Lives Saved from Age-Prioritized Covid-19 Vaccination

Joshua R. Goldstein[*1], Thomas Cassidy[†2], and Ayesha S. Mahmud[‡1]

[1]Department of Demography, University of California, Berkeley
[2]Department of Mathematics, Bucknell University

March 19, 2021

Abstract

BACKGROUND The criteria used to allocate scarce COVID-19 vaccines are hotly contested. While some are pushing just to get vaccines into arms as quickly as possible, others advocate prioritization in terms of risk.

OBJECTIVE Our aim is to use demographic models to show the enormous potential of vaccine risk-prioritization in saving lives.

METHODS We develop a simple mathematical model that accounts for the age distribution of the population and of COVID-19 mortality. This model considers only the direct live-savings for those who receive the vaccine, and does not account for possible indirect effects of vaccination. We apply this model to the United States, Japan, and Bangladesh.

RESULTS In the United States, we find age-prioritization would reduce deaths during a vaccine campaign by about 93 percent relative to no vaccine and 85 percent relative to age-neutral vaccine distribution. In countries with younger age structures, such as Bangladesh, the benefits of age-prioritization are even greater.

CONTRIBUTION For policy makers, our findings give additional support to risk-prioritized allocation of COVID-19 vaccines. For demographers, our results show how the age-structures of the population and of disease mortality combine into an expression of risk concentration that shows the benefits of prioritized allocation. This measure can also be used to study the effects of prioritizing other dimensions of risk such as underlying health conditions.

[*]josh.goldstein@berkeley.edu

[†]tcassidy@bucknell.edu

[‡]mahmuda@berkeley.edu

1

Introduction   The availability of highly effective COVID-19 vaccines creates the opportunity for rapidly lowering mortality. A simplified demographic model applied to the United States shows that age-prioritization would lower deaths by 85% relative to an age-neutral vaccine campaign. Even though optimal vaccination is not fully achievable and other considerations need to be taken into account, age-prioritization is an important life-saving benchmark against which alternative vaccination policy can be compared.

   The enormous life saving effect of age-prioritization is due to the joint effects of rapidly increasing mortality at older ages (Davies et al., 2020; Goldstein and Lee, 2020), and the relatively small number of elderly (Fig 1 A and B). In this paper, we provide a framework for measuring the benefits of age-prioritization. Our approach considers only the direct effects on saving lives of those who are vaccinated, and does not take into account the effects of reduced viral transmission. This simplified framework can be extended to apply to any other aspect of risk beyond just age.
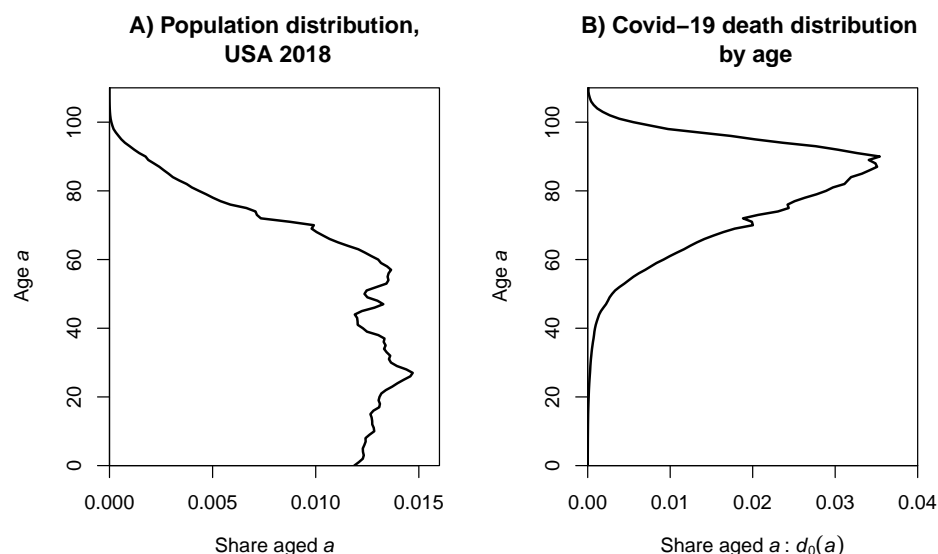

Figure 1: Age structure of the U.S. 2018 population and of Covid-19 deaths.

A model   To illustrate the life-saving potential of prioritizing vaccines by age, we consider a vaccine campaign that protects a constant share of the population each day until everyone has been vaccinated. We let p denote the share of the population that is protected from mortality. With constant shares vaccinated each day, p is also equal to fraction of time it takes to vaccinate everyone

   We further assume deaths at age a would remain at a constant level $d_0(a)$ in the

2

absence of vaccination.[1] During the vaccination campaign, mortality is determined by multiplying the share at each age that remains unvaccinated by the baseline age-specific COVID-19 death rates.

Under an age-neutral campaign, when a share p of the population is protected, deaths will be $(1 - p)$ of the baseline levels at all ages. Thus,

$$D_N(p) = (1 - p) \int_0^\infty d_0(a) \, da \qquad (1)$$

tells us that the ratio of deaths in an age-neutral campaign to baseline, $D_N(p)$, will decline linearly from unity to zero over the course of the campaign. The path of the age-neutral campaign is given by the dashed line Figure 2. The ratio of deaths over the entire campaign relative to baseline is the area under this line. This area, denoted $\mathcal{D}_N$, corresponds to one-half the deaths that would have occurred without vaccination.

Now consider an age-prioritized campaign, with deaths declining as the oldest become protected from risk. Let $A(p)$ be the age over which the share of the population is p. In this case, the deaths relative to baseline when a fraction p have been protected is

$$D_A(p) = \int_0^{A(p)} d_0(a) \, da. \qquad (2)$$

The curve labeled $D_A(p)$ in Figure 2 shows deaths during an age-prioritized campaign decline quickly at first as the most at-risk are vaccinated and then more slowly thereafter.

Over the entire course of the campaign, the ratio of deaths to baseline can be calculated by summing $D_A(p)$ over all values of p, which increases from 0 to 1 steadily as the campaign goes from start to finish.

$$\mathcal{D}_A = \int_0^1 D_A(p) \, dp \qquad (3)$$

This gives the deaths during an age-prioritized campaign relative to the baseline case without a vaccine. For the United States $\mathcal{D}_A = 0.075$, much smaller than the value of 0.5 of total deaths of a neutral vaccination campaign.

As is shown in the appendix, the effectiveness of age-prioritized vaccination can be re-expressed in terms of $d_0(a)$, the baseline age distribution of COVID-19 deaths, and $1 - F(a)$, the proportion of the population above age a:

$$\mathcal{D}_A = \int_0^1 D_A(p) \, dp = \int_0^\infty d_0(a)[1 - F(a)] \, da. \qquad (4)$$

Equation (4) has an intuitive interpretation. The middle term tells us that we can calculate the mortality over the entire time period by summing over the times

---

[1]This model could be extended to allow for a surface of deaths $d(a, t)$ that changes with both age and time. Similarly, our assumption that pace of vaccination is constant over the course of the campaign could be relaxed by expressing p as a function of time.
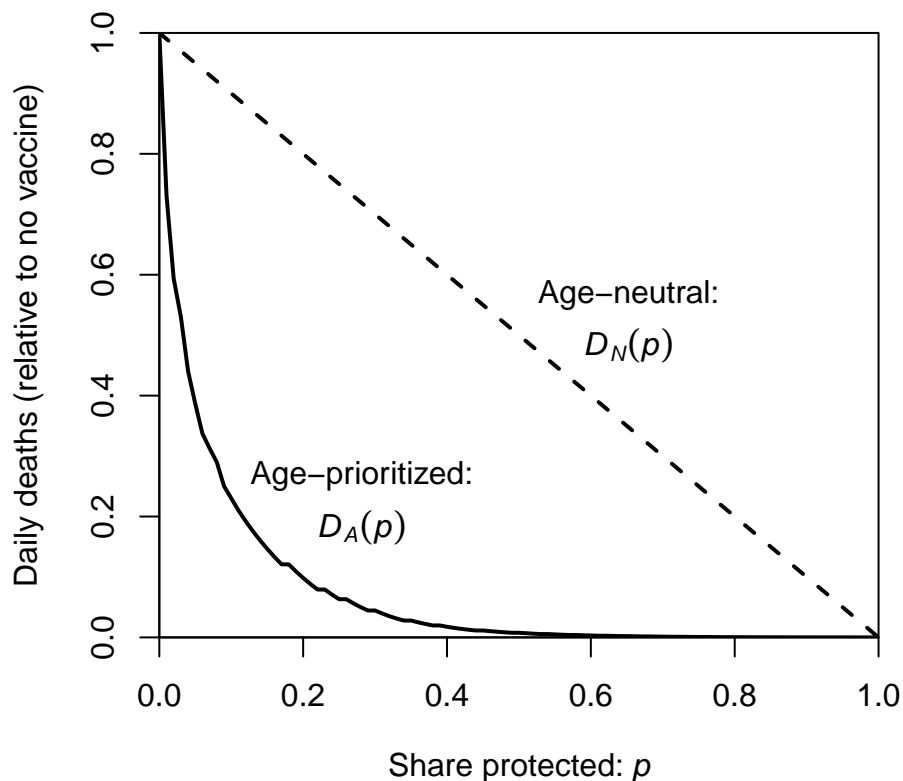
Figure 2: **The effectiveness of vaccine prioritization in reducing deaths**
The heights of the curves give the deaths, relative to baseline, when a fraction p is
protected. Under the assumptions of our model, the sum of the area under each
curve gives the total deaths over the full course of the campaigns, relative to what
would have been observed with no vaccine. The death toll relative to baseline of
an age-neutral campaign ($\mathcal{D}_N = 0.500$) shrinks dramatically ($\mathcal{D}_A = 0.075$) under
age-prioritization. Source: The $D_A(p)$ curve is estimated from equation (3) using
the U.S. age structure of population and Covid-19 mortality shown in Figure 1.

when a fraction p is protected. The right-most term shows that one can also sum
over age, with $1 - F(a)$ telling us the fraction of the campaign that people aged a
are exposed, and $d_0(a)$ giving the deaths at that age.

4

Results   We apply our model to estimates of age-specific COVID-19 death rates based on cumulative reports by CDC and population estimates from U.S. Census (Goldstein and Lee, 2020). A baseline schedule of daily deaths $d_0(a)$ at age a is obtained by applying these rates to the age-structured population (Fig 1 B).

In the United States, we calculate that $\mathcal{D}_A$ is about 0.075. This means that an idealized age-prioritized campaign would lower the lower the death toll by about 92.5% relative to no vaccine, and about 85% relative to the age-neutral campaign. The life-saving effects of age-prioritization can be seen in the much smaller area under the prioritized curve in Fig 2).

To make the magnitudes more concrete, consider a hypothetical campaign lasting 150 days that correspond very roughly to the United States case. Assume that deaths would remain at 2,000 per day over this time in the absence of vaccination. Over these 150 days, 300,000 people would die in the absence of a vaccine, 150,000 would die with age-neutral vaccination, and about 22,000 would die under age-prioritized vaccination.

As vaccination distribution expands in developing countries with younger age-structures, age-prioritization may prove even more effective. The age-structure of COVID-19 mortality risk is not thought to vary enormously across countries, but population pyramids do vary. Applying the same standard age-structure of COVID-19 mortality rates to the observed age-structures, we find that $\mathcal{D}_A/\mathcal{D}_N$ is 0.06 in Bangladesh, 0.26 in Japan (compared to 0.15 in the United States).

Much of the benefit of age-prioritization can be obtained by following strict-prioritization for the oldest segment of the population. For example, strict age-ordering for the first 20 percent of the population with age-neutral open eligibility thereafter is described by tracing along the age-prioritized $D_A(p)$ curve from $p = 0$ to $p = 0.2$ for the age-prioritized period and then drawing a straight-line (a "chord") from $D_A(p = .2)$ to $D_A(p = 1)$. The area under partially age-prioritized curve for the United States is about 0.100, as compared to 0.075 for full prioritization.

The model we present for age-prioritization is also relevant to any other ordered risk factor. The variable we use use for a could, for example, represent an ordered risk score that combines factors such as age, health status, and risk of exposure. The quantity $\mathcal{D}_A$ will be smaller when the concentration of risk is greater. The benefit of taking account of additional risk factors beyond age can be measured by recalculating alternative $\mathcal{D}_A$ values for different measures of risk.

Viewed in this way, the quantity $\mathcal{D}_A$ can be seen as a measure of the concentration of population risk. In fact, graphically it bears a strong resemblance to the Gini index and other measures of concentration, which researchers such as Shkolnikov et al. (2003) have looked at regarding age and mortality. A difference in this case is that it is not just the distribution of deaths that matters, but also the age-distribution of the population.

Discussion   Strict risk ordering of vaccination is optimal in that it most reduces deaths among those receiving the vaccine (assuming no indirect protection from vaccine roll-out). However, there are other considerations that can and arguably should be taken into account. For example protecting front-line workers, teachers,

and others may produce positive benefits for the larger society beyond protecting these individuals. The equity implications of risk-prioritization are also important to consider. More complete analysis accounting for transmission risks and indirect effects (e.g., Bubar et al. (2021)) is also necessary.

Generous eligibility for vaccines is an understandable popular demand, particularly in the face of the slow roll-out and vaccine hesitancy. However, opening eligibility too soon risks losing much of the enormous benefits of prioritization. Even worse, open eligibility in the face of vaccine shortages can result in the least at risk going first, as the same resources that protect individuals from infection and death are used to move to the front of the line. Such "reverse prioritization" would result in even more deaths than the theoretically "neutral" approach.

The politics and logistics of vaccine allocation are complicated. But the demography of the direct effect of vaccination is simple, particularly regarding age: the elderly are at much higher risk but represent only a small fraction of the population. A successful vaccination campaign is not just a question of getting needles into arms. It matters who is vaccinated when. Making every effort to vaccinate the elderly first can save an enormous number of lives.

Mathematical appendix

Definitions:

p The fraction of the population protected from mortality through vaccination, beginning at 0 and ending at 1. A general formulation could have $p(t)$ take any non-decreasing form. We consider the simple case of a constant rate of vaccination, or $p(t) = t$, with $0 \le t \le 1$.

$d_0(a)$ The number of people dying of Covid-19 at each age a per unit of time at baseline when no-one is protected. In its normalized form, with $\int_0^\infty d_0(a) = 1$, this gives the share of baseline deaths at age a.

$F(a)$ The cumulative distribution function of the population, i.e., the fraction of the population aged less than a. $1 - F(a)$ is the fraction aged more than a.

$A(p)$ The age above which there is a fraction p of the population. Under age-prioritization with constant shares of the population being protected over time, $A(p)$ is the age above which there is no mortality when a fraction p has been been protected.

$\mathcal{D}_A$ The ratio of deaths to baseline over the entire course of an age-prioritized vaccination campaign that protects older people first.

$\mathcal{D}_N$ The ratio of deaths to baseline over the entire course of an age-neutral vaccination campaign that protects an equal share of each age group each day.

We now calculate the deaths over the full time period of the campaign. Sums of normalized $d_0(a)$ (see above) give deaths as a fraction of the baseline total.

Substituting (2) into (3), we have

$$\mathcal{D}_A = \int_0^1 \int_0^{A(p)} d_0(a) \, da \, dp$$

Reversing the order of integration, noting that $p = 1 - F(a)$ is the inverse of $A(p)$, gives

$$\mathcal{D}_A = \int_0^\infty d_0(a) \int_0^{1-F(a)} 1 \, dp \, da = \int_0^\infty d_0(a)[1 - F(a)] \, da$$

the result in equation (4) of the text.

Similarly, we can consider the deaths-relative-to-baseline over the course of an age-neutral campaign

$$\mathcal{D}_N = \int_0^1 \int_0^\infty d_0(a)[1 - p] \, da \, dp,$$

which evaluates to $\int_0^\infty d_0(a) \, da/2$, or one-half of the baseline total.

References

Kate M Bubar, Kyle Reinholt, Stephen M Kissler, Marc Lipsitch, Sarah Cobey, Yonatan H Grad, and Daniel B Larremore. Model-informed covid-19 vaccine prioritization strategies by age and serostatus. *Science*, 2021.

Nicholas G. Davies, Petra Klepac, Yang Liu, Kiesha Prem, Mark Jit, and Rosalind M. Eggo. Age-dependent effects in the transmission and control of COVID-19 epidemics. *Nature Medicine*, pages 1–7, June 2020. ISSN 1546-170X. doi: 10.1038/s41591-020-0962-9.

Joshua R Goldstein and Ronald D Lee. Demographic perspectives on the mortality of covid-19 and other epidemics. *Proceedings of the National Academy of Sciences*, 117(36):22035–22041, 2020.

Vladimir M Shkolnikov, Evgueni E Andreev, and Alexander Z Begun. Gini coefficient as a life table function: computation from discrete data, decomposition of differences and empirical examples. *Demographic Research*, 8:305–358, 2003.