Qeios

Navigating the Herd Immunity Surface: A Novel Framework for Optimising Epidemic Response Strategies

Gabriele Oliva[1], **Sandro Bonfigli**[2,3,4], **Pierpaolo Cavallo**[5,3], **and Antonio Scala**[6,3,5,7*]

[1]University Campus Bio-Medico of Rome, via A. del Portillo 21, 00128 Rome, Italy
[2]Directorate General of Health Prevention, Ministry of Health, Rome, Italy
[3]Big Data in Health Society, Roma, Italy
[4]Global Health Security Alliance - GHSA Italy
[5]Department of Physics, University of Salerno, Fisciano, SA, Italy
[6]CNR-ISC, Applico Lab, Roma, Italy
[7]Centro Ricerche Enrico Fermi, Roma, Italy
[*]antonio.scala@cnr.it

ABSTRACT

Background: The typical reaction strategy to an epidemic involves the implementation of various pharmaceutical (e.g., vaccination) and non-pharmaceutical interventions (e.g., social distancing) to reach the so-called "herd immunity threshold," ensuring that new surges of the epidemic dampen out.
Aim: We introduce a novel concept: the "Herd Immunity Surface." Unlike traditional approaches, which focus on a single herd immunity threshold, our framework considers heterogeneous population classes, such as different age cohorts or geographical regions.
Methods: We demonstrate that multiple herd immunity thresholds can achieve equivalent epidemic-dampening outcomes, even when resulting from different strategies (e.g., uniform vaccination vs. prioritizing the elderly, generalized vs. selective lockdowns, etc.).
Results: This discovery opens the door to policy optimization, where the specific herd immunity threshold chosen becomes a strategic decision with profound economic, logistic, political, and ethical implications. Importantly, it facilitates informed decision-making regarding the selection of vaccination strategies, allowing for versatility in achieving effective epidemic control.
Conclusions: Our study introduces the concept of the Herd Immunity Surface, offering a novel framework that transcends traditional epidemic response strategies. By highlighting the redundancy in achieving herd immunity, our research provides a foundation for optimizing policy decisions, particularly in the context of vaccination strategies, with far-reaching implications for public health and policymaking.

Introduction

Mathematical and computational approaches are valuable tools for analyzing scenarios, making forecasts during epidemics, and addressing complex public health issues, including the development of effective control strategies (both pharmaceutical and not) during emergencies[1,2]. Such control strategies have gained new significance during the ongoing COVID-19 pandemic due to its global scale, particularly the challenge of fair and equitable vaccine distribution[3–5]. In many countries, heterogeneous vaccination policies have been implemented, prioritizing specific categories such as the elderly or exposed workers. However, the concept of heterogeneous vaccination policies has not been thoroughly investigated, except for some preliminary numerical analysis[6].

The concept of the "herd immunity threshold" (HIT) serves as the foundation for designing vaccination policies. The HIT represents the fraction r^* of immune individuals in the population required to dampen the spread of the epidemic. In this paper, we extend the HIT concept to a heterogeneous population, where the population is divided into distinct classes. We demonstrate that, rather than being a single point, the threshold becomes a surface $\mathscr{H}^*$ referred to as the "herd immunity surface" (HIS).

Each point on the HIS is a potential choice for a vaccination policy. Since each HIT corresponds to a different distribution of vaccines among the population classes, selecting the target HIT becomes a policy choice with economic, logistic, political, and ethical implications[7].

In this paper, by exploiting the possibility of reaching different HIT points on the HIS through non-pharmaceutical interventions, we show how such interventions can complement the shortage of vaccines, enabling epidemic control while

https://doi.org/10.32388/TKU022

vaccines are being produced and distributed. To illustrate the HIS concept effectively, we consider the SIR[8,9] and SEIR[10] models with heterogeneous population classes.

Methods

Epidemiological studies often rely on mathematical models to gain insights into the spread of diseases and assess the impact of control strategies. Two commonly employed modelling frameworks are the Susceptible-Infectious-Recovered (*SIR*) and Susceptible-Exposed-Infectious-Recovered (*SEIR*) models. These models provide valuable insights into the dynamics of epidemics by dividing the population into compartments representing different stages of infection. The simplicity of *SIR* models makes them a popular choice, while *SEIR* models are especially useful when accounting for significant incubation periods. It's worth noting that for both *SIR* and *SEIR* models, the 'Recovered' compartment can also account for individuals who have achieved immunity through a positive result from vaccination. In this section, we explore how these modelling approaches can lead to multiple choices when it comes to selecting a vaccination strategy. Specifically, we demonstrate that in a heterogeneous population, herd immunity may not be achieved at a single point but rather on a hyper-surface. We then discuss how to extend these results to other compartmental models.

In *SIR* models, the population is divided into compartments representing different stages of infection. Specifically, '*S*' stands for susceptible individuals (those who can contract the disease), '*I*' for infectious individuals (those who have contracted the disease and can transmit it to others), and '*R*' for recovered individuals that are now *R*esistant to being infected. To account for additional factors such as loss of immunity, births, deaths, or healthy carriers, more complex variations of the *SIR* models may introduce additional compartments[11].

Simple *SIR* models assume a uniform population and homogeneous mixing among infected and susceptible individuals. However, real-world epidemics often involve diverse or dispersed populations. SIR models can be extended to address this heterogeneity by dividing the population into different classes, as shown in Equation (1)[12,13]. These models require the estimation of contact matrices that quantify the frequency of interactions among different population classes. Such matrices have underscored the significance of social, demographic, and economic factors in influencing the actual spread of diseases among individuals[14].

It's worth noting that the results presented in this paper are also applicable to heterogeneous *SEIR* models. *SEIR* models are derived from the *SIR* framework and are particularly useful for cases with a substantial incubation period during which infected individuals are not yet infectious. This additional compartment denoted as E for exposed individuals is introduced to account for this delay. For the purposes of our discussion, we will collectively refer to "infected individuals" as either the I class in the *SIR* model or the combined compartments $E+I$ in the *SEIR* model.

In our notation, lowercase letters (i.e., s, i, e, r) represent the fractions of individuals within a specific class. For $S(E)IR$ models, recovered individuals are considered immune to the disease. Consequently, a key objective of vaccination strategies is to increase the fraction r of resistant individuals beyond the herd immunity threshold r^* (see Section Herd Immunity Surface and Next Generation Matrix).

We will represent vectors with bold symbols (i.e., $\boldsymbol{\nu}$ with components ν_k) and use uppercase letters for matrices (i.e., B as the matrix with elements B_{kl}) or significant quantities (i.e., A denoting the total number of affected, $\boldsymbol{N}$ as the vector where N_k represents the total number of individuals in class k). In terms of the fractions s_k, i_k, and r_k representing susceptible, infectious, and recovered individuals in each class k, the *SIR* model is described by a set of deterministic differential equations:

$$\begin{aligned}
\partial_t s_k &= -s_k \sum_l B_{kl} i_l, \\
\partial_t i_k &= s_k \sum_l B_{kl} i_l - \gamma_k i_k, \\
\partial_t r_k &= \gamma_k i_k.
\end{aligned} \tag{1}$$

Here, B is the *transmission matrix*, with B_{kl} being the rate at which a susceptible individual of class k encounters an infectious individual of class l and becomes infected, while γ_k represents the rate at which infectious individuals in class k are removed (i.e., recover) from the infection cycle.

In the *SEIR* model, the dynamics of the exposed fraction $\boldsymbol{e}$ is introduced:

$$\begin{aligned}
\partial_t s_k &= -s_k \sum_l B_{kl} i_l, \\
\partial_t e_k &= s_k \sum_l B_{kl} i_l - \mu_k e_k, \\
\partial_t i_k &= \mu_k e_k - \gamma_k i_k, \\
\partial_t r_k &= \gamma_k i_k.
\end{aligned} \tag{2}$$

In these equations, μ_k represents the rate at which exposed individuals of class k become infectious.

If we denote a_k as the fractions of *all* infected individuals in class k (i.e., $a_k = i_k$ in SIR and $a_k = i_k + e_k$ in SEIR), it follows from eq. (1) and eq. (2) that in both models:

$$\partial_t a_k = s_k \sum_l B_{kl} a_l - \gamma_k a_k. \tag{3}$$

Thus, indicating N_k as the number of individuals in the k^{th} class, the total number of infected $A = \sum N_k a_k$ evolves over time as:

$$\begin{aligned} \partial_t A &= \sum N_k \partial_t a_k \\ &= \sum_l N_l i_l \left(\sum_k \mathscr{B}_{lk} s_k - \gamma_l \right), \end{aligned} \tag{4}$$

where the matrix $\mathscr{B}_{lk} = N_k B_{kl} {N_l}^{-1}$, being by construction similar to B, is isospectral to B.

Herd Immunity Surface

Herd immunity is a condition whereby no epidemic outbreaks are possible, i.e., the number of infected individuals decreases over time. Since both $\boldsymbol{\gamma}$ and $\boldsymbol{i}$ are non-negative, in $S(E)IR$ models if $\sum_l \mathscr{B}_{kl} s_l - \gamma_k < 0$ (see eq. (4)), then the total number of infected individuals decreases (i.e., $\partial_t A < 0$) regardless of the fraction of infected individuals, i.e.:

$$\boldsymbol{\mathcal{R}s} < \mathbf{1} \tag{5}$$

where the *infection growth matrix* (IGM) $\boldsymbol{\mathcal{R}}$ has elements $\mathscr{R}_{kl} = \gamma_k^{-1} \mathscr{B}_{kl}$. Here, $\mathbf{1}$ is the vector with all components equal to 1 and we indicate with $\boldsymbol{u} < \boldsymbol{v}$ that $u_k < v_k$ for *all* classes.

Notice that, in the one-dimensional case (i.e., for a single class), denoting $\beta = \mathscr{B}_{11}$ and $\gamma = \gamma_1$, herd immunity is reached when $\mathscr{R}_0 s < 1$, where $\mathscr{R}_0 = \gamma^{-1}\beta$ is the so-called *basic reproduction number*, measuring the potential growth rate of an epidemic. Since $s \leq 1$, we recover the classical result that an epidemic does not spread unless $\mathscr{R}_0 > 1$. Moreover, in the one-dimensional case, a threshold $s^* = \mathscr{R}_0^{-1}$ separates the stable points with $s < s^*$ from the unstable points with $s > s^*$. Correspondingly, it is possible to define the so-called *herd immunity threshold* (HIT) as:

$$r^* = 1 - s^* = 1 - \mathscr{R}_0^{-1}, \tag{6}$$

i.e., the fraction of immune people (either due to acquired immunity or vaccination) above which no epidemic burst is possible.

In conclusion, the matrix $\boldsymbol{\mathcal{R}}$ plays the role of the basic reproduction number and defines the herd immunity surface (HIS) $\mathscr{H}^*$ as the boundary of the stable set $\mathscr{H}$, that is:

$$\mathscr{H} = \{\boldsymbol{s} : \boldsymbol{\mathcal{R}s} < \mathbf{1}\}. \tag{7}$$

In other words, $\mathscr{H}$ represents the set of states where no epidemic outbreaks can exist.

In Figure 1, we show the HIS (i.e., the green surface) for a three-class $S(E)IR$ model: all the points beyond the HIS are stable (i.e., the epidemic is dampened), while the unstable region (depicted in red) corresponds to the condition from which an epidemic can develop: clearly, instead of a single threshold, each individual point in the HIS represents a proper HIT and, consequently, represents a valid target for a vaccination policy.

Notice that such a formulation allows for a class-dependent definition of the HIS. Suppose that we have n classes, but we want to avoid epidemic outbreaks in a limited set of $m < n$ classes denoted as $\mathbb{C}$. The corresponding herd immunity surface $\mathscr{H}^*_{\mathbb{C}}$ is the boundary of the $\mathbb{C}$-stable set $\mathscr{H}_{\mathbb{C}}$, i.e.:

$$\mathscr{H}_{\mathbb{C}} = \{\boldsymbol{s} : (\boldsymbol{\mathcal{R}s})_k < 1 \; \forall k \in \mathbb{C}\}. \tag{8}$$

It has the property that $\mathbb{C} \subset \mathbb{D}$ implies that $\mathscr{H}_{\mathbb{D}} \subset \mathscr{H}_{\mathbb{C}}$. In Figure 2, we show the herd immunity surfaces for a $S(E)IR$ model with two classes: clearly, aiming for herd immunity in a single class enlarges the possible stable states.

Finally, for some systems, the contact matrix can assume a quasi-block-diagonal form[13]. As an example, if our classes represent geographical compartments like states or regions, the inter-compartment mobility (and hence the inter-compartment contacts) has been experimentally estimated to be of order $\varepsilon \sim 10^{-3}$, i.e., smaller than the mobility within the compartment[13]. Thus, assuming that we have m weakly interacting classes and denoting $\boldsymbol{v}^k$ as the vectors regarding the k^{th} class, and with $\boldsymbol{\mathcal{R}}^k$ as the diagonal blocks describing the interaction among the variables of the k^{th} class, the solution $\tilde{\boldsymbol{s}} = [\boldsymbol{s}^1, \cdots, \boldsymbol{s}^m]^T$ of the m equations:

$$\boldsymbol{\mathcal{R}}^k \boldsymbol{s}^k = \boldsymbol{u}^k \tag{9}$$

approximates the real solution $\boldsymbol{\mathcal{R}s} = \boldsymbol{u}$ to order ε. In other words, for such systems, $\|\tilde{\boldsymbol{s}} - \boldsymbol{s}\| = \mathscr{O}(\varepsilon)$, and the equations can be solved separately for each weakly interacting class, at the cost of introducing a small error only in the last significant digits.

Herd Immunity Surface and Next Generation Matrix

An alternative approach for defining the HIS could involve using the Next Generation Matrix (NGM) $\mathcal{K}(\boldsymbol{s})$ to describe the growth of the infected population in heterogeneous systems[11]. The NGM is a concept associated with the analysis of epidemic dynamics in the context of general compartmental models; it contains information about the rates of transition between compartments and is commonly used to predict the potential growth or decline of infectious individuals. The NGM helps determine whether an epidemic will grow or decline by analyzing the eigenvalues of this matrix and is instrumental in understanding the *state-dependent* basic reproduction number $\mathscr{R}_0(\boldsymbol{s})$ defined as the spectral radius[15]:

$$\rho(\mathcal{K}(\boldsymbol{s})) = \lim_{m\to\infty} \|\mathcal{K}^m(\boldsymbol{s})\|^{1/m}. \tag{10}$$

Thus, an alternative definition of the HIS would be to define $\mathscr{H}^*_{NGM}$ (the NGM-HIS) as the boundary of the NGM-stable set $\mathscr{H}_{NGM}$:

$$\mathscr{H}_{NGM} = \{\boldsymbol{s} : \rho(\mathcal{K}(\boldsymbol{s})) < 1\}. \tag{11}$$

However, for systems described by eqs. (1), (2), the NGM has elements $\mathscr{K}_{kl}(\boldsymbol{s}) = s_k \mathscr{R}_{lk}$. Thus, if a vector $\boldsymbol{s}$ is IGM stable (i.e., stable in the sense of eq. (5)), the column sums $\sum_k \mathscr{K}_{kl}(\boldsymbol{s}) = \sum_k \mathscr{R}_{lk} s_k$ are also bounded by 1 (see eq. 5). Since the spectral radius is upper-bounded by any matrix norm (e.g., see lemma 3.1.1 of ref.[16]), and since the maximum of the column sums (i.e., the one-norm) is bounded by one, it follows that $\boldsymbol{s}$ is also NGM stable (i.e., stable in the sense of eq. (11)). Thus, the NGM herd immunity surface $\mathscr{H}^*_{NGM}$, defined as the boundary of $\mathscr{H}_{NGM}$, encloses $\mathscr{H}$ since $\mathscr{H} \subset \mathscr{H}_{NGM}$; this implies that our IGM conditions are more stringent.

HIS as a Pareto Frontier

Due to the inherent trade-offs and optimization challenges involved in achieving herd immunity, the Herd Immunity Surface (HIS) can be conceptualized as a Pareto Frontier. The Pareto Frontier is a concept from economics that represents a situation where there are multiple solutions where no individual or group can be made better off without making someone else worse off.

In the context of vaccination and herd immunity, the HIS represents a spectrum of possibilities, each corresponding to a different combination of vaccination strategies and population coverage. Each point on the HIS reflects a specific trade-off between various dimensions, such as the percentage of the population vaccinated, logistical ease, cultural adaptation, and overall mortality. Thus, just like the Pareto Frontier, where improving the well-being of one individual may come at the expense of another, optimizing vaccination strategies on the HIS involves navigating complex trade-offs. Some strategies may excel in certain dimensions (e.g., minimizing logistical challenges) but fall short in others (e.g., achieving the desired herd immunity threshold). The Pareto Frontier nature of the HIS implies that there is no single, universally optimal solution; rather, the optimal strategy depends on the specific priorities and constraints at play.

Results

Obtaining precise human mixing patterns is essential for accurate disease modelling[14]. Contact matrices for various countries, as exemplified in the study by Mossong et al.[17], provide insights into the interactions between different age groups. In these contact matrices, the element C_{ij} signifies the likelihood that an individual in the i^{th} age class encounters someone from the j^{th} class. Effectively, the elements of the contact matrix B can be viewed as the product $B_{kl} = \Lambda_{kl} C_{kl}$, where C_{ij} represents contact rates related to social habits and interactions, and Λ_{ij} denotes the disease-dependent transmission probability, which measures the likelihood of infection transmission when there is a contact between i and j.

As a simplified model for illustrating the concept of the Herd Immunity Surface (HIS) involving several population classes, we focus on Italy's physical contact matrix[17]. For simplicity, we restrict our analysis to three age classes: Younger (Y), encompassing individuals aged ≤ 19 years, Median (M) covering those aged 20-69 years, and Elder (E) including individuals aged 70 years or older. Table 1 presents the values of the rescaled contact matrix $\mathscr{C}_{ij} = N_i^{-1} C_{ij} N_j$, while Table 2 provides information on the Italian population distribution across these three classes.

In this toy model, we make the simplifying assumptions of a constant transmission probability $\Lambda_{ij} = \Lambda$ and a uniform recovery rate $\gamma_i = \gamma$ across all classes. Consequently, the Infection Growth Matrix ($\boldsymbol{\mathcal{R}}$) for our simplified model can be expressed as $\boldsymbol{\mathcal{R}} = (\Lambda/\gamma) \cdot \boldsymbol{\mathcal{C}}$. To ensure a realistic representation of the model, we carefully select γ such as the ratio Λ/γ achieves a spectral radius $\|\boldsymbol{\mathcal{R}}\|$ of 3, a value typical of COVID-like epidemics.

We consider a scenario where we have a vaccine supply of V doses, each with an efficacy of δ, capable of immunizing a fraction $r = \delta \cdot v$ of the population (here $v = V/N$ is the fraction that can get the vaccine). We suppose the three age classes to be subject to the mortality rates μ_i of table 2. We can determine the optimal distribution of vaccines, denoted as v_i for the i-th

class, in a way that minimizes the mortality of susceptible individuals s_i, given by $s_i = 1 - r_i$ with $r_i = \delta \cdot v_i$. This optimization problem can be formulated as follows:

$$
\begin{array}{lll}
\min\limits_{\mathbf{v}} & \sum\limits_i \mu_i s_i N_i & \text{minimise the mortality of the susceptible population} \\
\text{s.t.} & \sum\limits_j \mathscr{R}_{ij} s_j \leq 1 & \text{ensuring herd immunity} \\
& s_i = 1 - \delta \cdot v_i \quad \forall i & \text{with a fraction of immunised equal to } \delta \cdot \mathbf{v} \\
& 0 \leq v_i \leq 1 \quad \forall i & \\
& \sum\limits_i v_i N_i = V & \text{respecting a given vaccine availability}
\end{array}
\tag{12}
$$

In Figure 4, we illustrate how the optimal fraction of vaccinated individuals per class changes as the fraction V/N of the population eligible for vaccination varies. It's important to note that, for a given vaccine efficacy, there exists a minimum threshold $\delta_{\min}$ below which no solution exists, meaning that achieving herd immunity is not possible.

The solution to this system corresponds to a *heterogeneous* vaccination policy (VP) that minimizes mortality for a given efficacy $\delta \geq \delta_{\min}$ by varying the fraction $v_i = V_i/N_i$ of vaccinated individuals per class. On the other hand, it is simpler to implement a *homogeneous* VP where the percentage of vaccinated individuals is the same for all classes. Notably, homogeneous VPs require larger quantities of vaccines than heterogeneous VPs (see fig. 5 for an illustration). For example, with an efficacy $\delta = 0.9$, using a heterogeneous VP can lead to saving up to approximately 15% more vaccines compared to a homogeneous VP.
‘

Discussion

Our exploration into vaccination policies for achieving herd immunity has unveiled a complex landscape marked by ethical intricacies and multifaceted considerations. A lens provided by the ethics of vaccination[18] allows us to scrutinize our findings with broader implications. The collective responsibility for herd immunity and ethical tensions surrounding individual rights and harm prevention[19] resonate in our results, especially amidst varying vaccine availabilities. The evolving discourse on conscientious objection to vaccination highlights the ethical debate on coercive versus non-coercive measures, particularly in emergencies[20–23].

Conceptualizing herd immunity as a spectrum, akin to a Pareto Frontier[24], aligns with the 'principle of least restrictive alternative' in public health ethics[25,26]. This perspective recognizes trade-offs, introducing an ethical conundrum in achieving herd immunity while respecting values.

The Herd Immunity Surface (HIS) presents choices for reaching multiple Herd Immunity Threshold (HIT). Different HITs exhibit varying characteristics, with practical constraints, logistical ease, or cultural adaptation possibly favouring some. Optimizing across HIS dimensions is challenging due to incomplete data, necessitating consideration of a broad set of elements. Importantly, the HIS is a Pareto Frontier, and HITs on the HIS are Pareto Optimal[24]. Since such Pareto optimal vaccination policies align with diverse stakeholder preferences, they can reflect nuanced interests, necessitating collaboration for collective benefit. In fact, even amidst efficient policies conflicts may arise, highlighting the need for consensus in shaping vaccination strategies.

Our exploration deepens understanding and emphasizes ethical considerations in shaping effective and morally defensible vaccination strategies. Intersections with public health ethics[27] provide a robust framework, ensuring a scientifically sound and ethically grounded pursuit of herd immunity.

The quest for herd immunity extends beyond models; it demands thoughtful consideration of individual liberties, collective responsibilities, and institutional frameworks[18,25,28]. Engaging with ethical dimensions and acknowledging complexities paves the way for an inclusive and responsible approach aligning with public health goals.

Advancements in data collection, analytics, and modelling techniques[29] are crucial. Real-time monitoring, improved data on vaccine acceptance, and AI insights empower policymakers for an adaptive, data-driven public health approach. Investing in data infrastructure and collaborative efforts align with this progress, navigating vaccination challenges effectively.

Predictive models considering epidemiological, logistical, cultural, and socioeconomic dimensions aid in identifying optimal vaccination thresholds[30,31]. Policymakers must leverage enhanced data and models through clear communication, education, and collaboration for evidence-based vaccination policies. However, while our study underscores the significance of multiple-choice decision-making within a complex landscape, it is imperative to recognize that other dimensions, like addressing social media dynamics[32,33], or fighting the systemic inequities driving mistrust within the information ecosystems[34], present an additional layer of consideration essential for comprehensive pandemic response and effective infodemic management.

Looking ahead, our exploration opens avenues for ongoing ethical scrutiny, sophisticated optimization tools, and innovative approaches to conscientious objection[35–37]. Integrating advanced modeling techniques and AI holds potential for precise and

effective vaccination policies, urging continued collaborative efforts in addressing emerging challenges and refining ethical frameworks for societal well-being.

References

1. Van Kerkhove, M. D. & Ferguson, N. M. Epidemic and intervention modelling–a scientific rationale for policy decisions? Lessons from the 2009 influenza pandemic. *Bull. World Heal. Organ.* **90**, 306–310, DOI: 10.2471/BLT.11.097949 (2012).
2. Santucci, F. *et al.* Evaluating the covid-19 impact in italian regions via multi criteria analysis. *PloS one* **18**, e0285452 (2023).
3. Moodley, K. *et al.* Ethical considerations for vaccination programmes in acute humanitarian emergencies. *Bull. World Heal. Organ.* **91**, 290–297, DOI: 10.2471/BLT.12.113480 (2013).
4. Jecker, N. S., Wightman, A. G. & Diekema, D. S. Vaccine ethics: an ethical framework for global distribution of COVID-19 vaccines. *J. Med. Ethics* **47**, 308–317, DOI: 10.1136/medethics-2020-107036 (2021).
5. Oliva, G., Schlueter, M., Munetomo, M. & Scala, A. Dynamical intervention planning against covid-19-like epidemics. *PloS one* **17**, e0269830 (2022).
6. Chang, M.-H. & Tassier, T. Spatially Heterogeneous Vaccine Coverage and Externalities in a Computational Model of Epidemics. *Comput. Econ.* **58**, 27–55, DOI: 10.1007/s10614-019-09918-7 (2019).
7. Stevens, J. B. *The economics of collective choice* (Routledge, 2018).
8. Kermack, W. O. & McKendrick, A. G. A contribution to the mathematical theory of epidemics. *Proc. royal society london. Ser. A, Containing papers a mathematical physical character* **115**, 700–721 (1927).
9. Marathe, M. V. Mathematical tools for understanding infectious disease dynamics. princeton series in theoretical and computational biology. by odo diekmann, hans heesterbeek and tom britton. xiv+ 502 pp. princeton, nj: Princeton university press. 2014. (2013).
10. Li, M. Y. & Muldowney, J. S. Global stability for the seir model in epidemiology. *Math. biosciences* **125**, 155–164 (1995).
11. Diekmann, O., Heesterbeek, H. & Britton, T. *Mathematical tools for understanding infectious diseases dynamics*. Princeton series in theoretical and computational biology (Princeton University Press, Princeton, 2013).
12. Hethcote, H. W. & Van Ark, J. W. Epidemiological models for heterogeneous populations: proportionate mixing, parameter estimation, and immunization programs. *Math. Biosci.* **84**, 85–118, DOI: 10.1016/0025-5564(87)90044-7 (1987).
13. Scala, A. *et al.* Time, space and social interactions: exit mechanisms for the Covid-19 epidemics. *Sci. Reports* **10**, 13764, DOI: 10.1038/s41598-020-70631-9 (2020). Number: 1 Publisher: Nature Publishing Group.
14. Mistry, D. *et al.* Inferring high-resolution human mixing patterns for disease modeling. *Nat. Commun.* **12**, 323, DOI: 10.1038/s41467-020-20544-y (2021). Number: 1 Publisher: Nature Publishing Group.
15. Diekmann, O., Heesterbeek, J. A. P. & Metz, J. A. On the definition and the computation of the basic reproduction ratio r0 in models for infectious diseases in heterogeneous populations. *J. mathematical biology* **28**, 365–382 (1990).
16. Bapat, R. B. & Raghavan, T. E. S. *Nonnegative matrices and applications*. No. v. 64 in Encyclopedia of mathematics and its applications (Cambridge University Press, Cambridge, UK ; New York, 1997).
17. Mossong, J. *et al.* Social Contacts and Mixing Patterns Relevant to the Spread of Infectious Diseases. *PLoS Medicine* **5**, e74, DOI: 10.1371/journal.pmed.0050074 (2008).
18. Giubilini, A. *The ethics of vaccination* (Springer Nature, 2019).
19. Giubilini, A. Vaccination ethics. *Br. Med. Bull.* **137**, 4–12, DOI: 10.1093/bmb/ldaa036 (2020).
20. Clarke, S., Giubilini, A. & Walker, M. J. Conscientious objection to vaccination. *Bioethics* **31**, 155–161, DOI: https://doi.org/10.1111/bioe.12326 (2017). https://onlinelibrary.wiley.com/doi/pdf/10.1111/bioe.12326.
21. Saunders, B. How Mandatory Can We Make Vaccination? *Public Heal. Ethics* **15**, 220–232, DOI: 10.1093/phe/phac026 (2022). https://academic.oup.com/phe/article-pdf/15/3/220/48944485/phac026.pdf.
22. Kärki, K. Listening to vaccine refusers. *Medicine, Heal. Care Philos.* **25**, 3–9 (2022).
23. Kowalik, M. Ethics of vaccine refusal. *J. Med. Ethics* **48**, 240–243, DOI: 10.1136/medethics-2020-107026 (2022). https://jme.bmj.com/content/48/4/240.full.pdf.
24. Mas-Colell, A., Whinston, M. D., Green, J. R. *et al. Microeconomic theory*, vol. 1 (Oxford university press New York, 1995).

25. Beauchamp, T. L., Childress, J. F. *et al. Principles of biomedical ethics* (Oxford University Press, USA, 2001).

26. Gostin, L. O. Beyond moral claims: A human rights approach in mental health. *Camb. Q. Healthc. Ethics* **10**, 264–274 (2001).

27. Childress, J. F. *et al.* Public Health Ethics: Mapping the Terrain. *J. Law, Medicine & Ethics* **30**, 170–178, DOI: 10.1111/j.1748-720X.2002.tb00384.x (2002).

28. Gostin, L. O. & Wiley, L. F. *Public health law: power, duty, restraint* (University of California Press, Oakland, California, 2016), third edition edn.

29. Teerawattananon, Y. *et al.* A systematic review of methodological approaches for evaluating real-world effectiveness of covid-19 vaccines: Advising resource-constrained settings. *PLoS One* **17**, e0261930 (2022).

30. of Vaccines, H. Cultural perspectives on vaccination (2021). Accessed on November 9, 2023.

31. Saha, S. & Annamaraju, P. Cultural competence and vaccination uptake: A systematic review (2021). Accessed on November 9, 2023.

32. Cinelli, M. *et al.* The COVID-19 social media infodemic. *Sci. Reports* **10**, 16598, DOI: 10.1038/s41598-020-73510-5 (2020). Number: 1 Publisher: Nature Publishing Group.

33. Briand, S. C. *et al.* Infodemics: A new challenge for public health. *Cell* **184**, 6010–6014 (2021).

34. Posada, A., Lopez Inigo, R. & Majid, B. Inequity driven mistrust (2023). Accessed on November 9, 2023.

35. Oliver, K. & Boaz, A. Transforming evidence for policy and practice: creating space for new conversations. *Palgrave Commun.* **5**, 60 (2019).

36. Rickinson, M. *et al.* Insights from a cross-sector review on how to conceptualise the quality of use of research evidence. *Humanit. Soc. Sci. Commun.* **8**, 141 (2021).

37. Maas, T. Y., Pauwelussen, A. & Turnhout, E. Co-producing the science-policy interface: towards common but differentiated responsibilities. *Humanit. Soc. Sci. Commun.* **9**, 93 (2022).

38. Rapid Risk Assessment: Coronavirus disease 2019 (COVID-19) in the EU/EEA and the UK– ninth update. https://www.ecdc.europa.eu/en/publications-data/rapid-risk-assessment-coronavirus-disease-2019-covid-19-pandemic-ninth-update. [Online; accessed 02-April-2021].

Competing interests

The authors declares no competing interests.

Data Availability

All data used to produce the examples of this paper are available in the main text. These data were derived from the following resources available in the public domain. The data on countries' contact matrices can be found in the paper by Mossong et al[17]; population data can be retrieved from the website of the Italian National Institute of Statistics (ISTAT) at the url `https://esploradati.istat.it/`; while the mortality data used to set up the optimization example can be found in the "9^{th} update of the Rapid Risk Assessment for COVID-19 in the EU and the UK" report[38].

Contact Matrix $\mathcal{C}$			
	Y	M	E
Y	2.350	1.694	0.455
M	0.122	0.590	0.139
E	0.736	1.980	0.800

Infection Growth Matrix $\mathcal{R}$			
	Y	M	E
Y	1.966	1.417	0.381
M	0.102	0.494	0.116
E	0.616	1.656	0.669

Table 1. Tables presenting the Contact Matrix $\mathcal{C}$ and the derived Infection Growth Matrix $\mathcal{R}$ (IGM). The Contact Matrix is based on empirical data, specifically from the study by Mossong et al.[17], providing insights into the frequency of interactions between different age groups or population classes. The Infection Growth Matrix (IGM) is computed from the Contact Matrix as $\mathcal{R} = (\Lambda/\gamma) \cdot \mathcal{C}$ by adjusting the parameter Λ/γ to achieve a realistic spectral radius of $\|\mathcal{R}\| = 3$, which is commonly observed in epidemics with COVID-like characteristics. The three age classes indicated as Y,M,E correspond to Young (0-19 years old), Median (20-69), and Elder (70+) individuals.

	Y (00-19)	M (20-69)	E (70+)
population N	11.316.741	40.776.591	10.297.032
mortality μ	0.001	0.008	0.148

Table 2. Population distribution and mortality rates for different age classes. The table presents the population counts N_i for each age class (Y: Younger, M: Median, E: Elder) and their corresponding mortality rates μ_i. Population data are available from the website of the Italian National Institute of Statistics (ISTAT) at the url https://esploradati.istat.it/ while mortality data are relative to the 2020's COVID-19 risk assessment[38].

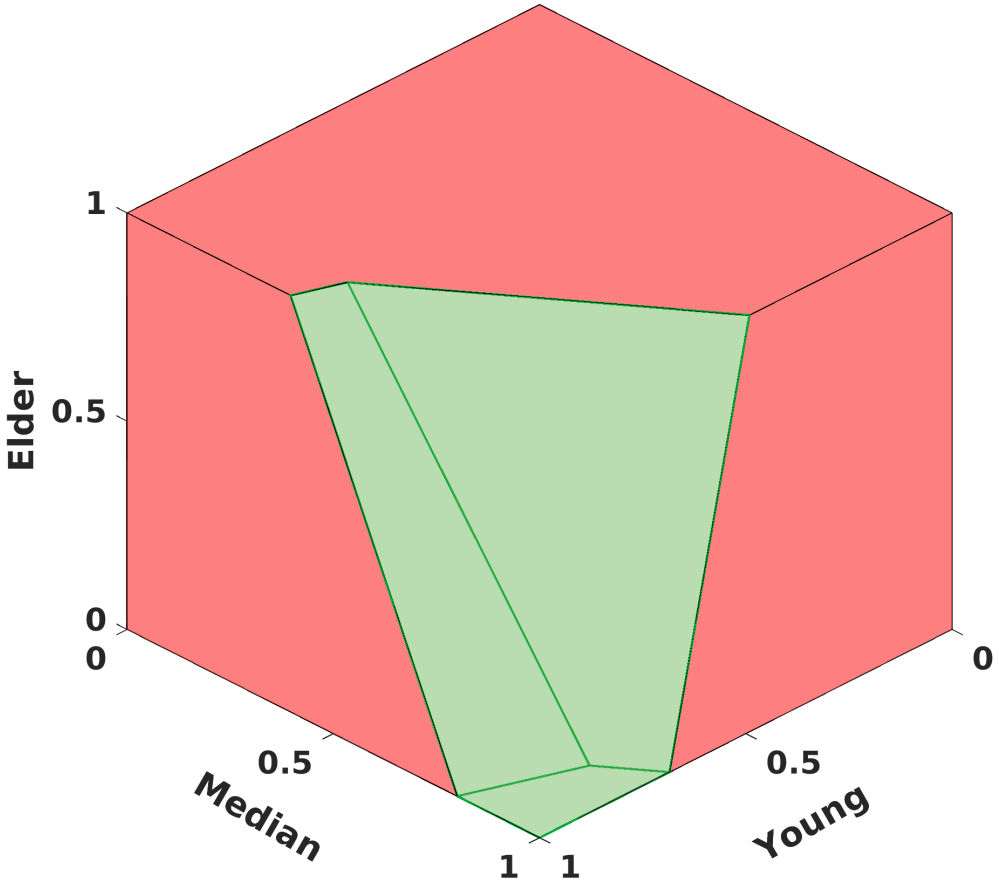

Figure 1. A 3D visualization of the Herd Immunity Surface (HIS) for a population divided into three classes: Young (0-19 years old), Median (20-69), and Elders (70+). The axes represent the fraction of individuals resistant to the epidemic, including those who are vaccinated or have recovered from the disease. In this visualization, the HIS (depicted in green) corresponds to a scenario with the Infection Growth Matrix (IGM) of Table 1. The red region represents unstable states where the epidemic grows. The HIS, as a complex multi-dimensional surface, highlights the existence of multiple herd immunity thresholds. Choosing a specific point on the HIS becomes a strategic decision with profound ethical, economic, and societal implications, demonstrating the intricate balance required in heterogeneous populations when designing vaccination strategies.

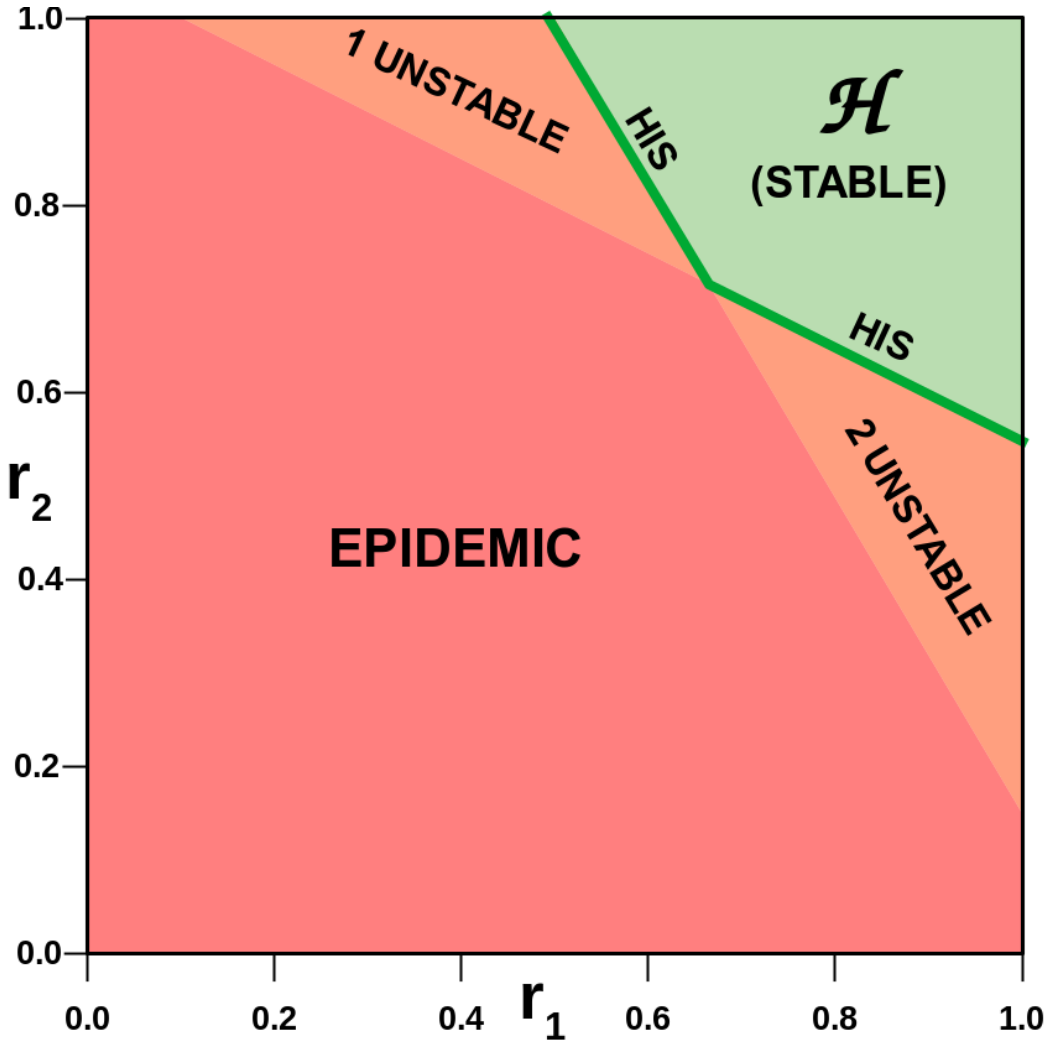

Figure 2. A depiction of a two-dimensional Herd Immunity Surface (HIS) represented as a line. The HIS corresponds to an SIR/SEIR model with two classes, characterized by specific parameters: $1/R_{11} = 0.5$, $1/R_{21} = 0.9$, $1/R_{22} = 0.45$, $1/R_{12} = 0.85$. The HIS effectively separates regions into stable and unstable categories. In the region labeled "1 UNSTABLE", infections in class 1 increase, while infections in class 2 decrease. This behavior suggests that class 2 has effectively reached its herd immunity threshold, rendering it more resistant to infections. Similarly, in the region labeled "2 UNSTABLE", infections in class 2 increase, while infections in class 1 decrease, signifying that class 1 has reached its herd immunity threshold. Using the notation of eq. 8, the HIS $\mathscr{H}^*$ represents the boundary of the stable set $\mathscr{H} = \mathscr{H}_{\{1,2\}}$ for all classes. The HIS $\mathscr{H}^*_{\{1\}}$ for class 1 demarcates the boundary of the stable set $\mathscr{H}_{\{1\}}$ encompassing the stable region $\mathscr{H}$ and the "2 UNSTABLE" region. Similarly, the HIS $\mathscr{H}^*_{\{2\}}$ for class 2 outlines the boundary of the stable set $\mathscr{H}_{\{2\}}$ encompassing the stable region $\mathscr{H}$ and extending into the "1 UNSTABLE" region.

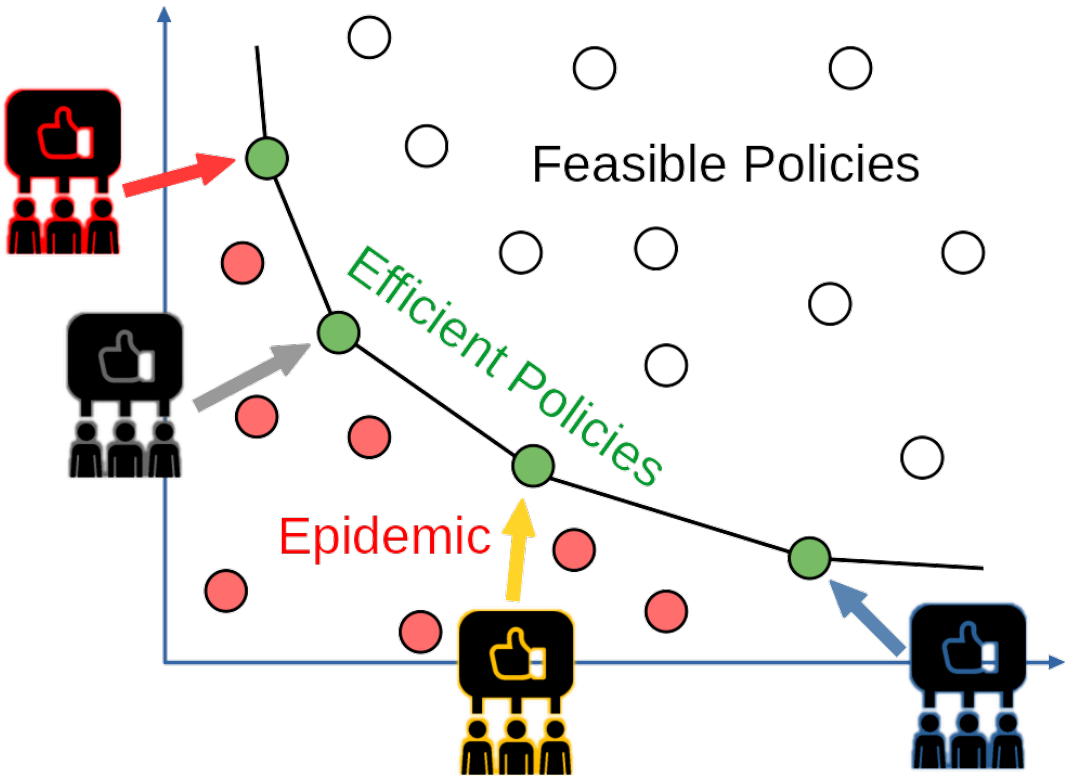

Figure 3. Optimizing Vaccination Policies for Stakeholder Preferences: the Pareto Frontier showcases efficient vaccination policies, catering to diverse stakeholder preferences. The Pareto Front's dynamic shapes reflect the nuanced interests of stakeholders, influencing policy outcomes. Efficient policies are the ones on the Heard Immunity Surface and represent the minimum threshold to be reached to achieve herd immunity. Collaborative efforts among stakeholders are necessary to drive the optimization of vaccination policies for collective benefit. Having many possible efficient policies, it is possible that conflicts set in among interest groups or stakeholders with different interests.

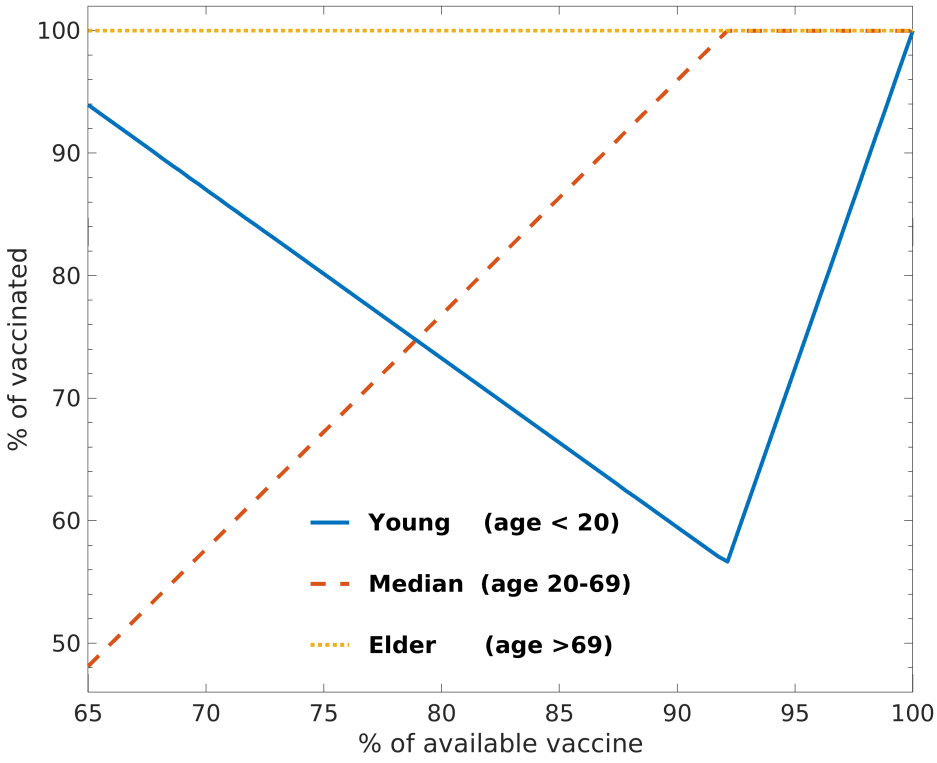

Figure 4. Vaccination policies optimized to reach herd immunity while minimizing the overall mortality for a vaccine with a 95% effectiveness rate, considering varying vaccine availability (no feasible solution can be found for a vaccine availability $\lesssim 64.58\%$). The horizontal axis represents the fraction v of the population eligible for vaccination, while the vertical axis indicates the percentage of individuals per class who need to be vaccinated to achieve the Herd Immunity Surface (HIS). The example demonstrates how vaccination policies differ based on vaccine availability, even when the goal remains the same. Given the higher mortality rates among the elderly population, the optimal strategy is to vaccinate all Elders (100%) to maximize the impact. However, for lower vaccine availability (less than $\sim 79\%$ of the population), prioritizing Younger individuals becomes more effective due to their stronger influence on the contact matrix. In contrast, with larger vaccine supplies, focusing efforts on middle-aged individuals becomes viable, as they are more numerous and exhibit higher mortality rates.

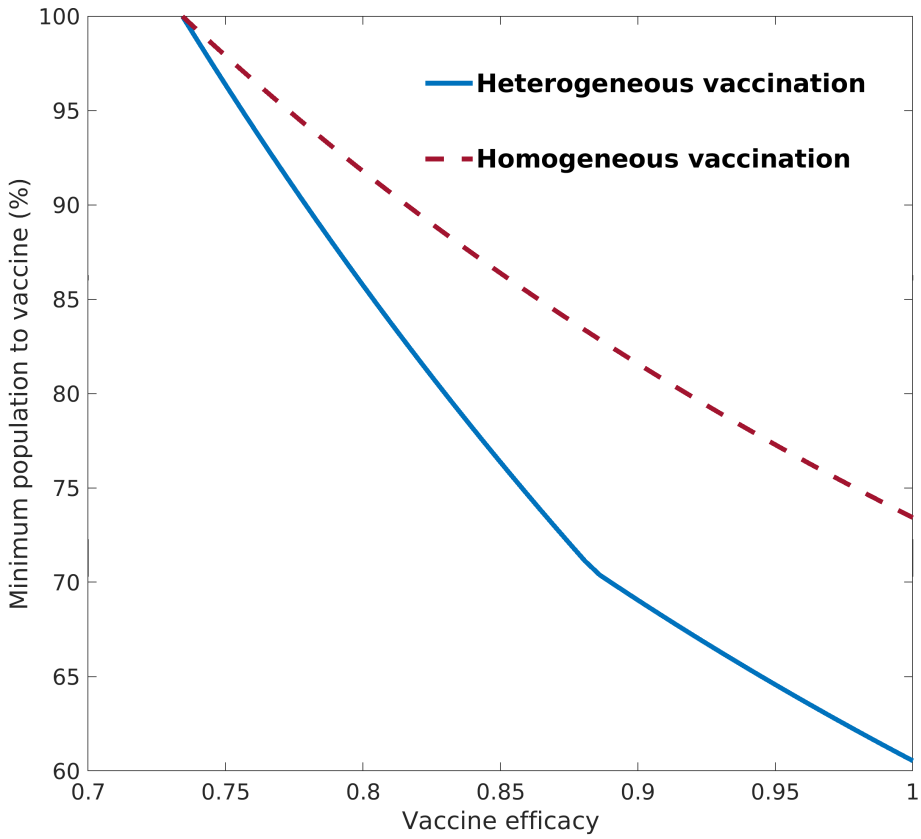

Figure 5. Minimum percentage of the population required to be vaccinated to achieve herd immunity, plotted against vaccine effectiveness. The blue solid line represents a heterogeneous vaccination policy, allowing different population classes to have varying percentages of vaccinated individuals, while the red dashed line represents a homogeneous vaccination policy, where the percentage of vaccinated individuals is the same for all classes. In general, due to the convex nature of the Herd Immunity Surface parameter region, implementing heterogeneous vaccination policies can help reduce the required resources (i.e., vaccines) compared to homogeneous policies. While the homogeneous vaccination may also be exceptionally the optimal policy, the convexity of the parameter region predominantly favors the advantages of heterogeneous approaches.