed* only during such limited durations. This result implies that complicated decision strategies involving arbitrary mixture of vaccination rate allocation across groups need not be considered; such policies can not usually be implemented any way.

Ideally, we would like the optimal vaccination strategy to be even simpler, e.g. one that starts with a group, vaccinates it in entirety and then moves on to another group. That is, a bang bang strategy with the number of jumps equaling the number of groups less 1. This is because only policies with a limited number of switches between groups are likely to be implemented in practice. Here we first obtain a negative result which shows that the optimal policy may have a greater number of switches. Subsequently, through an extensive numerical investigation we show that the negative result arises only for limited ranges of parameters, mostly there are no extra switches. Even when the number of switches exceeds the number of groups the difference is small. For our numerical results we focus on the simplest scenario that captures the heterogeneities that arise in practice. Specifically we focus on a scenario in which there are three distinct types of individuals: high risk (65+ yo, immunocompromised, etc.), high contact (essential workers, healthcare workers, etc.), and baseline.[1] These categories allow us to account for the disparate effects of COVID on people with underlying conditions [4] and individuals with high centrality and dynamic contacts who may act as super-spreaders [5]. This is especially important for settings like retirement homes where high risk individuals are overrepresented or LMICs which exhibit dramatically different contact networks and age structures [6].

The existing literature and contributions of our work are detailed in Section II. We then introduce our model and formulate our optimal control problem (Section III). In Section IV we employ PMP to determine structural properties of optimal policies, followed by numerical results (Section V) to corroborate and contextualize our findings. For brevity, additional numerical results and extensions can be found in [7].

[1]The group with high contacts and high risk is omitted because it is a relatively small demographic, and eliminates the need to tradeoff between prevention and protection.

II. RELATED WORKS

There is a long history of mathematical modelling applied to epidemiology and the control of epidemics [8], [9], [10], [11]. These efforts can broadly be categorized as either graph-based models or homogenized models [12]. Graph-based models have the flexibility to consider arbitrary topologies [13] and individual heterogeneity [14], but suffer from computational intractability. In fact, optimization of control measures over arbitrary graphs is provably NP-hard [15], leading to a focus on approximation algorithms [16], [17].

As a result, homogenized models, most notably the SIR compartmental model [18] and its variations [19], [20], are often used to study the control of epidemics. These models can yield elegant, interesting, and practical applications (e.g. [21]), but often lack the heterogeneity and flexibility to properly address the multitude of factors that affect the spread and control of infectious diseases. Augmented compartmental models, often incorporating disease stages or age stratification by expanding the state space [22], [23], incorporate particularly impactful forms of heterogeneity to chart a middle ground between graph-based models and full homogenization.

We build upon the framework proposed in [24], imbuing the compartmental model with generalized clusters of agents of different types, allowing us to capture heterogeneity in contact rates, risk factors, and network topology. We add a constraint on the number of vaccines applied at any time to account for shortages in vaccine supply or administration capacity. This constitutes a mixed path constraint on both the state and control [25] with a significant effect on optimal control structure and system dynamics [26], [27]. The combination of the mixed path constraint and the heterogeneity of our model allows us to answer the questions of vaccine prioritization, the crux of our modelling and analytical contributions.

Our application of this model to COVID-19 vaccination also constitutes a significant contribution. There have been a number of works studying COVID-19 vaccine allocation strategies [28]. [29] considers a few age-based vaccination policies in an augmented compartmental model. [30] considers a similar stratification while optimizing over initial static vaccine allocations. A number of other works have considered factors such as essential workers or geography, most still relying on simulation or a harsh restriction of the set of possible vaccine policies [31], [32], [33], [34]. By application of optimal control and PMP, we optimize over a much broader set of feasible policies while incorporating a more generalized framework for heterogeneity and clustering.

III. VACCINE PRIORITIZATION MODEL

Throughout this paper, we make use of the results of [35], allowing us to describe the evolution of a pandemic

via a system of ordinary differential equations. We will first describe a simplified SIRD model of a pandemic in a heterogeneous population. This will prove useful in deriving theoretical results before employing an expanded state space to accurately depict stages of COVID-19 infection for our numerical results in Section V. It should be noted that the simplified model maintains the essential components that dictate the structure of optimal policies broadly. The addition of intermediate states and asymptomatic infection does not significantly affect these findings.

A. System Dynamics

Assume a network of N individuals partitioned into three types of size N_1, N_2, N_3 corresponding to high risk, baseline, and high contact groups, respectively. Each individual exists in one of four states: susceptible (S), infected (I), recovered (R), and deceased (D).

A susceptible individual of type i becomes infected upon interacting with an infected person of type j at rate $\beta_{ij}/N = \beta_{ji}/N$ with $\beta_{1j} = \beta_{2j} \leq \beta_{3j}/C$ where C is the minimum ratio between high contact rates and baseline. The length of the infected period is an exponential random variable with mean $\frac{1}{\gamma_i}$ after which a person either dies with probability m_i or recovers with probability $(1 - m_i)$. It is assumed that $m_3 = m_2 < m_1$ to delineate the high risk group. Define $M = \frac{m_1}{m_2}$ to be the mortality ratio between the risk groups. Susceptible individuals of type i are also vaccinated at a rate of u_i, leading them directly to the recovered state. The maximum instantaneous vaccination rate is V_0. It is assumed that once in the recovered state reinfection is not possible and that the total number of individuals of each type remains constant. We denote the fraction of individuals who are susceptible, infected, recovered, and dead at time t by $S_i(t)$, $I_i(t)$, $R_i(t)$, and $D_i(t)$, respectively[2].

This engenders the following set of ordinary differential equation describing the system:

$$\dot{S}_i = -S_i \sum_{j=1}^{3} \beta_{ji} I_j - S_i u_i \tag{1a}$$

$$\dot{I}_i = S_i \sum_{j=1}^{3} \beta_{ji} I_j - I_i \gamma_i \tag{1b}$$

$$\dot{R}_i = S_i u_i + I_i \gamma_i (1 - m_i) \tag{1c}$$

$$\dot{D}_i = I_i \gamma_i m_i \tag{1d}$$

We will use $V(t) = \sum_{i=1}^{3} S_i u_i$ to denote the vaccine capacity used at time t. Note that, while we focus on three clusters corresponding to high contact, high risk, and baseline groups, this framework (as well as Thm. 1) generalizes to any number of clusters with arbitrary contact rates and risk parameters. Partitioning could more generally be based on geography, risk factors, contact rates, or contact tracing data. It can also be augmented with more stages of disease progression as in Section V.

[2]The time argument is often omitted for clarity.

B. Objective and Optimal Control Formulation

We adopt overall mortality as our objective function (including a terminal cost for remaining infected individuals) and assume a time horizon $T > \frac{S(0)}{V_0}$ so that there is sufficient time to vaccinate all susceptibles. This yields the following optimal control formulation

$$\begin{aligned} &\text{minimize } \sum_{i=1}^{3} D_i(T) + m_i I_i(T) \\ &\text{subject to (1)} \\ &\qquad V(t) \leq V_0 \quad \forall t \in [0, T] \\ &\qquad x(0) = x_0, \ x_0 \succeq 0 \end{aligned} \tag{2}$$

where x refers to the full state vector and x_0 its corresponding initial condition.

Remark 1: The system described by (1) with nonnegative initial conditions has a unique state solution which satisfies the initial condition and state constraints of (2) ([24], Theorem 1). This allows us to drop the state constraints in further considerations.

IV. OPTIMAL VACCINE PRIORITIZATION

Solving (2) directly is intractable as it would require optimizing over uncountably many potential control policies. Instead, we use PMP to obtain necessary conditions for an optimal control from which we discern structural properties [36], [37].

We first formulate the Hamiltonian and Lagrangian as follows:

$$\mathcal{H} := \sum_{i=1}^{3} (\lambda_i^S \dot{S}_i + \lambda_i^I \dot{I}_i - I_i \gamma_i m_i) \tag{3}$$

$$\mathcal{L} := \mathcal{H} - \mu \left(\sum_{i=1}^{3} S_i u_i - V_0 \right) \tag{4}$$

where the λ costate functions are absolutely continuous and satisfy

$$\dot{\lambda}_i^S = -\frac{\partial \mathcal{L}}{\partial S_i} \qquad \dot{\lambda}_i^I = -\frac{\partial \mathcal{L}}{\partial I_i} \tag{5}$$

$$0 = \lambda_i^S(T) \qquad -m_i = \lambda_i^I(T) \tag{6}$$

and μ is an integrable function satisfying

$$\mu \left(\sum_{i=1}^{3} S_i u_i - V_0 \right) = 0, \quad \mu(t) \geq 0 \ a.e. \tag{7}$$

$$n_i(t) = -\lambda_i^S S_i - \mu S_i \ a.e. \text{ where } n_i(t) \in N_{[0,1]}(u_i^*(t)) \tag{8}$$

where $N_{[0,1]}$ denotes the normal cone to $[0, 1]$.

This leads immediately to the following property of the costate variables which will prove useful while discerning the structure of optimal vaccination policies.

Lemma 1: For all $t \in [0, T)$ and $i = 1, \ldots, M$, $\lambda_i^I < 0$, $\lambda_i^S \leq 0$ and $\lambda_i^I - \lambda_i^S \leq 0$.

Intuitively, the costate variables can be thought of as shadow costs associated with their respective states (see [37]

Sec. 3.3.4). In this context, Lemma 1 states that marginal increases in the number of susceptible or infected individuals increase the objective function and the associated increase is greater for infected individuals.

PMP then states that $u^* \in \text{argmax}\ \mathcal{H}$ where u^* refers to an optimal control function. Plugging in expressions for $\mathcal{H}$ and (1a)-(1d), we obtain:

$$u^*(t) \in \text{argmax}_{u \in \mathcal{U}} \sum_{i=1}^{3} -\lambda_i^S S_i u_i \tag{9}$$

where $\mathcal{U} = \{v : 0 \preceq v \preceq 1, \sum_{i=1}^{3} S_i v_i \leq V_0\}$ denotes the set of admissible controls.

Note that (9) provides a large and convenient reduction in the set of potentially optimal vaccination policies which must be considered summarized in the following theorem.

Theorem 1: Suppose without loss of generality that $-\lambda_a^S S_a > -\lambda_b^S S_b > -\lambda_c^S S_c$ at time t where $a, b, c \in \{1, 2, 3\}$. Then there exists an optimal control which takes the following form:

$$u_a^*(t) = \min\left(1, \frac{V_0}{S_a(t)}\right)$$
$$u_b^*(t) = \min\left(1, \frac{V_0 - u_a^*(t)S_a(t)}{S_b(t)}\right)$$
$$u_c^*(t) = \min\left(1, \frac{V_0 - u_a^*(t)S_a(t) - u_b^*(t)S_b(t)}{S_c(t)}\right)$$

Further, if t_0 is the first time at which $S_i(t) < V_0$ for some i. Then

$$u_i^*(t) \in \left\{\frac{V_0}{S_i(t)}, 0\right\}\ \forall t < t_0$$

i.e. until t_0 the optimal controller devotes all vaccine capacity to one group at a time

Put simply, as much vaccine as possible is allocated to the highest priority group with the remainder going to the next highest priority until no vaccine capacity remains or no susceptible individuals remain.

To understand the importance and impact of the vaccine capacity constraint, it is useful to compare the result above to an unconstrained analogous result.

Lemma 2: If the vaccine capacity constraint is removed and replaced by a vaccine cost function of the form $\sum_{i=1}^{3} h_i(u_i)$ with h_i concave, then the optimal policy is of the following form for some $t_0 \geq 0$

$$u_i^*(t) = \begin{cases} 1 & t < t_0 \\ 0 & t \geq t_0 \end{cases}$$

If h_i is strictly convex, the transition between $u_i^* = 1$ and $u_i^* = 0$ is continuous and monotonic.

In this case, all clusters that can benefit from vaccination are vaccinated at full capacity. Without the mixed path constraint, the planner is not forced to prioritize any one group over another.

Although Thm. 1 does not provide as drastic a simplification of the optimal policy, it still retains the bang-bang-like structure due to the linearity of Hamiltonian in the control variable. By allowing us to ruling out mixed policies which split vaccine capacity between multiple groups, this yields a significant reduction in the set of potential optimal policies.

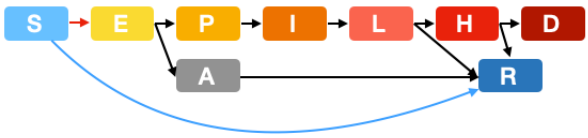

Fig. 1: **State diagram of disease progression.** Red arrow indicates exposure to an infectious individual, black arrows denote natural disease progression, and blue arrow denotes vaccination.

This reduction is also of practical importance: policies which focus on one group at a time are more easily implemented and align with the common practice of phased vaccine rollout to target groups. To this end, a particularly implementable class of optimal policies would be those corresponding to an ordering of the three types in which a group is fully vaccinated before moving on to vaccinate another. In such a policy, the number of times the highest priority vaccine target changes is exactly 2.

Such a structure is not guaranteed by Thm. 1, and, in fact, there are scenarios in which such policies are not optimal, however they are relatively uncommon (see Section V).

V. NUMERICAL RESULTS

In this section we detail our numerical investigations of both structural properties of the optimal policy as well as the impact of parameters on the optimal and implications for public health protocol.[3] To more accurately depict the progression and spread of COVID, we expand our state space from the simplified model presented in Section III to the richer one presented in Figure 1. Our state space now includes susceptible (S), exposed (E), presymptomatic (P), asymptomatic (A), infected (I), late-stage infected (L), hospitalized (H), recovered (R), and deceased (D). Intuitively, our findings for the model described in Section III generalize as the additional states (E, P, A, L, H) can be collapsed into the infected state. More specifically, this system retains the linearity of the Hamiltonian in the control variable.

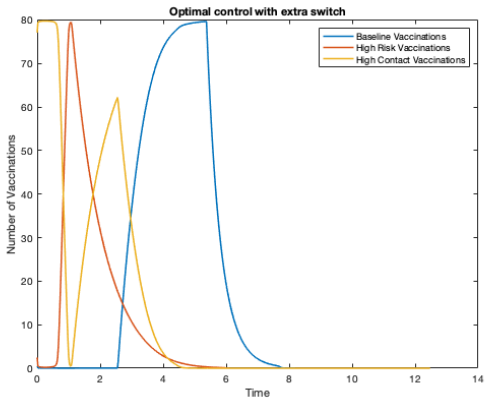

Fig. 2: Optimal policy with extra switch

[3]We use Yop [38] and CasADi [39] to obtain numerical results.

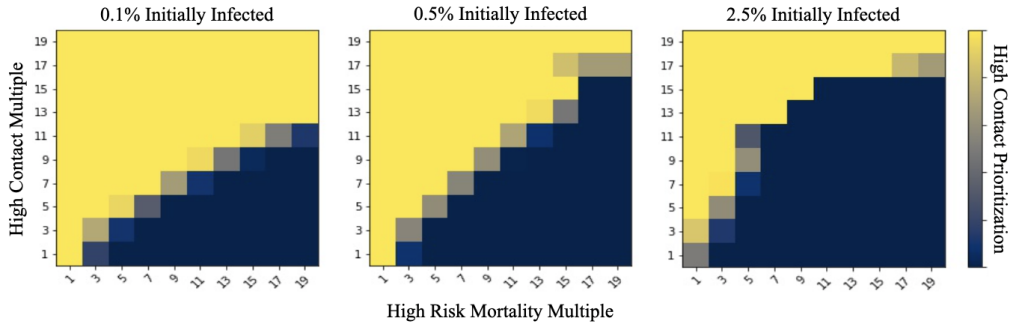

Fig. 3: Effects of initial infection prevalence on optimal policy

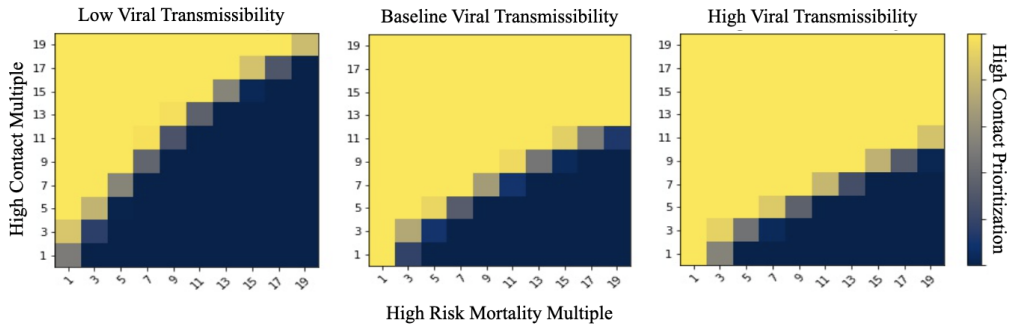

Fig. 4: Effects of baseline contact rate on optimal policy

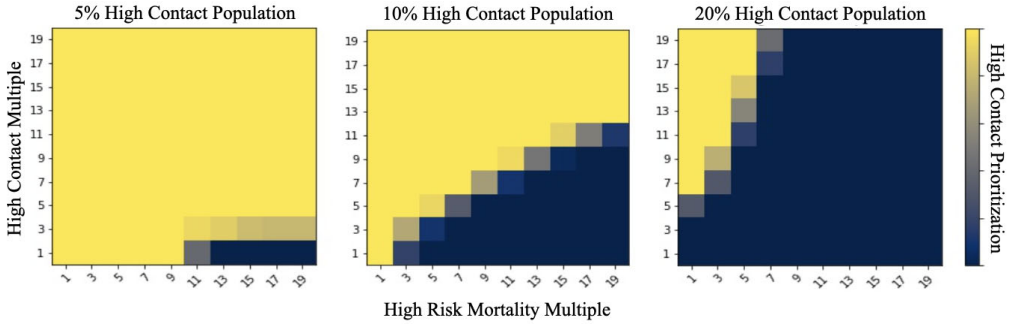

Fig. 5: Effects of high contact population size on optimal policy

A. *Structural Properties of Optimal Vaccination*

We solved the optimal control problem over a fine grid on contact rates, mortality levels, initial state, disease characteristics, population demography, and vaccine capacity, representing feasible ranges for various localities and COVID variants [2], [4], [6]. In all solutions, the baseline was vaccinated last as expected. Any fixed parameters, such as transition rates between disease states, were set based on best estimates from the Centers for Disease Control and Prevention [40].

In 94.3% of instances, the optimal solution found was a simple bang-bang policy with exactly 2 switches; however a minority of problems return solutions with an extra switch as depicted in Fig. 2 where the optimal policy switches from high contact vaccination to high risk before returning to the high contact. In these few cases, the closest performing simple policy incurred only 1.3% more deaths than the optimal. Based on Thm. 1, extra switches are indicative of two groups with similar vaccine priority, thus even when the optimal is not a simple policy, the reduction in deaths is marginal. This lends justification to the simple, phased vaccine rollout adopted by most countries for COVID – any marginal gain of a more complicated policy may even be countered by the increased logistical difficulty of implementation.

B. *Impact of Initial State, Disease, and Network Parameters*

Now we focus our attention on choosing between two simple bang-bang policies (high contact, high risk, then baseline or high risk, high contact, then baseline) and analyze the effect of certain key parameters and their broader implications.

1) Initial Infection Prevalence (see Fig. 3): An increase in initial infection prevalence decreases the prioritization of high contact individuals in optimal vaccination policies. Intuitively, the primary goal of vaccinating high contact individuals is to limit the proliferation of the pandemic. Thus, if the infection is already more widespread, the marginal benefit of vaccinating super-spreaders is lower.

Conversely, in the stylized example of a retirement community from Section I, if initial infections are low, vaccinating employees could mitigate the spread of the disease sufficiently early to curtail mortality among vulnerable residents.

2) Viral Transmissibility (see Fig. 4):
As viral transmissibility increases, the impact of superspreaders is exaggerated and, as such, the prioritization of high contact individuals increases. The effect of viral transmissibility is crucial as new variants emerge with disparate transmission characteristics. More or less virulent strains can call for dramatically different public health responses.

3) High Contact Population Size (see Fig. 5):
High contact prioritization decreases with increasing size of the high contact group. This result seems to conflict with Thm. 1 where the relative priority of a group is $-\lambda_i^S S_i$. Again it is useful to consider the ultimate goal of high contact vaccinations – to limit spread. As the high contact group grows, it is less possible to quickly vaccinate the group and cut off the vectors for disease spread. Thus, the crucial factor is the time taken to fully vaccinate the group and the disease spread that can happen in the interim.

In the case of LMICs, where contact rates even for more senior populations tend to be higher [6], this could require more targeted vaccination of individuals with particularly detrimental comorbidities.

VI. CONCLUSION

We present a generalized epidemic model that balances numerical tractability with the heterogeneity required to reason about vaccine prioritization. By casting vaccine prioritization as an optimal control problem and applying PMP, we prove that there exist optimal policies which are not mixed, instead focusing on one group at a time. Although we show that, generally, the optimal policy may include extra switches between priority groups, our extensive numerical results show that such regimes are uncommon and simple policies are near-optimal even in these cases. We also illicit the effects of several important parameters on the structure of the optimal policy and their implications for vaccine allocation.

REFERENCES

[1] *WHO Coronavirus (COVID-19) Dashboard*. URL: https://covid19.who.int/info/.

[2] Gang Lv et al. "Mortality rate and characteristics of deaths following COVID-19 vaccination". In: *Frontiers in Medicine* 8 (2021).

[3] Edouard Mathieu et al. "A global database of COVID-19 vaccinations". In: *Nature human behaviour* 5.7 (2021), pp. 947–953.

[4] Amir Emami et al. "The role of comorbidities on mortality of COVID-19 in patients with diabetes". In: *Obesity Medicine* 25 (2021), p. 100352.

[5] Ana Paula Schmitz Rambo et al. "Impact of superspreaders on COVID-19: systematic review". In: *Sao Paulo Medical Journal* 139 (2021), pp. 163–169.

[6] Andria Mousa et al. "Social contact patterns and implications for infectiousdisease transmission – A systematic review and meta-analysis of ContactSurveys". In: *eLife* 10 (2021). DOI: 10.7554/elife.70294.

[7] Raghu Arghal, Shirin Saeedi Bidokhti, and Saswati Sarkar. "Optimal Capacity-Constrained COVID-19 Vaccination for Heterogeneous Populations". In: *arXiv* ().

[8] Klaus Dietz and JAP Heesterbeek. "Bernoulli was ahead of modern epidemiology". In: *Nature* 408.6812 (2000), pp. 513–514.

[9] William O Kermack and Anderson G McKendrick. "Contributions to the mathematical theory of epidemics—II. The problem of endemicity". In: *Bulletin of mathematical biology* 53.1-2 (1991), pp. 57–87.

[10] Roy M Anderson and Robert M May. *Infectious diseases of humans*. en. Oxford science publications. London, England: Oxford University Press, Aug. 1992.

[11] Suresh P. Sethi and Preston W. Staats. "Optimal Control of Some Simple Deterministic Epidemic Models". In: *The Journal of the Operational Research Society* 29.2 (1978), pp. 129–136. ISSN: 01605682, 14769360. URL: http://www.jstor.org/stable/3009792.

[12] Cameron Nowzari, Victor M. Preciado, and George J. Pappas. "Analysis and Control of Epidemics: A Survey of Spreading Processes on Complex Networks". In: *IEEE Control Systems Magazine* 36.1 (2016), pp. 26–46. DOI: 10.1109/MCS.2015.2495000.

[13] Zoltán Dezs ő and Albert-László Barabási. "Halting viruses in scale-free networks". In: *Phys. Rev. E* 65 (5 May 2002), p. 055103. DOI: 10.1103/PhysRevE.65.055103. URL: https://link.aps.org/doi/10.1103/PhysRevE.65.055103.

[14] Romualdo Pastor-Satorras and Alessandro Vespignani. "Immunization of complex networks". en. In: *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* 65.3 Pt 2A (Mar. 2002), p. 036104.

[15] James Aspnes, Kevin Chang, and Aleksandr Yampolskiy. "Inoculation Strategies for Victims of Viruses and the Sum-of-Squares Partition Problem". In: *J. Comput. Syst. Sci.* 72.6 (Sept. 2006), pp. 1077–1093. DOI: 10.1016/j.jcss.2006.02.003.

[16] Po-An Chen, Mary David, and David Kempe. "Better Vaccination Strategies for Better People". In: *Proceedings of the 11th ACM Conference on Electronic Commerce*. EC '10. Cambridge, Massachusetts, USA: Association for Computing Machinery, 2010, pp. 179–188. ISBN: 9781605588223. DOI: 10.1145/1807342.1807370. URL: https://doi.org/10.1145/1807342.1807370.

[17] James Aspnes, Kevin Chang, and Aleksandr Yampolskiy. "Inoculation strategies for victims of viruses and the sum-of-squares partition problem". In: *Journal of*

Computer and System Sciences 72.6 (2006), pp. 1077–1093.

[18] Ronald Ross. "An application of the theory of probabilities to the study of a priori pathometry.Part I". In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 92.638 (1916), pp. 204–230.

[19] Boris Shulgin, Lewi Stone, and Zvia Agur. "Pulse vaccination strategy in the SIR epidemic model". In: *Bulletin of mathematical biology* 60.6 (1998), pp. 1123–1148.

[20] N.K. Gupta and R.E. Rink. "Optimum control of epidemics". In: *Mathematical Biosciences* 18.3 (1973), pp. 383–396. ISSN: 0025-5564. DOI: https://doi.org/10.1016/0025-5564(73)90012-6. URL: https://www.sciencedirect.com/science/article/pii/0025556473900126.

[21] M. H. R. Khouzani, Eitan Altman, and Saswati Sarkar. "Optimal Quarantining of Wireless Malware Through Reception Gain Control". In: *IEEE Transactions on Automatic Control* 57.1 (2012), pp. 49–61. DOI: 10.1109/TAC.2011.2150350.

[22] Stephane Verguet et al. "Controlling measles using supplemental immunization activities: a mathematical model to inform optimal policy". In: *Vaccine* 33.10 (2015), pp. 1291–1296.

[23] Fadoua Balabdaoui and Dirk Mohr. "Age-stratified discrete compartment model of the COVID-19 epidemic with application to Switzerland". In: *Scientific reports* 10.1 (2020), pp. 1–12.

[24] Soheil Eshghi et al. "Optimal Patching in Clustered Malware Epidemics". In: *IEEE/ACM Transactions on Networking* 24.1 (2016), pp. 283–298. DOI: 10.1109/TNET.2014.2364034.

[25] Francis Clarke and M do R De Pinho. "Optimal control problems with mixed constraints". In: *SIAM Journal on Control and Optimization* 48.7 (2010), pp. 4500–4524.

[26] Rachael Miller Neilan and Suzanne Lenhart. "An Introduction to Optimal Control with an Application in Disease Modeling." In: *Modeling paradigms and analysis of disease trasmission models*. 2010, pp. 67–81.

[27] Md Haider Ali Biswas, Luıs Tiago Paiva, and MDR De Pinho. "A SEIR model for control of infectious diseases with constraints". In: *Mathematical Biosciences & Engineering* 11.4 (2014), p. 761.

[28] Nuru Saadi et al. "Models of COVID-19 vaccine prioritisation: A systematic literature search and Narrative Review". In: *BMC Medicine* 19.1 (2021). DOI: 10.1186/s12916-021-02190-3.

[29] Kate M Bubar et al. "Model-informed COVID-19 vaccine prioritization strategies by age and serostatus". In: *Science* 371.6532 (2021), pp. 916–921.

[30] Laura Matrajt et al. "Vaccine optimization for COVID-19: Who to vaccinate first?" In: *Science Advances* 7.6 (2021), eabf1374.

[31] Jack H Buckner, Gerardo Chowell, and Michael R Springborn. "Dynamic prioritization of COVID-19 vaccines when social distancing is limited for essential workers". In: *Proceedings of the National Academy of Sciences* 118.16 (2021).

[32] Bjorn Goldenbogen et al. "Optimality in COVID-19 vaccination strategies determined by heterogeneity in human-human interaction networks". In: (2020).

[33] Shuli Zhou et al. "Optimizing spatial allocation of COVID-19 vaccine by agent-based spatiotemporal simulations". In: *GeoHealth* 5.6 (2021), e2021GH000427.

[34] Joseph Chadi Lemaitre et al. "Optimizing the spatiotemporal allocation of COVID-19 vaccines: Italy as a case study". In: *medRxiv* (2021).

[35] Thomas G. Kurtz. "Solutions of Ordinary Differential Equations as Limits of Pure Jump Markov Processes". In: *Journal of Applied Probability* 7.1 (1970), pp. 49–58. ISSN: 00219002. URL: http://www.jstor.org/stable/3212147.

[36] Lev Semenovich Pontryagin. *Mathematical theory of optimal processes*. CRC press, 1987.

[37] Dieter Grass et al. *Optimal control of nonlinear processes with applications in drugs,corruption, and terror*. Springer, 2010.

[38] Viktor Leek. *An optimal control toolbox for MATLAB based on CasADi*. 2016.

[39] Joel A E Andersson et al. "CasADi – A software framework for nonlinear optimization and optimal control". In: *Mathematical Programming Computation* 11.1 (2019), pp. 1–36. DOI: 10.1007/s12532-018-0139-4.

[40] *Covid-19 pandemic planning scenarios*. Mar. 2021. URL: https://www.cdc.gov/coronavirus/2019-ncov/hcp/planning-scenarios.html.