Optimal Dose Stretching Under Vaccine Supply Constraints

Michael Kremer
Kenneth C. Griffin Department of Economics, University of Chicago
E-mail: kremermr@uchicago.edu

Christopher M. Snyder
Department of Economics, Dartmouth College
E-mail: chris.snyder@dartmouth.edu

Witold Więcek
Development Innovation Lab, University of Chicago
E-mail: witold.wiecek@gmail.com

July 2026

Abstract: The urgency to quell an epidemic by vaccinating a large population as quickly as possible can lead to shortages, which can be alleviated with fractional doses. We study optimal fractional dosing in a tractable model capturing the essence of more complex epidemiological studies. Under a concave dose-response function, fractional dosing is always optimal if supplies are insufficient to deliver full doses to the entire population, and optimal dose strength is weakly increasing in vaccine supply. The same is often true under an s-shaped dose-response function, although exceptions can arise where a further tightening of extremely scarce supplies leads to more concentrated doses. The analysis then shifts to deriving quantitative results from a dynamic epidemiological model, allowing mortality and disease transmission to vary by age, calibrated to Covid-19 parameters. The gains from fractional dosing can be substantial, increasing the number of deaths averted by full doses by as much as 40%.

JEL Codes: I18 (Health: Government Policy, Public Health); O38 (Innovation: Government Policy); C63 (Simulation Modeling)

Keywords: vaccine, pandemic, epidemiology, dose response, supply

Acknowledgments: The authors are grateful to Rachel Glennerster, Nora Szech, and seminar participants at Harvard, M.I.T., University of Chicago, and the NBER Entrepreneurship and Innovation Policy and the Economy conference for insightful comments. Andreas Petrou-Zeniou, Alexandre Simoes Gomes, and Leonardo Wang provided excellent research assistance. Kremer and Snyder thank their colleagues at the Market Shaping Accelerator for advice and support. Funding from the Peter G. Peterson Foundation Pandemic Response Policy Research Fund is gratefully acknowledged.

1. Introduction

In an epidemic, infections and consequent demand for vaccines can grow exponentially. Production lines cannot be added instantaneously for any product, but lags in capacity expansion are particularly long for vaccines. Vaccines are complex biological formulations broadly administered to a healthy population and so receive unusually close regulatory scrutiny. The production process is closely monitored after licensure to ensure the purity and potency of doses in each lot match the approved formulation (Baylor and Marshall, 2013). While it may not be possible to substantially scale up the supply of the active ingredient (antigen) used in the vaccine to meet exploding demand during an epidemic, dose-stretching policies offer an alternative for obtaining more doses from existing supplies. Cutting doses in half could allow the production of twice as many doses from a given antigen supply, doubling the rate at which doses can cover the population. Cutting dose strength involves a tradeoff, however: less potent doses may cover more people but may be less effective for individuals.

In this paper, we evaluate the tradeoff between efficacy and coverage in our analysis of optimal dose stretching. We will label the optimal dose in an epidemic as "stretched" if it is smaller than the optimal dose in normal times outside of an epidemic, when there is no race to cover the population and thus antigen supplies are relatively unconstrained.

Dose-stretching strategies have been piloted in real-world epidemics. During the mpox outbreak in the United States in 2022, Emergency Use Authorization was granted for one fifth of the standard dose to be administered intradermally. Studies found that the lower dose retained substantial efficacy, estimating a 75% effectivenss against infection compared to 86% for the full dose (Dalton et al., 2023). Fifth doses were also used to stretch vaccine supplies during a yellow fever epidemic in Brazil in 2018 (Pan-American Health Organization, 2018) and to economize on stockpiles of inactivated polio vaccine used for routine immunizations in low-income countries (Okayasu et al., 2017).

Other ways of stretching doses besides fractional dosing have also been used. The World Health Organization recommended cutting the initially approved three-dose regimen of the human papillomavirus (HPV) vaccine down to one dose (World Health Organization, 2022b), citing numerous studies finding no loss of efficacy (World Health Organization, 2022a). During a cholera outbreak in 2016, Zambian officials stretched vaccine supplies by cutting the standard two-dose

regimen down to one dose (Leung, Eaton, and Matrajt, 2022). During the Covid-19 pandemic, the United Kingdom extended the time between first and second doses from three to 12 weeks to allow faster rollout of first doses to more people; Canada, Denmark, Norway, India, and South Africa followed suit with delayed second-dose strategies (Imai et al., 2023).

An extensive scientific literature has supported the potential benefits of dose stretching. The typical paper in this literature builds a complex epidemiological model, calibrated to the relevant parameters governing dose-response and transmissibility for the disease under study—whether cholera (Azman et al., 2015; Leung, Eaton, and Matrajt, 2022), Covid-19 (Matrajt et al., 2021; Więcek et al., 2022), influenza (Matrajt et al., 2015), mpox (Dimitrov, Adamson, and Matrajt., 2023), or yellow fever (Wu et al., 2016). They generally conclude that dose stretching can improve upon full doses in an epidemic, reducing infection and death rates, even if dose stretching entails some decrease in vaccine effectiveness.

With few exceptions, the literature compares the performance of full doses to an exogenous alternative dose-stretching policy, often the policy that was actually implemented. The analysis of isolated alternatives begs the question of whether other dose stretching policies would be better still, and which would be optimal among those. Riley, Wu, and Leung (2007) stands out from much of the literature in simulating a fine grid of dose-stretching policies and selecting the optimal grid point—in the specific application, the fractional dose that minimizes influenza infections. As path-breaking as this paper is, it is still based on simulations of a complex epidemiological model and so remains something of a "black box." It is difficult to distill principles behind why one grid point is optimal and not another and difficult to judge whether the insights are general or depend on the specific disease parameterization.

We seek to fill this gap in our understanding by constructing a model that preserves the crucial elements of the complex epidemiological models but is tractable enough to provide analytical solutions for general functional forms. With these analytical solutions in hand, we can proceed to synthesize the principles behind the optimal policy and understand the comparative statics of how the policy prescription varies across different environments.

We show how the optimal fractional dosing policy varies with the scarcity of supplies and the shape of the dose-response function, which traces out the gain in efficacy for larger fractional doses. Two shapes are most commonly considered in the pharmacological literature for the dose-response function: concave and sigmoid (i.e., s-shaped). We find that the results differ qualitatively

across these two shapes. The contrast is particularly stark in the case of costless vaccination, which we begin the analysis with. With costless vaccination and a concave dose-response function, the whole population should be vaccinated with full doses if possible; but if supplies of the active ingredient are too scarce for that, fractional doses should be spread to the whole population. This prescription leads to a straightforward formula for optimal dosage in the region of scarce supplies: dose size should vary linearly with the supply of the active ingredient. This is exactly the relationship that Riley, Wu, and Leung (2007) find in every case (see their Figure 2)—a result we could have predicted without resorting to simulations, knowing only that their calibrated dose-response functions are concave.

A sigmoid shape embodies threshold effects in the immune response, documented by the medical literature for certain vaccines. If the dose-response function is sigmoid rather than concave, we find that the optimal dose fraction no longer converges to zero as supplies become scarce but is bounded below. In the simple case of costless vaccination, the bound is the point of unit elasticity, where a percentage increase in dose size yields the same percentage increase in vaccine efficacy, recovering a result appearing in Riley, Wu, and Leung (2007) and Wu et al. (2016).

While a model that abstracts from side effects (along with other costs of administering stronger doses and covering more of the population) provides a useful starting point for the analysis, our main analysis builds those costs back in. With a concave dose-response function, we continue to find that the optimal fractional dose is weakly increasing in vaccine supplies. With a sigmoid function, we provide necessary and sufficient conditions for there to be a region of extremely scarce supplies in which further decreases in supply leads the optimal fractional dose to be stronger, not weaker. The intuition for this surprising result is that when supplies are extremely scarce, protecting vaccinated individuals becomes the paramount concern since limited vaccine supplies will lead to soaring infection rates. Concerns that administering a more potent dose to them might entail worse side effects become relatively less important.

To assess the quantitative importance of the results from the reduced-form model, the second part of the paper turns to constructing and simulating a full-blown epidemiological model. We adapt the epidemiological model from our own previous work Więcek et al. (2022), which allows for heterogeneity in mortality risk and a flexible contact matrix across seven age cohorts and is calibrated to the Covid-19 pandemic. The dynamic vaccination process distributes doses to cohorts prioritized by risk in a flow constrained by available capacity. Like Riley, Wu, and Leung (2007),

we solve for the optimal fractional dose and determine how it varies with vaccine supplies. We add the novel element of calibrating both concave and sigmoid dose-response functions and seeing how the optimal fractional dose differs across those functional forms. Complementing our analysis of optimal uniform fractional doses, we solve for the optimal age-varying fractional doses, a level of detail that is novel in the literature. The few previous studies that allowed any variation in dosing by age considered a single age threshold between those receiving full doses and those receiving an exogenous stretched dose (Wu et al., 2016; Matrajt et al., 2021; Leung, Eaton, and Matrajt, 2022).

The discussion so far has traced the connection between our paper and the biomedical literature. The paper also continues a line of theoretical research integrating epidemiology into economic models, which gained attention during the Covid-19 pandemic. See Philipson (2000) and Philipson (2016) for surveys of the early literature and McAdams (2021) and Avery (2024) for more recent surveys. Recent work has studied optimal subsidy policy (Goodkin-Gold et al., 2024), the relative merits of vaccines versus therapeutics (LaTorre, Malik, and Marsiglio, 2020; Toxvaerd and Rowthorn, 2022), the substitutability between vaccination and social distancing (Atkeson, 2021; Avery, 2024), and the formation of robust policy in the face of model and parametric uncertainty (Manski, 2020; Gans, 2022). Our paper contributes the analysis of a new issue, fractional dosing of vaccines.

The paper is organized as follows. Section 2 introduces the reduced-form model. We derive the optimal fractional dose in this model under costless vaccination in Section 3 and under costly vaccination in Section 4. Section 6 provides calibrations in a detailed epidemiological model. Section 7 concludes. A series of appendixes provide proofs of propositions, details behind the calibrations, and additional calibrations for alternative parameters.

2. Model

While we will eventually move to calibrating a richer epidemiological model, the analysis up to that point will use the model presented in this section. The model is simplified along a number of dimensions: it is static, involves homogeneous consumers, and embodies complex epidemiological processes in a reduced-form benefit function.

The model involves three sets of actors: a planner, suppliers, and consumers. The planner is the main decision maker in the model. The planner formulates a vaccine using ingredients purchased from suppliers and distributes it for free to consumers (as most countries did during the Covid-19

pandemic and other epidemics).

A key variable in our model is the strength of a vaccine course. In practice, there are various ways to increase a vaccine's strength: incorporating more doses in a full course, changing the interval between doses, or using more of the active ingredient (antigen) in each dose. Our static model glosses over such distinctions, since different ways of packaging a given total quantity of the active ingredient entail the same total supply in a static setting. We therefore simplify the exposition by taking a course to be a single dose of strength $x \geq 0$, determined by the quantity of the active ingredient used in it. However, x can be thought of as proxying for any of the various ways of increasing the strength of a vaccine course that involve greater demand for the active ingredient.

We will model an *all-or-nothing* vaccine, providing perfect protection against infection, illness, transmission, and death to those for whom it works but no protection to those for whom it does not (Halloran, Haber, and Longini, Jr., 1992).[1] Whether the vaccine works for an individual is a Bernoulli random variable, identically and independently distributed in the population. Let $\phi(x) \in [0,1]$ denote the probability that the vaccine works for an individual, termed the vaccine's *efficacy*. As we do all the functions introduced in this model section, we take $\phi : [0,\infty) \to [0,1]$ to be real-analytic on its domain, meaning that it converges to its Taylor series expansion for any point in the domain.[2] Assume that a vaccine with no active ingredient is useless, i.e., $\phi(0) = 0$. We initially assume that stronger doses are more effective, i.e., $\phi'(x) > 0$ for all $x > 0$, although we allow for peaked efficacy functions in a later extension.

We take ϕ to be a *sigmoid* function. Informally, this means that the graph of ϕ is s-shaped. Formally, a sigmoid function is associated with a point x^i such that it is convex for all $x < x^i$ and concave for all $x > x^i$. For differentiable functions like ϕ, this definition is equivalent to its derivative $\phi'(x)$ having a unique mode at x^i. The given definition of a sigmoid function encompasses the two subscases shown in Figure 1. The panels on the left illustrate the subcase of a *strictly sigmoid* efficacy function, whose first derivative has a strictly positive mode, $x^i > 0$, which is also its inflection point. The panels on the right illustrate the subcase of a concave efficacy function.

[1] The alternative, what Halloran, Haber, and Longini, Jr. (1992) dub a *leaky* vaccine, provides the same imperfect level of protection to everyone vaccinated. For example, Farrington (2003) classifies the measles and rubella vaccines as all-or-nothing and the pertussis vaccine as leaky.

[2] The assumption that modeled functions are real-analytic is a technical condition, ensuring that if the optimal-dosing problem has a set of solutions, the set must be finite and thus have a well-defined minimum and maximium. Real-analytic functions are smooth (i.e., have derivatives of all orders), but the reverse is not always true.

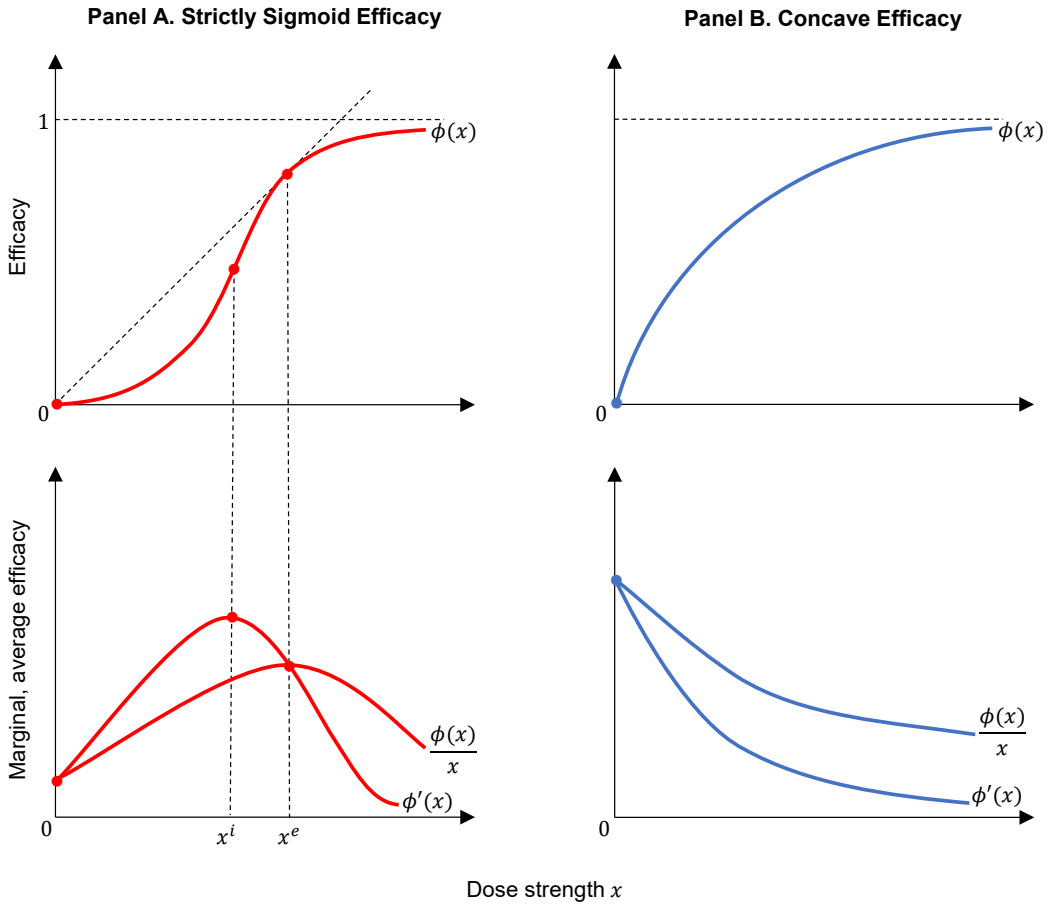


Figure 1: Subcases of a Sigmoid Efficacy Function. The panels on the left exhibit the case of strictly sigmoid ϕ, having an interior inflection point $x^i > 0$ and interior point of unit elasticity, $x^e > 0$. The panels on the right exhibit the case of concave ϕ. A concave ϕ can equivalently be thought of as a weakly sigmoid curve, the first derivative of which has a mode at 0.

This function still qualifies as sigmoid since its first derivative is unimodal. But it is a special case in that the mode of its first derivative is 0, the lower endpoint of its domain, rather than an interior point. The two subcases generate qualitatively similar results as long as vaccination involves nontrivial costs, which we will assume for most of the analysis. The qualitative results can diverge when vaccination is costless, a special case studied in Section 3.

Sigmoid dose-response curves are widely used in the pharmacometric modeling of vaccines (Rhodes et al., 2022), receiving both theoretical and empirical support in the scientific literature. The fact that the increase in efficacy is initially slow at low levels of the active ingredient is consistent with threshold effects in immune response. For example, a T-cell requires a cetain number of receptor signals to trigger cytokine response (Itoh and Germain, 1997) and proliferation (An-Yeung et al., 2014). That ϕ eventually plateaus reflects the saturation of the immune system (Borghans

et al., 1999) and homeostatic mechanisms such as regulatory T-cells that prevent excessive immune responses leading to autoimmune disease (Sakaguchi et al., 2008).[3] Empirical estimates of dose-response functions find a sigmoid shape fits the data well in both animal and human studies (Benest et al., 2020; de Jonge et al., 2020; Roozendaal et al., 2021).

Let $s(x) \geq 0$ denote the expected harm from side effects for an individual vaccinated with a dose of strength x. We will take the function $s : [0,\infty) \to [0,\infty)$ to be real-analytic on $[0,\infty)$. Assume that a vaccine with no active ingredient produces no side effects, i.e., $s(0) = 0$. Assume further that side effects are nondecreasing in dose strength at a nondecreasing rate: i.e., $s'(x) \geq 0$ and $s''(x) \geq 0$ for all $x \geq 0$.

The population consists of a continuum of homogeneous individuals of mass N. Letting $n \in [0,N]$ denote the mass of individuals receiving a vaccine course and $v \in [0,N]$ denote the mass of successfully immunized people, we have $v = n\phi(x)$. Denote the social benefit from preventing disease harm via successful vaccination by $B(v)$, capturing both the direct protective effects and the indirect epidemiological effects of reducing transmission from immunized individuals. While there is a natural reason to restrict the domain of the benefits function to $[0,N]$ since successfully vaccinating more than 100% of the population does not have a physical meaning, for technical reasons we will take $B : [0,\infty) \to [0,\infty)$ to be a real-analytic extension of the benefit function over the entire nonnegative real line.[4] Normalize benefits to 0 in the absence of vaccinations, i.e., $B(0) = 0$. Assume benefits are increasing and strictly concave, i.e., for all $v \geq 0$, $B'(v) > 0$ and $B''(v) < 0$. The assumption of concave benefits follows previous economic analyses of the market for epidemic vaccines. Ahuja et al. (2021) justify their use of a concave benefit function in their calibrations of the Covid-19 pandemic on the grounds that prioritizing initial doses to the most vulnerable (aged, immunocompromised, and health workers) generates steep initial returns, which fall as less vulnerable people are vaccinated, eventually reaching a plateau as the population reaches herd immunity.[5] We will assume the real-analytic extension satisfies $\lim_{v\uparrow\infty} B'(v) = 0$.

[3]Biology aside, as a mathematical principle, any increasing, bounded function converges to its limit from below, implying that it must eventually plateau.

[4]Our specification accommodates the possibility that vaccinating the last person in a 100% successful universal vaccination campaign has positive marginal benefit (a case that could arise if vaccination prevents individual harm but not transmission). Alternative extensions of B have drawbacks. One alternative extension would be to set $B'(N) > 0$ but $B'(v) = 0$ for all $v > N$. This alternative is not differentiable at N and thus not real-analytic on $[0,N]$. Another alternative would be to specify $B'(v) = 0$ for all $v \geq N$, but this restricts the value and shape of benefits $B(v)$ in a neighborhood below N. We address the concern that extending the benefit function above N could lead to planner solutions outside of the range with physical meaning by including constraint $n \leq N$ in the planner's problem.

[5]In their model of homogeneous consumers, Goodkin-Gold et al. (2024, Proposition 7) provide necessary and

A social planner chooses the vaccine's strength x and number of people vaccinated n to maximize the population's health—the social benefit from the vaccine less harm from side effects—net of expenditures. The planner's expenditures include spending on the vaccine's active ingredient (antigen), for which the planner pays an exogenous unit price $p \geq 0$.[6] This price may equal unit cost c or may reflect a margin $m \equiv (p-c)/p$ above c.

The planner also expends $k \geq 0$ per dose to administer the vaccine, independent of the dose's strength. This includes the cost of filling and finishing the vaccine vial, transporting the vials to dispensaries, and labor and facilities involved in administering a dose to a consumer. If the social planner internalizes the consumer's cost of traveling to the dispensary and the pain of being jabbed, those costs could be included in k as well.

The planner's objective is thus to maximize the value of its objective

$$B\big(n\phi(x)\big) - n\left[s(x) + px + k\right]. \tag{1}$$

Our specification nests the case of a perfectly benevolent planner who internalizes suppliers' profits. That case can be recovered by substituting $p = c$ into equation (1), which then becomes equivalent to social welfare. In practice, many health agencies do not internalize firm profits, so the model is specified to allow for that possibility.

Let $g(x) \equiv s(x) + px + k$ denote the planner's generalized cost of vaccinating an individual, embodying both the expected harm the individual experiences from side effects as well as fixed and variable expenditures on a dose. The function g inherits several key properties of s: namely, for all $x \geq 0$, it is weakly increasing, as $g'(x) = s'(x) + p \geq 0$; it is convex, as $g''(x) = s''(x) \geq 0$; and it is real-analytic, as its component functions are. The planner's objective (1) can then be rewritten more succinctly as

$$B\big(n\phi(x)\big) - ng(x). \tag{2}$$

sufficient conditions on the disease's effective reproductive number for the vaccine benefits to be concave.

[6] The choice to model p as a constant, independent of the quantity demanded by the vaccine program, can be justified on several grounds. On theoretical grounds, price-taking is relevant to the case in which the planner controls vaccine production and internalizes unit cost c as the relevant price p. Price-taking is also a reasonable assumption if the vaccine program is only a small part of the overall vaccine market. On practical grounds, during the Covid-19 pandemic, vaccine manufacturers appear to have been content to ration supplies to countries rather than letting rising prices clear the market. Castillo et al. (2021) found that vaccine courses (generally two doses) were being sold for between \$6 and \$40, compared to an estimated social value in the thousands of dollars. Vaccine manufacturers might fear that raising prices during a crisis would be judged repugnant (Kahneman, Knetsch, and Thaler, 1986) and generate a backlash.

The planner faces several constraints on this problem. Let $Q > 0$ denote the supply of the active ingredient. The total quantity used cannot exceed this supply:

$$nx \leq Q. \tag{3}$$

Further, no more than the total population can be vaccinated:

$$n \leq N. \tag{4}$$

Finally, the planner's choice variables must be nonnegative:

$$n, x \geq 0. \tag{5}$$

As usual, elasticities will play a key role in the analysis. Let $\eta_f(x) \equiv xf'(x)/f(x)$ denote the elasticity of an arbitrary function f at point x. In particular, $\eta_\phi(x)$ is the elasticity of efficacy at dosage x and $\eta_g(x)$ is the elasticity of generalized marginal cost. It will also be useful to have notation for a function's per-unit average, $\alpha_f(x) \equiv f(x)/x$. In particular, $\alpha_\phi(x)$ is the per-unit efficacy of dosage x and $\alpha_g(x)$ is the per-unit generalized cost.

3. Optimal Dosing with Costless Vaccination

We begin the analysis by abstracting from any costs of vaccination, setting $g(x) = 0$. This is a good place to begin for several reasons. First, the analysis in this case is relatively simple and generates clear results on fractional dosing. Second, this is the case carried into the calibrations for the Covid-19 pandemic in Section 6. That section derives the social benefit of vaccination from a complex epidemiological model and calibrates various functional forms for the efficacy function based on data for Covid-19 vaccines. However, that section takes a cruder approach to costs, setting $g(x) = 0$, in part because of a paucity of information on costs but mainly reflecting our view that the social costs of vaccination were eclipsed by the social benefits in the Covid-19 pandemic (see footnote 6).

3.1. Unit-Elastic Benchmark

The key benchmark with costless vaccination is provided by the dosage x^e at which efficacy is unit elastic. This point can be read off the graphs in Figure 1. For both functional forms, the origin provides one unit-elastic point. To see this, note that $\lim_{x\downarrow 0} \alpha_\phi(x) = \phi'(0)$ by l'Hôpital's Rule, impliying $\lim_{x\downarrow 0} \eta_\phi(x) = 1$ since $\eta_\phi(x) = \phi'(x)/\alpha_\phi(x)$. While $\alpha_\phi(x)$ and thus $\eta_\phi(x)$ are undefined at $x = 0$, having the $0/0$ form, it is natural to extend these functions to equal their right-hand limits at 0: thus, $\alpha_\phi(0) = \phi'(0)$ and $\eta_\phi(0) = 1$. According to the last equation, the origin provides one unit-elastic point on the efficacy function.

A strictly sigmoid efficacy function has another point of unit elasticity, the interior point where the chord from the origin is tangent to the curve. The analysis will focus on this nontrivial unit-elastic point when it exists but settle for the trivial point 0 when ϕ is concave and has no positive unit-elastic point, reflected in the following definition:

$$\bar{x}^e \equiv \max\left\{x \geq 0 \mid \eta_\phi(x) = 1\right\}. \tag{6}$$

Let X^e denote the set of unit-elastic points, i.e., $X^e = \{0\}$ for concave ϕ and $X^e = \{0, \bar{x}^e\}$ for strictly sigmoid ϕ. Let x^e be an arbitrary element of X^e.

The following lemma states that $\bar{x}^e$ is well defined and characterizes additional of its useful properties. The proofs of this and all subsequent results are provided in the appendix.

Lemma 1. *There exists a unique $\bar{x}^e \geq 0$ satisfying equation* (6)*. When ϕ is strictly sigmoid, $\bar{x}^e > x^i > 0$; when ϕ is concave, $\bar{x}^e = 0$. By definition, $\eta_\phi(\bar{x}^e) = 1$. For all $x \in (0, \bar{x}^e)$, $\eta_\phi(x) > 1$; for all $x > \bar{x}^e$, $\eta_\phi(x) < 1$.*

The unit-elastic dosage $\bar{x}^e$ provides a key benchmark for the analysis. As the next lemma states, $\bar{x}^e$ maximizes per-unit efficacy, $\alpha_\phi(x) \equiv \phi(x)/x$, which we will find is the optimal policy when vaccination is costless.

Lemma 2. *Per-unit efficacy $\alpha_\phi(x)$ is strictly quasiconcave. It is strictly increasing for $x \in (0, \bar{x}^e)$, strictly decreasing for $x > \bar{x}^e$, and reaches a maximum at $\bar{x}^e$.*

The relatively lengthy proof using only basic mathematics is supplemented with a shorter alternative leveraging a result in Karlin and Rubin (1956) that certain integral transforms of a function have a variation-diminishing property, experiencing no more sign changes than the function itself.

3.2. Optimal Vaccine Program

With the analysis of the unit-elastic benchmark x^e in hand, we can proceed to solve for the planner's optimum. Substituting $g(x) = 0$ into equation (2), with costless vaccination the objective function simplifies to

$$B\big(n\phi(x)\big). \tag{7}$$

It is evident that capacity constraint (3) binds so all capacity Q is used in an optimum. For any policy (n,x) that does not exhaust capacity, it can be improved upon by using some of the residual capacity to vaccinate the same individuals with a stronger dose. Stronger doses increase social benefits since $\phi(x)$ and $B(v)$ are increasing in their arguments but entail no greater social costs.

We thus have $n^*x^* = Q$. Rearranging this equation as $n^* = Q/x^*$ and substituting into (7) yields $B\big(Q\phi(x^*)/x^*\big) = B\big(Q\alpha_\phi(x^*)\big)$. Since B is increasing in its argument and Q is a constant, we see that the x^* should maximize per-unit efficacy $\alpha_\phi(x)$ and thus equal x^e by Lemma 2. This is the optimal dose unless unless it would leave some capacity unused even if all N individuals were vaccinated with it. In that case, the optimum simply vaccinates the entire population ($n^* = N$) with the largest possible dose subject to available capacity ($x^* = Q/N$). The following proposition synthesizes this analysis.

Proposition 1. *Suppose vaccination is costless to the planner, i.e., $g(x) = 0$. The optimal policy vaccinates $n^* = \min\{Q/\bar{x}^e, N\}$ individuals with dose $x^* = \max\{\bar{x}^e, Q/N\}$.*

The optimal policy can be visualized with the help of Figure 2. Focus first on the top panel, covering the subcase of strictly sigmoid ϕ. Given vaccination is costless and benefits and efficacy do not hit peaks, there is no largest optimal dose. For very high levels of capacity Q, the optimal dose vaccinates the whole population with the largest possible dose allowed by available capacity: $x^* = Q/N$. No matter how high Q is, the optimal dose can still be considered "fractional" relative to the optimal dose for yet higher capacity. Fractional dosing is the rule, not the exception.

As capacity shrinks from some initial high level, the optimal dose falls linearly until a kink is hit at $Q = N\bar{x}^e$. Below that threshold capacity, rather than continuing to shrink the size of the dose linearly, it is better for population health to keep the dose at the level $\bar{x}^e$ maximizing per-unit efficacy and ration it in the population. The unit-elastic dose $\bar{x}^e$ provides a floor on optimal dose size.

The lower panel covers the subcase of concave ϕ. In that subcase, $\bar{x}^e = 0$, so the unit-elastic

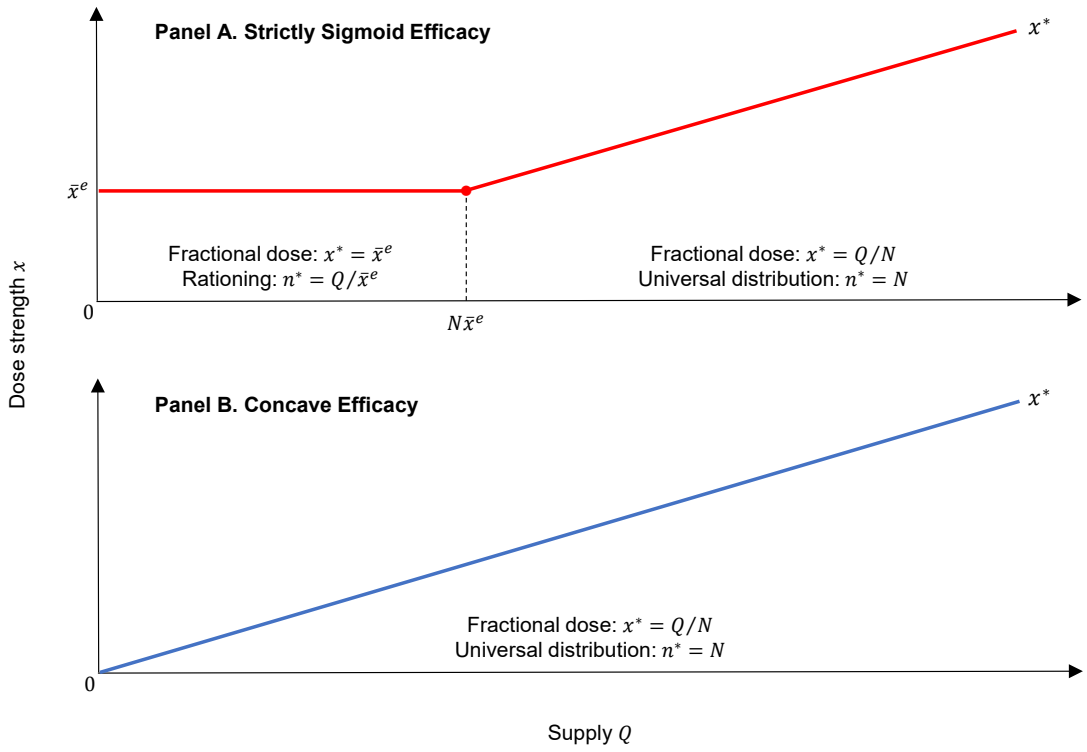


Figure 2: Optimal Dosage with Costless Vaccination. The panels show the planner's optimal vaccine policy for different shapes of a sigmoid efficacy function, strictly sigmoid in the top panel and concave in the bottom panel.

dose does not provide a floor on optimal dose size. Population health is maximized by spreading available capacity evenly in the population. Thus, the optimal dose continues to shrink linearly with capacity Q all the way to the origin.

A formal proof of Proposition 1 is provided in the appendix. The planner's problem of choosing x and n to maximize objective (7) subject to inequality constraints (3)–(5) can be solved using the Karush–Kuhn–Tucker (KKT) method. In the appendix, we prove that the problem satisfies linear independence constraint qualification (LICQ), a textbook condition ensuring the KKT conditions are necessary for an optimum (Solodov, 2011). The KKT conditions narrow the optimum down to two potential solutions—an interior solution involving unit-elastic dose $\bar{x}^e$ and a corner solution—which can be compared. The interior solution solves the relaxed problem omitting a constraint and thus provides the optimum in the original problem unless it violates the omitted constraint, in which case the corner solution provides the optimum.

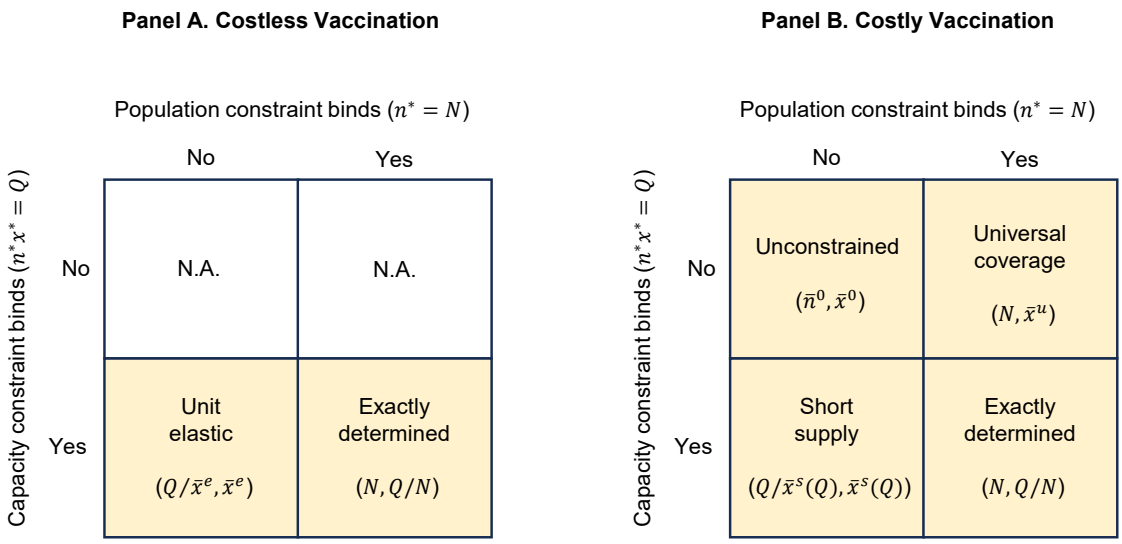


Figure 3: Benchmarks Comprising Optimal Solution. The planner's optimum, solved using the KKT method, is comprised of the benchmarks in the matrix entries, each characterized in a subsequent subsection.

4. Optimal Dosing with Costly Vaccination

In the case of costless vaccination just studied, there is no reason not to use all available capacity, so the capacity constraint (3) always binds. This reduces the complexity of the optimal policy. As the first panel of Figure 3 indicates, the optimal policy can take one of only two forms depending on whether the remaining constraint—the population constraint (4)—binds. If the population constraint does not bind, the optimal solution is interior, involving the unit-elastic dose x^e; if the population constraint bids, the two variables in the optimal policy n^* and x^* are exactly determined by the two equations provided by the binding constraints.

This section studies the case of costly vaccination, assuming $p, k > 0$ and $s(x) \geq 0$, implying $g(x) > 0$ for all $x > 0$. Now there is a reason not to use all available capacity. The decreasing marginal social benefit of employing the last units of capacity may not justify the increasing marginal social cost. The optimal policy becomes more complex, taking on four possible forms delineated in the second panel of Figure 3, depending on whether capacity constraint and/or the population constraint bind. Each of the four boxes provides a benchmark characterized in a series of subsections below. These benchmarks become building blocks in constructing the planner's optimum, derived using the KKT method.

4.1. Nontrivial Policy

With costless vaccination, the optimal policy is always nontrivial, i.e., involving positive vaccinations ($n^* > 0$) and positive dosage ($x^* > 0$). Indeed, it is optimal to allocate the whole capacity Q across some or all individuals in the population. By contrast, the optimal policy may be trivial with costly vaccination—namely when the cost of every nontrivial vaccine policy exceeds its benefit. The following lemma provides a necessary and sufficient condition for a nontrivial optimum under costly vaccination.

Lemma 3. *The planner's optimum is nontrivial if and only if* $B'(0)\phi(x) > g(x)$ *for some* $x > 0$.

To gain some intuition for the proof of the lemma, suppose $B'(0)\phi(x) > g(x)$ for some $x > 0$. Then a policy of vaccinating a few people with that dosage generates a greater health benefit than generalized cost. This policy thus provides the planner with positive surplus, strictly dominating the trivial policy of not vaccinating anyone. On the other hand, suppose $B'(0)\phi(x) \leq g(x)$ for all $x > 0$. Then there is no way for the planner to obtain positive surplus by vaccinating a few people. Vaccinating more people would be even worse for the planner because the benefit function B is concave. So the trivial policy of not vaccinating anyone is optimal in that case.

The condition for a nontrivial optimum stated in Lemma 3 is rather amorphous and may be difficult to verify. We can provide a more pointed, equivalent condition; but since it requires notation developed in the next subsection, we postpone its statement until then.

4.2. Unconstrained Benchmark

The first benchmark we analyze involves maximizing the planner's objective (2) removing the capacity and population constraints, retaining only the nonnegativity constraints:

$$\max_{n,x\geq 0} \left[B\big(n\phi(x)\big) - ng(x)\right]. \tag{8}$$

Denote the solution by (n^0, x^0). In this subsection, we provide a characterization of this solution and show it is well defined under the condition in Lemma 3 ensuring the planner's optimum is nontrivial.

Considerable insight can be gained by examing the problem that is equivalent to (8) but makes the change of variables $v = n\phi(x)$ and solves for the optimal number of successful immunizations

v rather than the optimal number of vaccinations n. Equation (8) becomes

$$\max_{v,x\geq 0}\left[B(v)-v\cdot\frac{g(x)}{\phi(x)}\right]. \tag{9}$$

It is immediate that the unconstrained-optimal dose x^0 maximizes the ratio $\phi(x)/g(x)$. This ratio is not guaranteed to have a unique maximizer (though it will in many applications). Let X^0 denote the set of maximizers,

$$X^0 \equiv \operatorname*{argmax}_{x\geq 0}\left[\frac{\phi(x)}{g(x)}\right], \tag{10}$$

and x^0 be an arbitrary element of X^0. The analysis will often focus on the largest maximizer: $\bar{x}^0 \equiv \max(X^0)$. The largest maximizer forms a solution pair $(\bar{n}^0,\bar{x}^0)$ most likely to respect the omitted constraints and thus provide a feasible solution to the unconstrained problem (a detail covered in the relevant appendix proof).

Taking the first-order condition with respect to x in equation (10) and rearranging yields a necessary condition for any maximizer x^0:

$$\eta_\phi(x^0)=\eta_g(x^0). \tag{11}$$

The optimal dose in the unconstrained problem balances the percentage increase in efficacy against the percentage increase in generalized cost entailed by a dose increase. Note that the social benefit generated by the vaccine does not appear in this equation, only its efficacy. Below, we will see that the social benefit factors into the determination of the optimal number of vaccinations n^0, but the optimal dose they are vaccinated with only depends on efficacy and generalized cost.

The preceding sketch of the analysis glossed over some technical details. It remains to be verified that X^0 is nonempty and finite so that its largest element $\bar{x}^0$ is well defined. The following lemma fills in those technical details. It also provides a useful result, positioning $\bar{x}^0$ relative to the unit-elastic benchmark x^e: if $\bar{x}^0 \geq x^e$, then X^0 is a singleton. This result follows since if any solution x^0 lies above x^e, it lies above the point of inflection for ϕ, in the range where ϕ is a well-behaved concave function, generating a unique optimum.

Lemma 4. *There exists a unique* $\bar{x}^0>0$ *satisfying* $\bar{x}^0=\max(X^0)$. *If* $\bar{x}^0\geq x^e$, *then* $X^0=\{\bar{x}^0\}$.

The ratio of the social benefit to cost of vaccinating the nth individual with the unconstrained optimal dose $\bar{x}^0$ will play an important role in the subsequent analysis, so we introduce notation

for it:

$$\rho^0(n) \equiv \frac{B'\big(n\phi(\bar{x}^0)\big)\phi(\bar{x}^0)}{g(\bar{x}^0)}. \tag{12}$$

When $\rho^0(n) > 1$, a case can be made for vaccinating the nth person—at least in a counterfactual unconstrained world. The inequality tells us that the social benefit exceeds the cost of vaccinating the nth person with the dose $\bar{x}^0$ that, like all elements of X^0, has the greatest potential to generate higher benefits than costs. That is not to say that $\bar{x}^0$ will be the optimal dose when supply and population constraints operate or that vaccinating the nth person will be optimal then. The inequality provides a necessary but not sufficient condition for the nth person to be vaccinated in the planner optimum. When $\rho^0(n) \leq 1$, on the other hand, vaccinating the nth person cannot be strictly optimal.

The derivative of (12) with respect to n is

$$\frac{B''\big(n\phi(\bar{x}^0)\big)\phi(\bar{x}^0)^2}{g(\bar{x}^0)}, \tag{13}$$

which is negative since $B'' < 0$. This negative derivative implies that the case for vaccinating the next individual is increasingly hard to make as more individuals are vaccinated. The marginal social cost of vaccinating each individual remains the same but the marginal social benefit falls.

We can put this notation to immediate use in restating the condition for a nontrivial planner optimum more succinctly.

Lemma 5. *The planner's optimum is nontrivial if and only if* $\rho^0(0) > 1$.

Intuitively, if $\rho^0(0) > 1$, a case can be made for vaccinating a few individuals with dose $\bar{x}^0$. Whatever the available supply and whatever the dosage, there is always enough to vaccinate a sufficiently small number of individuals with that dosage. If $\rho^0(0) \leq 1$, on the other hand, it is not socially worthwhile to vaccinate even the first individual, let alone an individual added to the n that have already been vaccinated.

It remains to derive n^0. Substituting x^0 into (9) and taking the first-order condition with respect to v yields $B'(v^0) = g(x^0)/\phi(x^0)$. Changing back to the original choice variables, $v^0 = n^0\phi(x^0)$, and solving for n^0 yields

$$n^0 = \frac{1}{\phi(x^0)}(B')^{-1}\left(\frac{g(x^0)}{\phi(x^0)}\right). \tag{14}$$

If the planner's optimum is nontrivial, it can be shown that the inverse $(B')^{-1}$ exists and thus

equation (14) has a solution, as the next lemma states.

Lemma 6. *As long as the planner's optimum is nontrivial, there exists a unique* $n^0 > 0$ *satisfying* (14).

When X^0 includes several maximizers, different $x^0 \in X^0$ pair with different n^0 solving equation (14). In particular, larger x^0 are associated with smaller n^0.[7] Let $\bar{n}^0$ denote the solution to equation (14) when $x^0 = \bar{x}^0$.

4.3. Universal-Coverage Benchmark

We next analyze what we call the universal-coverage benchmark, maximizing the planner's objective (2) given the entire population must be vaccinated:

$$\max_{x \geq 0} \left[B\big(N\phi(x)\big) - Ng(x) \right]. \tag{15}$$

As before, the maximizer of this objective is not guaranteed to be unique. Let X^u denote the set of maximizers,

$$X^u \equiv \operatorname*{argmax}_{x \geq 0} \left[B\big(N\phi(x)\big) - Ng(x) \right], \tag{16}$$

and x^u denote an arbitrary element of X^u. While we focused on the largest maximizer in the previous benchmark, here we will switch the focus to the smallest maximizer $\underline{x}^u \equiv \min(X^u)$. The smallest maximizer relaxes the omitted capacity constraint to the extent possible and thus is the most likely maximizer to constitute a feasible solution.

It is not a foregone conclusion that the objective function (15) has a nontrivial maximizer and if there is a set of maximizers, that this set is finite so has a well-defined largest element. The next lemma provides a condition guaranteeing these conclusions are true.

Lemma 7. *If* $\rho^0(N) \geq 1$*, there exists a unique* $\underline{x}^u > 0$ *satisfying* $\underline{x}^u = \min(X^u)$.

The proof, while similar to that of Lemma 4, is not identical, so we provide a separate proof in the appendix.

Taking the first-order condition and rearranging, we have that any maximizer x^u must satisfy

$$B'\big(N\phi(x^u)\big)\phi'(x^u) = g'(x^u). \tag{17}$$

[7] To see this, note that the ratio in the argument of $(B')^{-1}$ is independent of which x^0 is selected because all of them maximize the ratio $\phi(x)/g(x)$, but the leading factor $1/\phi(x^0)$ is decreasing in x^0.

The left-hand side is the marginal social benefit (per vaccinated individual) of increasing dose strength. The optimal dose in the universal-coverage benchmark equates this marginal social benefit with the marginal generalized cost of that dosage increase on the right-hand side of the equation.

4.4. Short-Supply Benchmark

The next benchmark is the mirror image of the one just analyzed. The universal-coverage benchmark maximized the planner's objective omitting the constraint on supply Q, fixing the number of people to be vaccinated N. The benchmark analyzed here drops the constraint on the number of vaccinations and instead fixes the supply of vaccine substance that must be utilized at Q. We will call this the short-supply benchmark. The label anticipates the role this benchmark will play in the planner's optimum, providing the formula for the optimum when supplies Q become so tight that the supply constraint (3) takes over as the only active constraint; supplies are too constrained to consider vaccinating anywhere near the total population, allowing us to safely ignore constraint (4) on the number of vaccinations.

The binding constraint that Q must be fully utilized entails $nx = Q$. Substituting $n = Q/x$ into the planner's objective function (2), the maximization problem can be written

$$\max_{x \geq 0} \left[B\big(Q\alpha_\phi(x)\big) - Q\alpha_g(x) \right]. \tag{18}$$

Let $X^s(Q)$ denote the set of maximizers,

$$X^s(Q) = \operatorname*{argmax}_{x \geq 0} \left[B\big(Q\alpha_\phi(x)\big) - Q\alpha_g(x) \right], \tag{19}$$

and $x^s(Q) \in X^s(Q)$ denote an arbitrary maximizer. For this benchmark, the analysis will focus on the largest maximizer, $\bar{x}^s(Q) \equiv \max[X^s(Q)]$, since this relaxes the omitted population constraint to the extent possible and thus is the most likely maximizer to constitute a feasible solution. We have emphasized the dependence of these maximizers on Q by adding it as an argument; benchmarks analyzed in previous subsections did not include a supply constraint, so those benchmark optima are independent of Q.

If Q is large, there is no guarantee forcing the population to consume that much vaccine substance will generate positive surplus. If Q is sufficiently small, however, allocating that supply to the population in doses of size $\bar{x}^s(Q)$ is guaranteed to generate positive surplus if the planner's

optimum is nontrivial, as the next lemma states.

Lemma 8. *Suppose the planner's optimum is nontrivial. For sufficiently small $Q > 0$, there exists a unique $\bar{x}^s(Q) > 0$ satisfying $\bar{x}^s(Q) = \max[X^s(Q)]$ and delivering a positive value of the objective* (18).

Taking the first-order condition in equation (18) and rearranging, all solutions $x^s(Q)$ must satisfy

$$B'\big(Q\alpha_\phi(x^s(Q))\big)\,\alpha'_\phi\big(x^s(Q)\big) = \alpha'_g\big(x^s(Q)\big). \tag{20}$$

This rearranged first-order condition resembles the analogous one for the universal-coverage benchmark in (17). Without a constraint on supplies, there, the planner compared marginal efficacy to marginal generalized cost *in level terms*. Here, the fixed supply (Q) forces the planner to compare marginal efficacy to marginal generalized cost *in per-unit terms*, as the inclusion of the factors $\alpha'_\phi(x)$ and $\alpha'_g(x)$ suggests.

It is natural to think that a reduction in Q would reduce $\bar{x}^s(Q)$: a tighter supply constraint would be accommodated by smaller fractional doses. In fact, the comparative-statics is more complicated: $\bar{x}^s(Q)$ may be increasing, constant, or decreasing in Q. The next proposition provides a sufficient statistic for the direction of the comparative-static effect, depending on whether generalized cost is more or less elastic than efficacy where efficacy is unit elastic.[8]

Proposition 2. $\operatorname{sgn}\big(d\bar{x}^s(Q)/dQ\big) = \operatorname{sgn}\big(\eta_\phi(\bar{x}^e) - \eta_g(\bar{x}^e)\big) = \operatorname{sgn}\big(1 - \eta_g(\bar{x}^e)\big)$.

We have accorded this comparative-static result the status of a proposition since, as we will see in the next subsection, it dictates the shape of the overall planner optimum at low levels of Q. It is therefore worth developing some intuition this result. Note that no matter how the planner responds to a reduction in Q—whether reducing n, x, or both—it cannot prevent the level of effective vaccination v from falling when Q falls, leading to an increase in marginal benefit $B'(v)$ since B is concave. This inclines the planner to wring more efficacy out of the limited supply by moving the dose toward the unit-elastic point $\bar{x}^e$, which maximizes per-unit efficacy, $\alpha_\phi(x)$. Under severe capacity constraints, the planner is willing to accept greater marginal harm to vaccinated individuals to avoid an even greater reduction in the social benefit of effective vaccinations. Whether

[8] The sign function in the statement of the proposition has the standard definition:

$$\operatorname{sgn}(x) = \begin{cases} 1 & x > 0 \\ 0 & x = 0 \\ -1 & x < 0. \end{cases}$$

moving toward $\bar{x}^e$ entails an increase or decrease in x depends on whether $\bar{x}^s(Q)$ is greater or less than $\bar{x}^e$, which in turn depends on the relative elasticities of g and ϕ. When $\eta_g(\bar{x}^e) > \eta_\phi(\bar{x}^e)$, the benchmark dose $\bar{x}^s(Q)$ is kept below the unit-elastic point $\bar{x}^e$ to avoid generating too much marginal harm. In this case, moving $\bar{x}^s(Q)$ toward $\bar{x}^e$ entails an increase in dose size, explaining the surprising comparative-static result.

If $\eta_g(\bar{x}^e) < \eta_\phi(\bar{x}^e)$, then $\bar{x}^s(Q) > \bar{x}^e$. In that case, moving $\bar{x}^s(Q)$ toward $\bar{x}^e$ in response to a reduction in Q entails a decrease in dose size, the anticipated comparative-static result. In the edge case where $\eta_g(\bar{x}^e) = \eta_\phi(\bar{x}^e)$, $\bar{x}^s(Q) = \bar{x}^e$, where it remains as Q falls.

In a variety of important cases, it is straightforward to bound $\eta_g(\bar{x}^e)$ and make definitive claims about when the surprising comparative-static result emerges.

Proposition 3. *Suppose ϕ is concave. Then $d\bar{x}^s(Q)/dQ > 0$.*

Proposition 4. *Suppose the side-effect function is linear, i.e., $s(x) = \sigma x$ for some $\sigma \geq 0$. Then $\partial \bar{x}^s(Q)/\partial Q > 0$.*

The proofs of these two propositions are almost immediate from Proposition 2. To see the first result, if ϕ is concave, then $\bar{x}^e = 0$ by Lemma 1. Then $\eta_g(\bar{x}^e) = \eta_g(0) = 0/k = 0$. Therefore, $1 - \eta_g(\bar{x}^e) = 1 > 0$, implying $d\bar{x}^s(Q)/dQ > 0$ by Proposition 2. To see the second result, if $s(x) = \sigma x$, then $1 - \eta_g(\bar{x}^e) = k/g(\bar{x}^e)$, which is positive for all $k > 0$, implying implying $d\bar{x}^s(Q)/dQ > 0$ by Proposition 2.

According to Propositions 3 and 4, the surprising comparative-static result $\partial \bar{x}^s(Q)/\partial Q < 0$ never emerges when ϕ is concave or $s(x)$ is linear. Evidently, the complex interplay between the strictly sigmoid shape of ϕ and the convexity of $s(x)$ is key to generating the surprising result.

When $s(x)$ is strictly convex, the sign of $d\bar{x}^s(Q)/dQ$ can be tied to the level of the fixed vaccination cost, k.

Proposition 5. *Suppose the side-effect function is strictly convex, i.e., $s''(x) > 0$ for all $x > 0$. Then*

$$\operatorname{sgn}\left(\frac{d\bar{x}^s(Q)}{dQ}\right) = \operatorname{sgn}(k - \hat{k}), \tag{21}$$

where $\hat{k} \equiv \bar{x}^e s'(\bar{x}^e) - s(\bar{x}^e)$. If ϕ is concave, then $\hat{k} = 0$; if ϕ is strictly sigmoid, then $\hat{k} > 0$.

The proof is again almost immediate from Proposition 2. Using the expression for $\hat{k}$ from the statement of the proposition, for any $s(x)$, we can write $1 - \eta_g(\bar{x}^e) = (k - \hat{k})/g(\bar{x}^e)$. Hence, by Proposition 2, the sign of $k - \hat{k}$ determines the sign of $d\bar{x}^s(Q)/dQ$.

Regarding the value of $\hat{k}$, if ϕ is concave, then $\bar{x}^e = 0$ by Lemma 1, implying $\hat{k} = 0$. In that case, $k - \hat{k}$ is positive for all $k > 0$, implying $d\bar{x}^s(Q)/dQ > 0$ for all $k > 0$, reprising the result for concave ϕ from Proposition 3. If ϕ is strictly sigmoid and $s(x)$ is strictly convex, then

$$s(\bar{x}^e) = \int_0^{\bar{x}^e} s'(x)dx < \int_0^{\bar{x}^e} s'(\bar{x}^e)dx = \bar{x}^e s'(\bar{x}^e), \tag{22}$$

implying $\hat{k} > 0$. In that case, there is a nontrivial constellation of parameters, namely $k \in (0, \hat{k})$, for which the surprising comparative-static result $d\bar{x}^s(Q)/dQ < 0$ emerges. The surprising comparative-static result emerges if k is sufficiently small; if k is large, the comparative-static effect goes in the anticipated direction. High values of k "wash out" the convexity of $s(x)$ required for the surprising comparative static to emerge.

4.5. Optimal Vaccine Program

The various benchmarks just analyzed can be combined to characterize the planner's optimum under costly vaccination, as stated in the next proposition. The proof provided in the appendix derives the optimum from the associated KKT conditions.

Proposition 6. *Suppose the planner's optimum is nontrivial. It can be characterized by one of two cases, depending on* $\rho^0(N)$.

- *When* $\rho^0(N) > 1$, $n^* = \min\{N, Q/\bar{x}^s(Q)\}$ *and*

$$x^* = \begin{cases} \underline{x}^u & Q \geq N\underline{x}^u \\ Q/N & Q \in (N\bar{x}^s(Q), N\underline{x}^u) \\ \bar{x}^s(Q) & Q \leq N\bar{x}^s(Q). \end{cases}$$

- *When* $\rho^0(N) \leq 1$, $n^* = \min\{\bar{n}^0, Q/\bar{x}^s(Q)\}$ *and* $x^* = \max\{\bar{x}^0, \bar{x}^s(Q)\}$.

The dense notation in the proposition obscures the rather straightforward structure of the optimal policy, which can be visualized with the help of Figure 4. Start with the upper panel, showing the optimal policy when $\rho^0(N) > 1$. This inequality ensures that the planner prefers to vaccinate the whole population if supplies allow. For sufficiently high supply Q in this panel, there is no stretching or rationing. The entire population is vaccinated with the universal-vaccination benchmark dosage, $\bar{x}^u$.

As Q falls below $N\bar{x}^u$, the capacity constraint begins to bind. The optimal dose is initially stretched so all individuals can still be vaccinated. For sufficiently tight capacity, below threshold

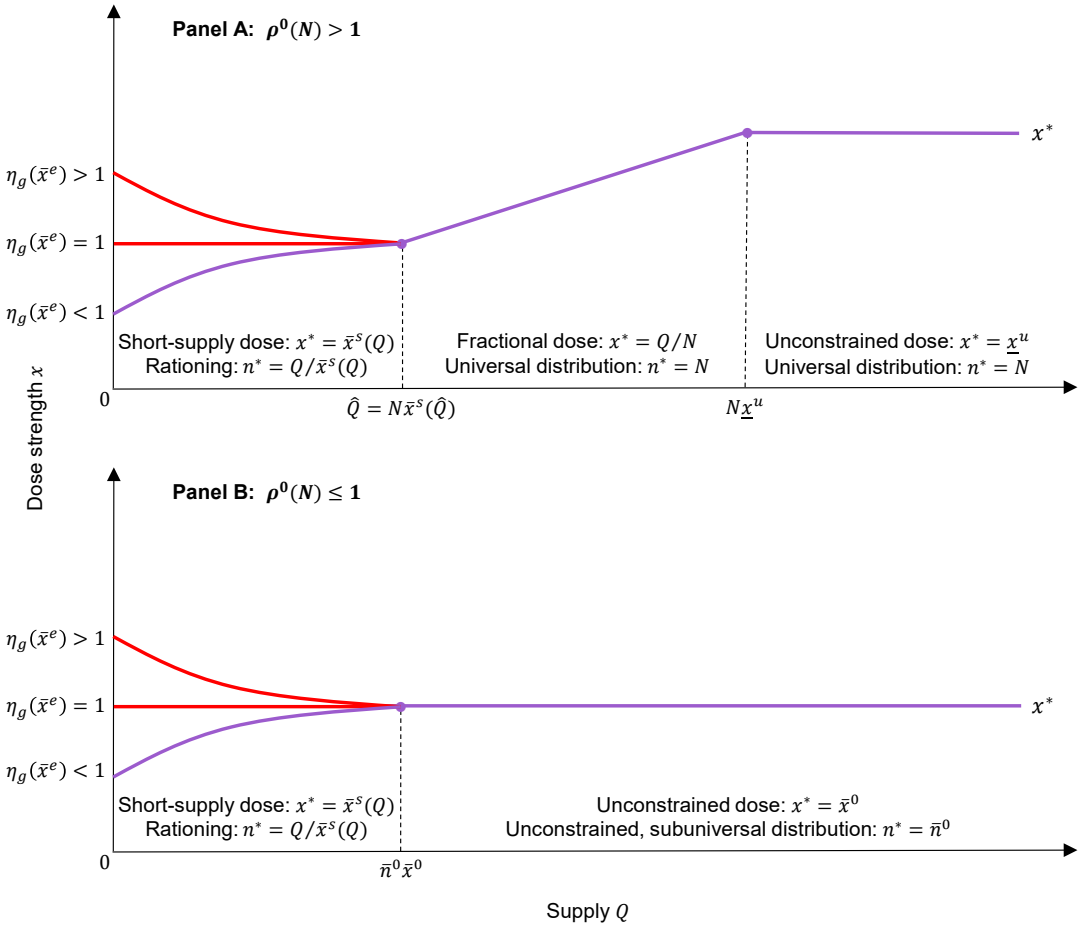


Figure 4: Optimal Dosage for Costly Vaccination. The panels show the planner's optimal vaccine policy for various values of $\rho^0(N)$. Unlike previous figures, here, panels distinguish between different parameter values, not different shapes for ϕ. In the interval closest to the origin, the optimal dose corresponds to the short-supply dose $\bar{x}^s(Q)$, which may slope up, down, or be constant. The figure shows three example slopes, labeled with the necessary and sufficient condition from Proposition 2 for that slope. Purple curve segments (a blend of the colors used in previous figures) apply to both concave and strictly sigmoid ϕ; blue segments apply only to strictly sigmoid ϕ.

$\hat{Q}$, the optimum switches to the short-supply dose, $\bar{x}^s(Q)$. As we saw in Proposition 2, the comparative statics here are ambiguous: $\bar{x}^s(Q)$ may increase, decrease, or remain constant as Q falls, depending on whether generalized cost is more or less elastic than efficacy at $\bar{x}^e$. These comparative statics are inherited by the optimal dosage x^* since $x^* = \bar{x}^s(Q)$ in this region. It is possible that x^* increases as Q falls. According to Propositions 3–5, the necessary and sufficient conditions for this outcome are that ϕ is stricly sigmoid, s is convex, and k is sufficiently low.

Extreme examples can be constructed in which the curve turns so sharply upward as Q approaches 0 that $\bar{x}^s(Q)$ is higher in that limit than the unconstrained dose $\bar{x}^u$. In such extreme cases, x^* is highest when supply constraints are the tightest.

The lower panel of Figure 4 illustrates the case of a low benefit/cost ratio, $\rho^0(N) \leq 1$. In this case, for sufficiently high supply Q, unconstrained doses are distributed, but to less than the entire population. It is optimal to leave some individuals unvaccinated not because of scarce supplies but because the marginal benefit of vaccinating them is less than the marginal generalized cost. The optimal vaccine policy unconstrained by supplies or population happens to involve an interior number of vaccinations. As capacity is tightened and Q falls below $\bar{n}^0\bar{x}^0$, the supply constraint starts to bind, and the optimal policy switches to short-supply dosing $\bar{x}^s(Q)$. Decreasing Q in this interval can increase, decrease, or leave the optimal dose constant as in the upper panel of the figure.

5. Extensions

This section analyzes two extensions to the original model. First, we consider an alternative efficacy function that has an interior peak instead of being everywhere increasing. Second, we consider allow for uncertainty in the level of efficacy at the time the planner sets the dosage.

5.1. Peaked Efficacy

In the model analyzed thus far, we took efficacy to be an increasing function over its whole domain. Marginal efficacy remained positive for any finite dose, no matter how large. In this subsection, we consider an alternative model of peaked efficacy. In this alternative illustrated in Figure 5, efficacy increases up to a peak, denoted x^p, and decreases for larger doses.

We have several motives for analyzing peaked efficacy. First, it is consistent with some vaccinology studies. Experiments with animal models found statistically significant evidence of a peak in the dose-response functions (Pompano et al., 2014; Rhodes et al., 2016; Billeskov et al., 2018). Second, we will adopt peaked efficacy functions in our calibrations. As we will see, the peak x^p provides a well-defined notion of a full dose, facilitating comparisions in the calibration section. Before turning to the calibrations, it is useful to work out the implications of peaked efficacy functions in the simpler model here first.

The following proposition characterizes the planner optimum integrating peaked efficacy into the model with costless vaccination.

Proposition 7. *Suppose vaccination is costless to the planner, i.e., $g(x) = 0$, and efficacy has a peak at x^p. The optimal policy vaccinates $n^* = \min\{Q/\bar{x}^e, N\}$ individuals with dose $x^* =$*

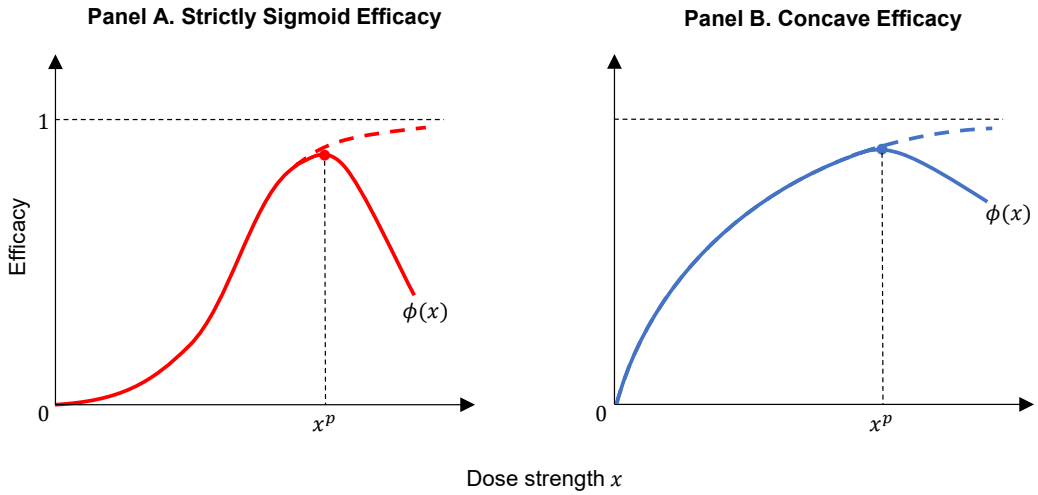


Figure 5: Peaked Efficacy Function. In the original model, ϕ continues to increase for the largest doses, illustrated by the dashed curves. The alternative model in this section involves peaked ϕ, which turns down after reaching peak x^p, experiencing reductions in efficacy beyond that dosage.

$\max\{\bar{x}^e, \min\{Q/N, x^p\}\}$.

Figure 6 helps visualize how the planner optimum under costless vaccination changes with peaked efficacy. The optimum is the same as in the original model except for the imposition of a cap x^p on the optimal dose. The planner has no reason to go beyond x^p since this only reduces social surplus by reducing efficacy.

As straightforward as it was to integrate peaked efficacy into the model with costless vaccination, integrating it into the model with costly vaccination is even more so. The positive marginal generalized cost of larger doses already provided a reason for the planner to endogenously cap the dose below where efficacy flattens out. Thus, the optimum prescribed in Proposition 6 never ventures into dosages where efficacy peaks and turns down. The optimal policy is the same whether efficacy is everywhere-increasing or peaked, as the next proposition states.

Proposition 8. *Suppose vaccination is costless to the planner, i.e., $g(x) > 0$, and efficacy has a peak at x^p. The optimal policy remains unchanged from Proposition 6.*

5.2. Uncertain Efficacy

In this subsection, we extend the model to allow for residual uncertainty in the efficacy function when the optimal dosage is set. This formulation of the model captures the reality that manufacturers seldom run late-stage clinical trials for more than one dosage. Experimentation on dosage,

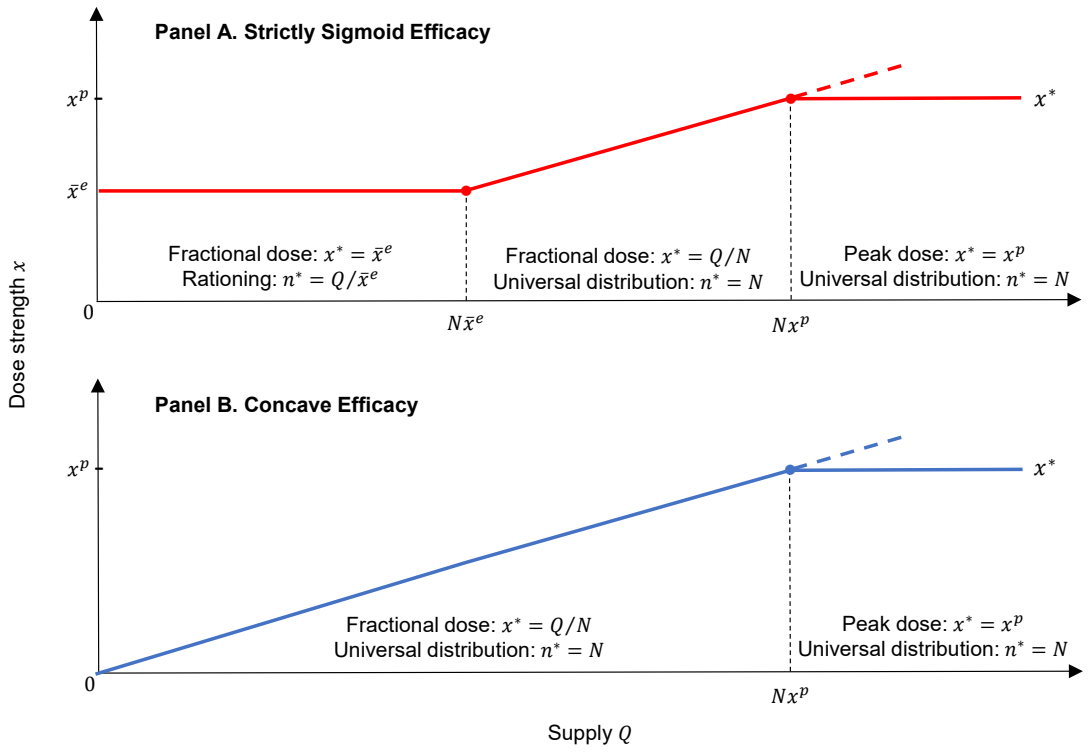


Figure 6: Optimal Dosage with Costless Vaccination and Peaked Efficacy. Moving from everywhere-increasing to peaked efficacy caps the planner optimum at x^p but otherwise leaves it unchanged. The dashed line shows the optimum in the original model with everywhere-increasing efficacy for comparision.

if done at all, is restricted to low-powered, early-stage trials and tends to be limited, only investigating a handful of alternative doses. This was especially true in the rush to develop Covid-19 vaccines (Dunn et al., 2023).

Respecify the efficacy function as $\phi(x,\theta)$, adding a second argument θ representing the unknown state of the world when dosage x is set. Assume high values of θ represent "good news" about efficacy and marginal efficacy: letting subscripts denote partial derivatives with respect to those arguments, we have $\phi_\theta > 0$ and $\phi_{x\theta} > 0$. Once x is set, assume it is difficult to change, perhaps requiring new clinical trials or adjustments to manufacturing processes. Ex post, clinical trials and medical experience with chosen dosage x may reveal information about θ that is not available ex ante, when x is set. Assume θ is randomly distributed on support $[0,\infty)$ with probability density function $f(\theta)$ and cumulative distribution function $F(\theta)$.

With the addition of uncertainty complicating the analysis, we seek to streamline the analysis by considering the simplest variant of the model along other dimensions. In particular, throughout this subsection, we will maintain the assumptions that efficacy is concave in x, i.e., $\phi_{xx} < 0$, and

that vaccination is costless, i.e., $g(x) = 0$. Under these assumptions, were efficacy not peaked, the optimal dosage would simply be $x^* = Q/N$ with or without uncertainty over θ. Uncertainty would play no role in the analysis.

Uncertainty does play a role when efficacy is peaked, motivating us to adopt this the variant of the model throughout this subsection. Intuitively, the optimal dose set under ex ante uncertainty x^* may turn out to be too strong in the light of further information about θ. If it is possible to adjust x^* ex post, moving to a fractional dose could increase efficacy. Thus, uncertainty provides another rationale for fractional dosing apart from epidemic supply constraints.

To integrate uncertainty and peaked efficacy in the same model, assume for all $\theta \in (0,\infty)$, there exists $x^p(\theta) > 0$ such that $\phi_x\big(x^p(\theta),\theta\big) = 0$. Totally differentiating with respect to θ and rearranging yields

$$\frac{dx^p(\theta)}{d\theta} = -\frac{\phi_{x\theta}\big(x^p(\theta),\theta\big)}{\phi_{xx}\big(x^p(\theta),\theta\big)} > 0. \tag{23}$$

Since $\phi_{xx} < 0$, we have $\phi_x(x,\theta) > 0$ for all $x \in \big[0, x^p(\theta)\big)$ and $\phi_x(x,\theta) < 0$ for all $x > x^p(\theta)$.

To highlight the possibility that uncertainty can provide a rationale for fractional dosing apart from supply constraints, we will abstract from any constraint on Q in this analysis. In the absence of a constraint on supply, since vaccination is costless and more vaccinations produce a positive benefit—i.e., $B'(v) > 0$ for all $v \geq 0$—it is immediate that the planner's optimum involves vaccinating as many individuals as possible: $n^* = N$. The optimal ex ante dose x^* maximizes the expected social benefit of the universal vaccine campaign,

$$\max_{x \geq 0} \left[\int_0^\infty B\big(N\phi(x,\theta)\big) f(\theta) d\theta \right], \tag{24}$$

satisfying the first-order condition

$$\int_0^\infty B'\big(N\phi(x^*,\theta)\big)\phi_x(x^*,\theta) f(\theta) d\theta = 0. \tag{25}$$

Since $B' > 0$ and $f > 0$, the left-hand side of (25) can only equal 0 if it integrates some positive and negative values of ϕ_x. Moreover, since $\phi_{x\theta} > 0$, equation (25) implies the existence of $\hat{\theta}(x^*) > 0$ such that $\phi_x\big(x^*, \hat{\theta}(x^*)\big) = 0$, $\phi_x(x^*,\theta) < 0$ for all $\theta \in \big(0, \hat{\theta}(x^*)\big)$, and $\phi_x(x^*,\theta) > 0$ for all $\theta > \hat{\theta}(x^*)$.

Focusing on those states of the world θ below $\hat{\theta}(x^*)$, we have that the ex ante optimal dose turns out to be beyond the efficacy peak in those states: formally, $x^* > x^p(\theta)$. If the true state θ

were learned and it was possible to reduce dosage ex post, a fractional dose would improve social welfare in those states. The probability that the ex ante optimal dose turns out overly strong equals $F\left(\hat{\theta}(x^*)\right)$. If the efficacy peak were known, it would be optimal to set x^* to that peak. Given the peak is uncertain when x^* is set in the model, the optimal value under uncertainty balances undershooting the peak versus overshooting, and thus optimally entails a positive probability of overshooting. It may not be possible to adjust x^*, but if it is, there is a positive probability $F\left(\hat{\theta}(x^*)\right)$ that moving to a fractional dose improves efficacy. This finding is just a consequence of uncertainty, independent of supply constraints that provide additional rationale for fractional dosing.

6. Calibrations of Epidemiological Model

This section moves from a static, reduced-form model to a detailed epidemiological model. The analysis of the simpler model had the advantage of producing analytical results that hold generally rather than just for calibrated parameters. The analysis in this section has a number of offsetting advantages.

The model incorporates dynamics in both the spread of the disease and the distribution of vaccine to the population. The dynamics are not merely a better approximation of reality; they bring forth new economic insights. In the static model, it was possible to install enough capacity Q to relieve the constraint on vaccine supplies. In a dynamic model, the constraint always binds. Whatever the installed capacity Q, were it doubled, the time to complete the vaccination campaign would be cut in half. The richer model also incorporates heterogeneity in mortality risk and contact rates across age cohorts. Beyond improving the model's realism, this addition opens up the possiblity of a new policy to investigate, moving from uniform to age-specific doses.

We calibrate the model with parameters drawn from the Covid-19 pandemic. The analysis thus far has suggested the theoretical possibility that fractional improves on full dosing. Here, we are interested in learning whether the theoretical possibility is relevant for realistic parameters and, if so, how small optimal fractional doses are and how much surplus is gained from that policy.

While most parameters are calibrated to the Covid-19 pandemic, we leave some free for comparative-statics analysis. The main one is the constraint on vaccine supply. We consider possibilities ranging from a small fraction to several multiples of the supply available during the Covid-19 pandemic. We seek to determine how optimal fractional dosing varies with the tightness of supply. We also consider both concave and sigmoid shapes for the efficacy function. To

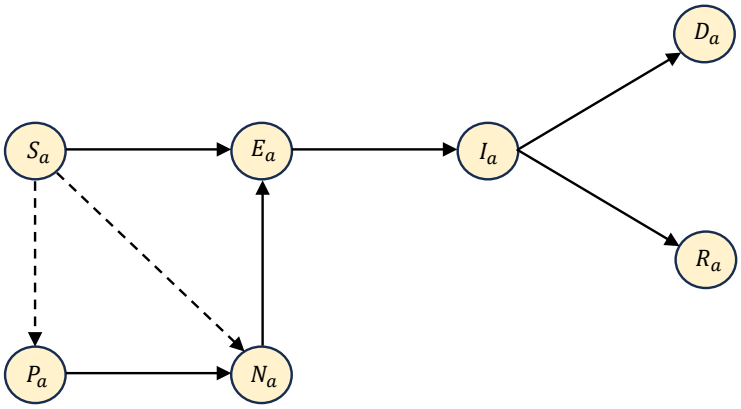


Figure 7: Schematic Diagram of SEIRDV Model. Arrows indicate flows between compartments. Dashed arrows represent vaccinations. The model allows for heterogeneity by age cohort, as indicated by the a subscripts.

determine which was relevant for Covid-19 vaccines would require data from the low end of the dose-response curve, which is mostly lacking for Covid-19 vaccines. In place of a known functional form, we explore the sensitivity of the results to various alternatives. In addition, given the considerable space allocated to separately analyzing each functional form, we are interested in seeing whether the functional forms lead to any practical difference in policy calibrations.

6.1. SEIRDV Model

To account for the dynamics of infection, mortality, and vaccination, we adopt a compartmental epidemiological model, the workhorse of mathematical epidemiology (Anderson and May, 1992). In particular, we adopt an extension of the classic and commonly used susceptible-exposed-infected-recovered model that includes vaccination and mortality dynamics, abbreviated SEIRDV. The model used here is a slight modification of the one described in detail in in own prior work (Więcek et al., 2022). Here, we give only a short exposition.

The population is stratified into age cohorts indexed by $a = 1, \ldots, A$. Each age cohort is further subdivided into the compartments shown in the schematic diagram in Figure 7. Susceptible individuals (S_a) who come into contact with someone who is infectious may be exposed to the disease (E_a) and themselves become infected (I_a). Some infected individuals die (D_a); the others recover (R_a), obtaining natural immunity. Susceptible individuals may be vaccinated. If the imperfectly effective vaccine happens to work, the individual is protected from infection (P_a); otherwise they are not protected (N_a).

More formally, members of each age cohort move between compartments according to the

following differential equations:

$$\dot{S}_a(t) = -\lambda_a(t)S_a(t) - q\,v_a(t)\left[\frac{S_a(t)}{S_a(t)+R_a(t)}\right] \tag{26}$$

$$\dot{E}_a(t) = \lambda_a(t)[S_a(t)+N_a(t)] - \varepsilon_a E_a(t) \tag{27}$$

$$\dot{I}_a(t) = \varepsilon_a E_a(t) - \xi_a I_a(t) \tag{28}$$

$$\dot{R}_a(t) = (1-\mu_a)\xi_a I_a(t) \tag{29}$$

$$\dot{D}_a(t) = \mu_a \xi_a I_a(t) \tag{30}$$

$$\dot{P}_a(t) = q\,v_a(t)\phi(x_a)\left[\frac{S_a(t)}{S_a(t)+R_a(t)}\right] - \psi_a P_a(t) \tag{31}$$

$$\dot{N}_a(t) = q\,v_a(t)[1-\phi(x_a)]\left[\frac{S_a(t)}{S_a(t)+R_a(t)}\right] + \psi_a P_a(t) - \lambda_a(t)N_a(t). \tag{32}$$

Susceptible individuals, in compartment $S_a(t)$, become infected at rate $\lambda_a(t)$, which depends on contacts with infectious individuals in various cohorts:

$$\lambda_a(t) = \beta \sum_{a'=1}^{A} c_{a,a'} I_{a'}(t). \tag{33}$$

Parameter β captures the virus' transmissibility; $c_{a,a'}$ is an element of the contact matrix, equal to the number of contacts made by an individual in cohort a with members of cohort a'.

Upon infection, individuals move from the susceptible to the exposed compartment E_a, where they remain while the virus incubates, moving to the infectious compartment I_a at a rate ε_a determined by the virus' incubation period. Individuals in the infectious compartment are able to transmit the virus to others. Individuals exit the infectious compartment at rate ξ_a determined by the virus' infection period. They then either die—entering compartment D_a with probability μ_a determined by the age-specific infection fatality rates—or recover—entering compartment R_a with probability $1-\mu_a$. Recovered individuals obtain natural immunity, which for most of the analysis we assume does not wane.[9]

Each unit of time, a constant proportion q of the total population is vaccinated, using all available vaccine capacity. The prioritization strategy is captured by $v_a(t)$, an indicator function that turns on when the cohort is eligible to receive the vaccine and remains on until the campaign is

[9]Assuming natural immunity does not wane is a reasonable approximation given the calibration traces a single epidemic outbreak over an interval of months, not years.

completed for that cohort. As in the reduced-form model, we assume that the vaccine is all-or-nothing, providing perfect protection (perfectly preventing transmission, disease, and death) for an individual with a probability given by the vaccine's efficacy, $\phi(x_a)$, moving them to compartment P_a. We abstract from the time required to develop immunity following vaccination. With probability $1-\phi(x_a)$, the vaccine does not work for the person, so they are not protected, moving to compartment, N_a. Individuals exit P_a when they lose vaccine-induced immunity, which occurs at rate ψ_a, moving them to compartment N_a. Individuals exit N_a when they are infected, which occurs at the same rate $\lambda_a(t)$ as for susceptible individuals.

Possible imperfections in the vaccine campaign are captured via three model assumptions. First, as mentioned, we assume that the vaccine is imperfectly effective, i.e., $\phi(x_a) < 1$. Second, we assume that recovered and susceptible individuals cannot be distinguished, so both compartments end up being vaccinated. Since recovered individuals already have natural immunity, the vaccine is effectively wasted on them. This imperfection explains the factor in square brackets in equation (26), the share left for the susceptible compartment after the "leakage" to the recovered compartment. Third, we assume that loss of vaccine-induced immunity is unobservable, so the move from P_a to N_a cannot be prevented by revaccinating the individual.

6.2. Calibration Parameters and Methods

We calibrate model parameters to the outcomes of initial clinical trials of the first effective Covid-19 vaccines and mass vaccination campaigns that began in late 2020 and continued through 2021, when all countries faced shortage of vaccines. Our goal is to guide the selection of roughly plausible parameters. Providing an accurate portrayal of the Covid-19 pandemic is beyond the scope of this exercise for several reasons. First, outcomes varied widely across countries; in taking the global average for a parameter, we may fail to reflect the experience in any given region. Second, the all-or-nothing nature of the vaccine in the model allows for a single notion of the efficacy of the vaccine, which works equally well to prevent infection, disease, severe disease, or death. In reality, Covid-19 vaccines were more effective at preventing deaths and hospitalizations than preventing transmission, and this distinction only grew with emergence of variants, which we do not model. Most importantly, our specification of the SEIRDV model has a single epidemic outbreak. This may not be unrealistic if interpreted as applying to the first wave caused by the ancestral strain of Covid-19 but becomes increasingly unrealistic with the emergence of Covid-19 variants, leading

to infection cycles, which the model is not designed to capture.[10]

Age Cohorts. We partition the population into eight age cohorts: a cohort of children under 19 and seven cohorts of adults, aged 20–29, 30–39, 40–49, 50–59, 60–69, 70–79, and 80 and above.

Virus Transmission. We calibrate the contact matrix $(c_{a,a'})$ to the frequency of interactions between different age groups from the large-scale surveys of European countries in Mossong et al. (2008). The basic reproductive number, $\mathscr{R}_0$, represents the average number of new infections generated by a single infected individual in a fully susceptible population. We calibrate the transmissibility parameter β so that the model delivers a value of $\mathscr{R}_0 = 2$. This round value of $\mathscr{R}_0$ is not meant to be a precise estimate[11] but illustrative of a relatively fast-moving epidemic like Covid-19. The appendix provides a calibration for a slower growing epidemic, with $\mathscr{R}_0 = 1.1$.

Fatality Rate. We calibrate the infection fatality rate μ_a using estimates from Verity et al. (2020), ranging from 0.03% for the youngest adult cohort (20–29 year-olds) up to 9.3% for the oldest cohort (80 and above).[12] These calibrated values of μ_a entail that the population mortality rate is 0.55% in the absence of vaccination in the model; roughly 1% of infected individuals die.

Vaccine Supply. Turn now to calibrating the vaccine side of the model, starting with vaccine supply. We measure vaccine supply in terms of the percentage of the whole population that can be vaccinated with full doses per day. We allow vaccine supply to range over several orders of magnitude (0.01%, 0.1%, 1%, and 10%), spanning possibilities from negligible to virtually unconstrained supplies. For example, a daily 0.1% supply achieves 50% coverage of the population with full doses in just over 16 months. We take 0.1% as the benchmark supply, since the associated time to achieve 50% coverage is close to the median 14 months across countries during the Covid-19 pandemic.[13]

[10] Więcek et al. (2023) analyzes a model allowing for infection cycling, finding that estimates of vaccine benefits are sensitive to fine model details concerning the partial immunity afforded by the primary vaccine series against future variants, the timing of boosters, and so forth.

[11] In reality, $\mathscr{R}_0$ did not have a single, well-defined value for Covid-19 but varied widely across regions and time in response to social distancing, lockdowns, and continual emergence of new variants of the virus.

[12] The Verity et al. (2020) estimates are based on large-scale surveys in Europe of mortality from ancestral Covid-19. While Covid-19 mortality rates varied across countries and over time, particularly as new variants emerged, the log-linear increase of mortality with age found in Verity et al. (2020) is consistent across studies.

[13] We calculated this figure using data from Our World in Data (2024), taking the start of the vaccination period as December 2, 2020, the first date vaccines were distributed in any country. Supplies varied dramatically across countries. Gibralter was the fastest country to 50% coverage with full doses, which it achieved in about three months; the United States and median European Union country took about eight months; 38% of countries never achieved that

Vaccine Prioritization. To calibrate the vaccine prioritization embodied in $v_a(t)$, note that, while the details varied from country to country, most followed the general principle that individuals at greatest risk of severe disease or death should be vaccinated first. Advanced age was typically a leading factor in the prioritization (World Health Organization, 2020b; Williams et al., 2021). Since our model only considers heterogeneity in age and not other characteristics like comorbidities or healthcare work, we specify a strict age-based prioritization, where each cohort is vaccinated in descending age order. The cohort whose turn it is to receive the vaccine is vaccinated with the entire available supply until all willing members are vaccinated. We assume 80% of each cohort are willing to be vaccinated and 20% not (perhaps due to logistical frictions or vaccine hesitancy).[14] We assume that the cohort of children cannot be vaccinated, echoing the fact that Covid-19 vaccines were not initially approved for children.

Dose-Response. Our calibrated dose-response curves are shown in Figure 8. We started with flexible sigmoid and concave functions, having three free parameters each. For each curve, we solve for the value of the three parameters that allow the curve to fit the three data points observed for Covid-19 vaccines, labeled *A*, *B*, and *C* in the figure. Details behind the specific functional forms used and parameters obtained are provided in the appendix.

Point *A* corresponds to the full doses of approved vaccines used during the pandemic. We normalize the dose strength at point *A* to to $x = 1$ and set the efficacy to $\phi(1) = 0.95$, approximating the efficacy of the mRNA vaccines against symptomatic infection in early clinical trials (Baden et al., 2020; Polak et al., 2020). Point *B* corresponds to a half-strength dose. Evidence summarized in Więcek et al. (2022) puts a reasonable estimate of the efficacy of a half dose at $\phi(0.5) = 0.8$. Point $C = (0,0)$ embodies the assumption that a dose with no active ingredient is useless.

The specifications we chose for the efficacy functions happen to pass through point *C* for all parameters, so an additional condition is required to pin down the three parameters. The additional condition is derived from the assumption that efficacy peaks at *A* and there is no marginal efficacy beyond that dosage. Formally, $x^p = 1$, implying $\phi'(1) = \phi'(x^p) = 0$. Assuming efficacy peaks at *A* preserves the internal consistency of the calibration. Given that we will abstract from the

coverage.

[14]Overall, 80% was the upper limit of vaccination coverage at the end of the pandemic (by 2023 United States reached 69% fully vaccinated, Germany 76%, United Kingdom 75%). In 2021 it was typical for countries with good supply to reach over 80% in highest risk groups and much lower in general population. This pattern differed considerably across countries (by 2023 Nigeria reached 30%, Ethiopia 35%, Democratic Republic of Congo below 10%; all figures according to Our World in Data (2024)).

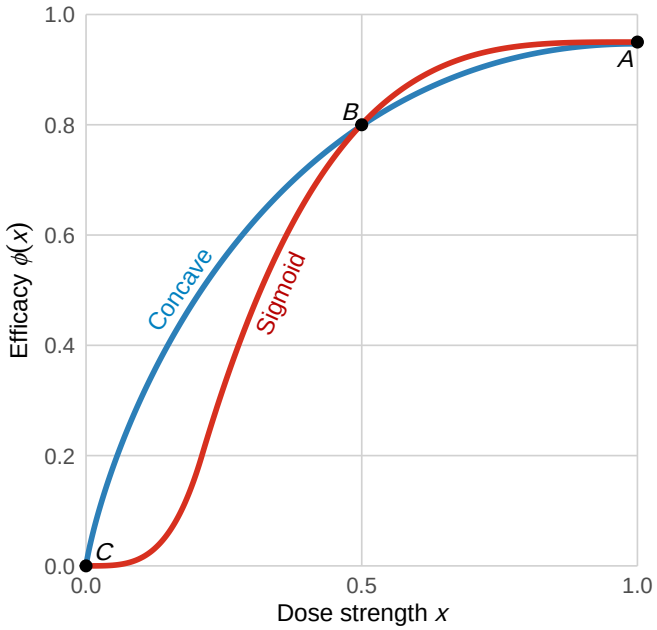


Figure 8: Calibrated Dose-Response Functions. The general functional forms for the curves are provided in equations (A100) and (A101) in the appendix. Curves are calibrated to pass through points A, B, and C and satisfy derivative condition $\phi'(1) = 0$, ensuring A is the peak of the efficacy function, thus representing a full dose.

generalized cost of vaccination, setting $g(x) = 0$, the only reason for limiting the strength of full doses is because efficacy peaks there.

Generalized Costs. The calibration abstracts from planner expenditures on vaccines and side effects, setting generalized costs, $g(x)$, to zero. We have practical and pedagogical grounds for this abstraction. On practical grounds, setting $g(x) = 0$ is a reasonable approximation, since serious side effects from Covid-19 vaccines were shown to be very rare (Faksova et al., 2024) and social benefits were estimated to be two orders of magnitude greater than the price countries paid for the vaccine (Castillo et al., 2021). On pedagogical grounds, setting $g(x) = 0$ absolves us from having to monetize mortality and morbidity losses to weigh against expenditures.

Calibration Methods. While some simple compartmental models admit closed-form solutions (see, e.g., Goodkin-Gold et al., 2024), complex models like the one here—with more compartments and heterogeneous cohorts—have no known analytical solutions. Therefore, we rely on numerical methods to simulate the time path of the model's compartments and evaluate the objective function behind the dose-optimization problem.[15] We run all simulations for 365 days.

[15] We use the *odin* package in R to solve the differential-equation system and the *nloptr* package in R (with tolerance set to 1e−6) for dose optimization. All of our code is available at https://github.com/DevInnovationLab/optimal_dose.

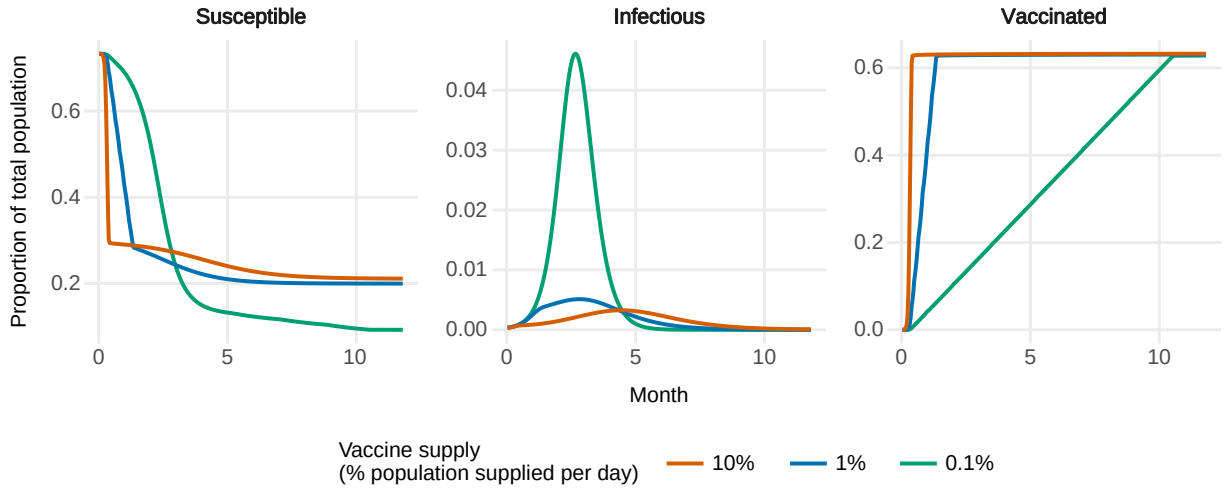


Figure 9: Evolution of Compartments in Calibrated Model. Simulates outcome from SEIRDV model for various levels of vaccine capacity. Simulation allows proportions in compartments to vary across age cohorts; graphs here aggregate age cohorts, showing proportions of total population.

Epidemic Diagram. Figure 9 shows the evolution of compartments in the model for various levels of capacity to vaccinate a flow of the population with full doses. At the calibrated value of the basic reproductive number, $\mathscr{R}_0 = 2$, the epidemic grows relatively fast. In the absence of vaccines, infections peak at about four months and about 55% of population are eventually infected.

6.3. Calibration Results

This subsection presents the results for the planner's optimal dosing policy in the calibrated model. Given that we have abstracted from generalized costs of vaccines, $g(x)$, the planner's objective reduces to choosing dose strength to maximize health benefits, or equivalently, to minimize epidemic harms. We will simplify further by abstracting from morbidity losses and consider only mortality losses. We also abstract from discounting and years of lost life, measuring mortality losses simply by the total number of deaths.

We solve for two versions of the optimal fractional-dosing policy, an age-invariant (uniform) fractional dose and age-varying fractional doses. The uniform fractional-dosing policy involves optimizing over a single variable $x \in [0,1]$. The age-varying fractional-dosing problem involves optimizing over a seven-dimensional vector of dose strengths (x_a).

The performance of these fractional-dosing policies will be gauged against a no-vaccine benchmark and a full-dose benchmark. Formally, let $AD = \sum_{a=1}^{A} D_a(T)$ denote aggregate deaths caused

by the epidemic by the end of the calibration summed across age cohorts. Adding an argument $AD(X)$ allows us to distinguish deaths under various vaccine policies X. Thus, $AD(0)$ denotes aggregate deaths under no vaccine, $AD(x^p)$ under full doses, $AD(x^*)$ under optimal uniform fractional dosing, and $AD(x_a^*)$ under optimal age-varying fractional dosing. The performance of policy X will be measured in two ways: the fraction of deaths averted relative to a no-vaccine benchmark,

$$\frac{AD(0) - AD(X)}{AD(0)}, \tag{34}$$

and relative to a full-dose benchmark,

$$\frac{AD(x^p) - AD(X)}{AD(x^p)}. \tag{35}$$

Table 1 presents the results for various levels of vaccine capacity and for both shapes of the dose-response function. Focus first on the results for the benchmark supply (0.1% of the population vaccinated each day) and sigmoid dose-response function. In that case, full doses avert 32% of the deaths under no vaccination. To stretch such a modest capacity, the optimal uniform dose ends up being about half of a full dose. Distributing these smaller doses more quickly allows the planner to avert 44% of deaths. With ten times more vaccine supply than the benchmark (1% of the population vaccinated each day), full doses avert most deaths (88%). Capacity of this magnitude speeds doses to the population so quickly that there is little gain from speeding the process further with fractional doses. Uniform fractional doses are 70% of the strength of full doses, averting 91% of deaths, only three percentage points more than full doses. When vaccine supply is increased another ten fold—to 10% of the population vaccinated each day—at that point supply is all but unconstrained, essentially leaving no room for fractional doses to generate improvements. Full doses already avert 95% of deaths, the maximum possible given the vaccine is calibrated to be 95% effective.

Moving in the other direction, considering capacity only a tenth of the benchmark (0.01% of the population vaccinated each day), full doses avert just 4% of deaths. The optimal uniform fractional dose averts 7% of deaths. Dose strength converges to the unit-elastic point ($\bar{x}^e = 0.4$) as predicted by the relevant proposition (Proposition 1) for the case of sigmoid efficacy and costless vaccination.

The results from the calibration are roughly similar with a concave rather than sigmoid dose-

Table 1: Calibration Results for Optimal Dosing and Averted Deaths

	Fraction of full dose X/x^p			Fraction of no-vaccine deaths averted $[D(0)-D(X)]/D(0)$		
Vaccine capacity (fraction of population supplied per day)	Full dose $X = x^p$	Uniform fractional dose $X = x^*$	Age-varying fractional dose $X = (x_a^*)$	Full dose $X = x^p$	Uniform fractional dose $X = x^*$	Age-varying fractional dose $X = (x_a^*)$
Sigmoid Dose Response						
0.01%	100%	42%	Fig. 11	4%	7%	7%
0.1%	100%	51%	Fig. 11	32%	44%	44%
1%	100%	70%	Fig. 11	88%	91%	91%
10%	100%	91%	Fig. 11	95%	95%	95%
Concave Dose Response						
0.01%	100%	7%	Fig. 11	4%	11%	11%
0.1%	100%	35%	Fig. 11	32%	45%	46%
1%	100%	73%	Fig. 11	88%	90%	90%
10%	100%	98%	Fig. 11	95%	95%	95%

Notes: Reports dose size under three different vaccination policies: full doses, optimal uniform fractional doses, and optimal age-varying fractional doses. Reported policy outcome is the percentage of deaths that would have occurred without a vaccine that are averted under the indicated vaccine policy given by equation (34). The optimal age-varying dosage is not a single number but the vector plotted in Figure 11.

response function. The main difference is that the optimal uniform fractional dose varies in a wider band. The dose is a larger share of the full dose for the largest capacity levels and a smaller share of the full dose for the smallest capacity levels. The dose converges to zero as vaccine supplies converge to zero, as predicted by Proposition 1 in the case of concave efficacy.

Figure 10 helps visualize the improvement of fractional over full doses, graphing the metric in equation (35). The percent of full-dose deaths averted is hump-shaped. When supplies are very low, fractional dosing cannot generate a big improvement because few deaths can be averted with any dosing policy. When supplies are very high, almost all deaths can be averted with full doses, so little room for improvement is left for fractional dosing. The percent improvement is greatest for moderate supplies. In particular, the curves peak at about 0.6% of the population supplied per day. At this supply level, fractional dosing increases averted deaths by 30–35% depending on the shape of the dose-response function.

The additional flexibility of allowing fractional doses to vary with age does not offer much benefit over uniform fractional doses. The dashed curves for age-varying dosing in Figure 10 are hardly above the solid curves for uniform dosing. The biggest departure is about 3 percentage

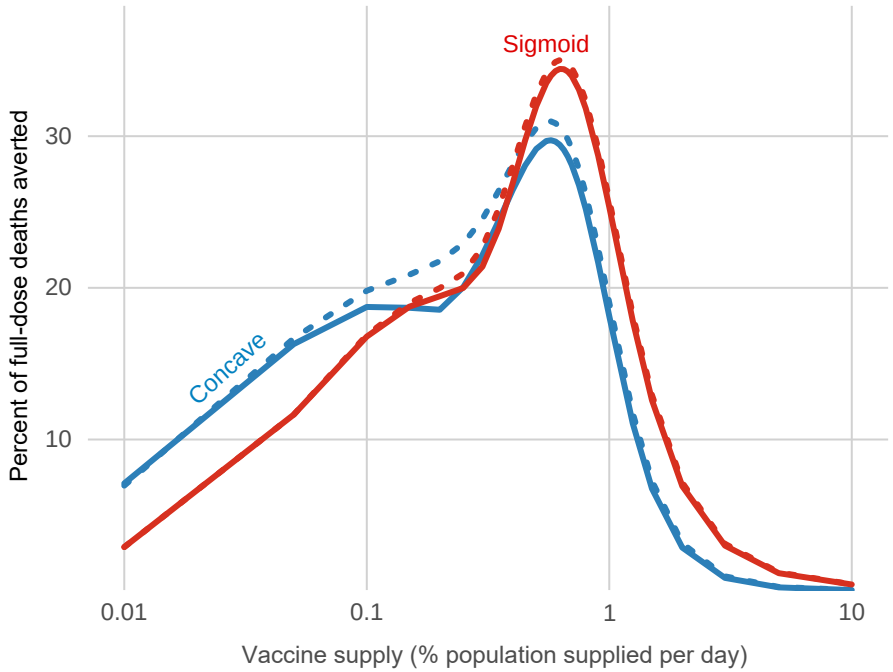


Figure 10: Deaths Averted Moving from Full to Fractional Dosing. Solid curves represent percentage of deaths under full dosing averted under optimal uniform fractional dosing. Dashed curves represent move from full to optimal age-varying fractional doses. Two sets of curves are provided for the two shapes of the dose-response function: sigmoid in red and concave in blue. The formula for the curves is provided in equation (35).

points, at supply of 0.2% per day. Age-varying doses hardly improve on uniform doses since vaccines are already prioritized to the most vulnerable age cohorts first. There is little to be gained from reducing the strength of doses going to an older cohort to slightly speed up the vaccination of the next, less vulnerable cohort.

Figure 11 plots the fractional dose received by each age cohort in the optimal age-varying policy. In cases in which all age cohorts receive doses, so the full age pattern can be observed, we see that dose strength is a U-shaped function of age. It relatively strong for the oldest cohorts because they are the most vulnerable. It is also relatively strong for the youngest cohorts because few cohorts, if any, remain to speed doses to. Middle-age cohorts receive the weakest doses, but only modestly less than the oldest or youngest cohorts. The variation in dose strength is particularly small under a sigmoid dose-response curve. For this curve, the "sweet spot" for fractionating doses without losing too much efficacy is small, and age-specific doses tend to stay inside this range.

Appendix A3 reprises all the results here for the case of a slow-growing epidemic, reducing the calibrated value of $\mathscr{R}_0$ from 2 to 1.1, holding other parameters constant. At the lower value of $\mathscr{R}_0$, the epidemic resembles a typical seasonal flu more than the Covid-19 pandemic. Infections

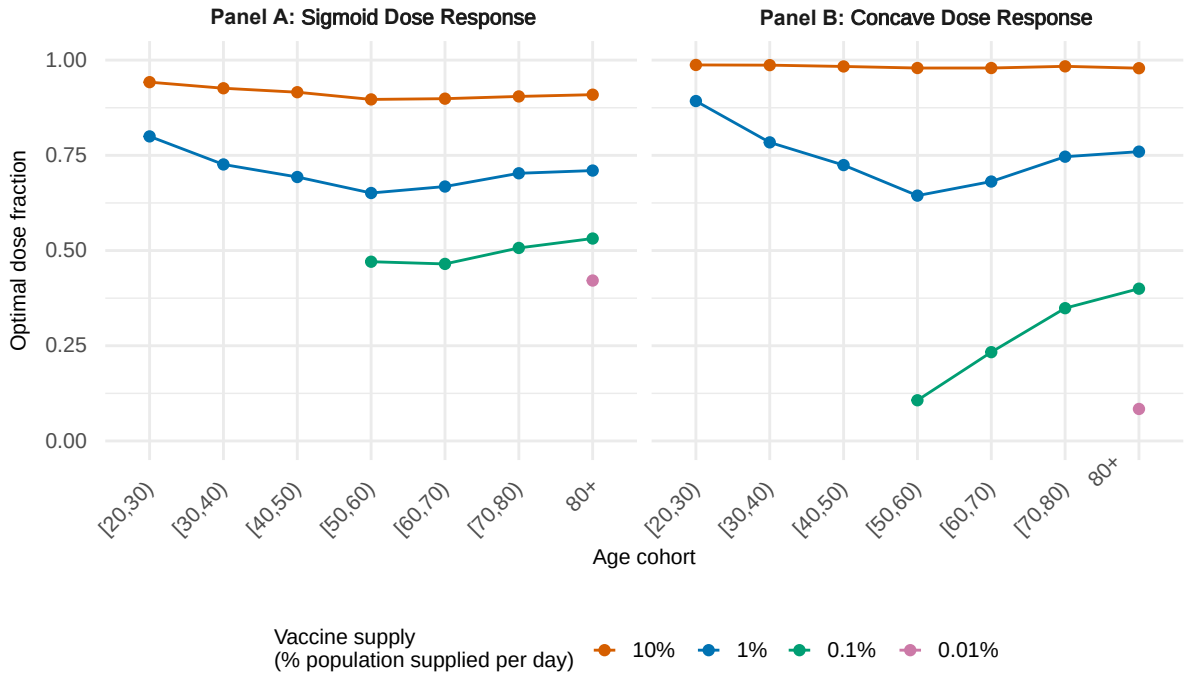


Figure 11: Optimal Age-Varying Fractional Doses. Each curve corresponds to a different order of magnitude of vaccine supply. Tick marks on horizontal axis correspond to different age cohorts. Variation in strength across cohorts is layered on top of vaccine prioritization by age, completing the vaccination campaign for an older cohort before moving to a younger cohort. We omit points when age cohort does not receive supply during the simulation period.

rise slowly over about six months in the absence of vaccination; fewer than 15% of population are eventually infected. The key insights are preserved from the case of a fast-growing epidemic: it remains true that optimal fractional doses can be considerably smaller than full doses and can result in substantial percentage-point improvements in averted deaths. Rather surprisingly, fractional dosing may be an even more powerful tool in a slow-growing epidemic. The reason is that herd immunity, an unreachable goal with a fast-growing epidemic, may be possible with vaccination in a slow-growing epidemic. The best strategy for reaching herd immunity—and heading off much of the mortality from the pandemic—may be to use fractional doses to speed vaccines to all age cohorts. Fractional dosing can improve on full dosing by as much as 40 percent in this calibration.

7. Conclusion

The first part of the paper studied a theoretical model designed to capture the essence of more complex epidemiological models in a tractable form, allowing analytical solutions to be derived. The analysis makes a clear case for fractional dosing under a concave dose-response function. For

that shape, if supplies are insufficient to deliver full doses to the entire population, fractional dosing is always optimal. Furthermore, the strength of the optimal dose is weakly increasing in vaccine supply.

The same results often hold under a sigmoid dose response. However, there are limited conditions under which a decrease in supply leads to stronger optimal doses. This counterintuitive result emerges when side effects are convex and fixed administration costs are sufficiently high. Under these conditions, as vaccine supplies become extremely scarce, the relatively unprotected population faces a rising infection rate. It becomes more important to provide effective protection with a concentrated dose than to moderate the side effects caused by concentrated doses.

The paper then moved to calibrating a richer epidemiological model to parameters from the Covid-19 pandemic. The goal was to gauge the relevance of the theory for realistic supply constraints (and a range well above and below), not to prescribe policy for Covid-19 or future pandemics. To prescribe policy, some limiting assumptions would have to be relaxed, including assuming vaccines are all-or-nothing (Covid-19 vaccines were in fact more effective in preventing serious illness and death than infection and transmission), assuming efficacy is constant (the emergence of different Covid-19 strains and viral drift lead the effectiveness of vaccines to wane over time), and assuming that the onset and duration of immunity are independent of dose strength.

In calibrations with the concave dose-response function, the optimal fractional dose spanned almost the entire unit interval for the capacity levels we investigated. The optimal dose varied in a narrower range for the sigmoid function, corresponding to its narrower "sweet spot" for maintaining efficacy, which drops precipitously at low dosages for this shape. We found that the gains from fractional dosing can be substantial. With enough supply to vaccinate 0.6% of the population daily (achieving 50% coverage with full doses in about three months, the fastest any country was able to achieve that coverage during the Covid-19 pandemic), mortality under full doses can be reduced 30% by moving to optimal fractional doses when the dose-response function is concave. The gains can be even larger for a sigmoid dose-response function and for the slow-growing epidemic studied in the appendix.

We also calibrated the performance of a program allowing the optimal dose to vary by age group. We found that this produced minimal gains over the optimal uniform fractional dose (same across age groups). The fact that the calibration prioritizes older cohorts to receive vaccines earlier precludes much benefit from varying dosage by age. The benefit would have to come from accel-

erating doses distributed later in the vaccination campaign by shrinking doses distributed earlier. But shrinking doses given to the most vulnerable cohorts vaccinated early in the campaign can be counterproductive.

Our calibration may considerably understate the potential for gains from fractional dosing in many real-world cases. The introduction mentioned that fifth doses were used in the 2018 yellow fever and 2022 mpox outbreaks. In both cases randomised trials suggested immunological non-inferiority (Juan-Giner et al., 2021; Centers for Disease Control and Prevention, 2026). In the context of our model this suggests that full doses for these vaccines must be very far out on the long plateau of dose-response functions. This is in contrast to our calibration example which assumed a large drop-off in efficacy at low doses.

That the full dose was five times more concentrated than needed for the same efficacy is remarkable. Such an outcome may emerge from way clinical trials are staged in vaccine development. Early-stage trials are often run escalating the antigen concentration while an acceptable safety and reactogenicity profile is retained. Facing substantial uncertainty about the dose-response curve before committing to costly, late-stage efficacy trials, a profit-maximizing developer may choose to run the late-stage trials on the strongest dose, offering the best chance of generating a detectable immune response. For some vaccines, this process can lead the full dose as initially approved to be much larger than needed to generate a similar efficacy.

Our extension to uncertain efficacy in Section 5.2 helps formalize these ideas. The extension (which posited negligible side effects and a peaked efficacy function) explored the case in which losses from missing an unknown peak on the low side are roughly symmetric to the losses from missing on the high side. If, alternatively, the efficacy function has a long plateau rather than a sharp peak, losses from missing the optimal full dose on the low side could be asymmetrically severe, biasing the developer toward full doses that are stronger than necessary, thus increasing the gains from fractional dosing when an emergency leads to shortages of those vaccines.

Appendix A1. Proofs

This appendix provides proofs of the propositions and lemmas stated in the text.

Proof of Lemma 1

After establishing some preliminary results, we prove the lemma for the subcase of concave ϕ separately from the subcase of strictly sigmoid ϕ.

Right limit. The text defined $\alpha_\phi(0) \equiv \phi'(0)$ and $\eta_\phi(0) \equiv 1$. Since the present proof relies on those definitions, we need to verify that they are consistent with the right limits of those functions at 0, as the text claims. Since $\phi(0) = 0$, $\alpha_\phi(0) = \phi(0)/0$ has the $0/0$ form. Applying l'Hôpital's Rule,

$$\lim_{x \downarrow 0} \alpha_\phi(x) = \lim_{x \downarrow 0} \frac{\phi(x)}{x} = \lim_{x \downarrow 0} \frac{\phi'(x)}{1} = \phi'(0). \tag{A1}$$

Since $\eta_\phi(x) = \phi'(x)/\alpha_\phi(x)$,

$$\lim_{x \downarrow 0} \eta_\phi(x) = \lim_{x \downarrow 0} \frac{\phi'(x)}{\alpha_\phi(x)} = \frac{\phi'(0)}{\phi'(0)} = 1. \tag{A2}$$

Left limit. We use a sandwiching argument to show

$$\lim_{x \uparrow \infty} \alpha_\phi(x) = 0. \tag{A3}$$

Now

$$0 < \alpha_\phi(x) = \frac{\phi(x)}{x} < \frac{1}{x}, \tag{A4}$$

where the last step follows from $\phi(x) < 1$. Taking limits,

$$0 \leq \lim_{x \uparrow \infty} \alpha_\phi(x) \leq \lim_{x \uparrow \infty} \frac{1}{x} = 0, \tag{A5}$$

which proves (A3).

Useful alternative expressions. Next, we establish some useful alternative expressions for average efficacy $\alpha_\phi(x)$ and its derivative $\alpha'_\phi(x)$. By the Fundamental Theorem of Calculus,

$$\alpha_\phi(x) = \frac{\phi(x)}{x} = \frac{1}{x} \int_0^x \phi'(\omega) d\omega. \tag{A6}$$

Differentiating,

$$\alpha'_\phi(x) = \frac{z(x)}{x^2}, \tag{A7}$$

where

$$z(x) \equiv x\phi'(x) - \int_0^x \phi'(\omega) d\omega \tag{A8}$$

$$= \int_0^x \omega \phi''(\omega) d\omega. \tag{A9}$$

To see (A8) equals (A9), integrate (A9) by parts:

$$\int_0^x \omega\phi''(\omega)d\omega = \int_0^x v(\omega)du(\omega) \tag{A10}$$

$$= u(\omega)v(\omega\Big|_{\omega=0}^{\omega=x} - \int_0^x u(\omega)dv(\omega) \tag{A11}$$

$$= x\phi'(x) - \int_0^x \phi'(\omega)d\omega, \tag{A12}$$

where we have defined $du(\omega) \equiv \phi''(\omega)d\omega$ and $v(\omega) \equiv \omega$, implying $u(\omega) = \phi'(\omega)$ and $dv(\omega) = d\omega$.

Concave ϕ. We next prove the lemmas in the subcase of concave ϕ. In this subcase,

$$\alpha_\phi(x) = \frac{1}{x}\int_0^x \phi'(\omega)d\omega \tag{A13}$$

$$> \frac{1}{x}\int_0^x \phi'(x)d\omega \tag{A14}$$

$$= \phi'(x), \tag{A15}$$

where (A14) follows from (A6) and (A14) follows from $\phi''(x) < 0$, implying $\phi'(\omega) > \phi'(x)$ for $\omega \in (0,x)$. Thus, $\eta_\phi(x) = \phi'(x)/\alpha_\phi(x) < 1$ for all $x > 0$, implying $\{x \geq 0 \mid \eta_\phi(x) = 1\} = \{0\}$, implying $\bar{x}^e = 0$ by (6). Since $\bar{x}^e = 0$, $\eta_\phi(x^e) = \eta_\phi(0) = 1$ by definition. The set $(0,\bar{x}^e)$ is empty. We have shown $\eta_\phi(x) < 1$ for all $x > 0 = \bar{x}^e$. This completes the proof of Lemma 1 when ϕ is concave.

Strictly sigmoid ϕ. Next, we prove the lemmas in the subcase of strictly sigmoid ϕ. By definition of this shape, there exists $x^i > 0$ such that $\phi''(x) > 0$ for all $x \in (0,x^i)$ and $\phi''(x) < 0$ for all $x > x^i$.

To show that $\alpha_\phi(x)$ is strictly quasiconcave, we proceed in several steps. For all $x \in (0,x^i)$, $\phi''(t) > 0$ for all $t \in (0,x)$, implying $z(x) > 0$ by (A9), implying $\alpha'_\phi(x) > 0$ by (A7). Consider $x' > x^i$ such that $\alpha'_\phi(x') < 0$. Such x' must exist since $\alpha_\phi(x)$ is positive for all $x > 0$ but approaches 0 in the limit for large x by (A3), so $\alpha_\phi(x)$ must eventually decrease for large x. For all $x'' > x'$,

$$z(x'') = z(x') + \int_{x'}^{x''} \omega\phi''(\omega)d\omega \tag{A16}$$

$$< z(x') \tag{A17}$$

$$< 0, \tag{A18}$$

where (A16) follows from (A9); (A17) follows from $\phi''(\omega) < 0$ for all $\omega > x' > x^i$; and (A18) follows from $0 > \alpha'_\phi(x') = z(x')/(x')^2$ by (A7), implying $z(x') < 0$. Hence, $\alpha_\phi(x'') = z(x'')/(x'')^2 < 0$. This proves that $\alpha_\phi(x)$ is strictly quasiconcave because we have shown that the function is initially upward sloping, is eventually downward sloping, and once it slopes downward, it continues to slope downward for all larger x.

Since ϕ is real-analytic, α_ϕ is real-analytic for all $x > 0$, being the ratio of nonzero real-analytic functions. Hence, α_ϕ is smooth, meaning it is continuously differentiable of all orders, implying α'_ϕ is continuous. Since $\alpha'_\phi(x) > 0$ for all $x < x^i$ and $\alpha'_\phi(x) < 0$ for all $x > x'$, by continuity $\alpha'_\phi(\tilde{x}) = 0$

for some $\tilde{x} \in (x^i, x')$. This first-order condition is necessary and sufficient for $\tilde{x}$ to be the unique maximizer of the quasiconcave function α_ϕ. Differentiating,

$$\alpha'_\phi(x) = \frac{\phi'(x)x - \phi(x)}{x^2} = \frac{\alpha_\phi(x)}{x}[\eta_\phi(x) - 1]. \tag{A19}$$

Therefore, the first-order condition $\alpha'_\phi(\tilde{x}) = 0$ is equivalent to $\eta_\phi(\tilde{x}) = 1$.

Since $\alpha_\phi(x)$ is strictly quasiconcave for all $x > 0$ and $\alpha'_\phi(\tilde{x}) = 0$, $\alpha'_\phi(x) > 0$ for all $x \in (0, \tilde{x})$, implying $\eta_\phi(x) > 1$ for all $x \in (0, \tilde{x})$ by (A19). Similarly, $\alpha'_\phi(x) < 0$ for all $x > \tilde{x}$, implying $\eta_\phi(x) < 1$ for all $x > \tilde{x}$. Therefore, $\eta_\phi(x) = 1$ for just two values of $x \geq 0$: those in the set $\{0, \tilde{x}\}$. Therefore,

$$\bar{x}^e \equiv \max\left\{x \geq 0 \mid \eta_\phi(x) = 1\right\} \tag{A20}$$

$$= \max\{0, \tilde{x}\} \tag{A21}$$

$$= \tilde{x}. \tag{A22}$$

To summarize, for strictly sigmoid ϕ, we identified $\tilde{x} > x^i$ satisfying $\eta_\phi(\tilde{x}) = 1$ and showed $\eta_\phi(x) > 1$ for all $x \in (0, \tilde{x})$ and $\eta_\phi(x) < 1$ for all $x > \tilde{x}$. Since $\bar{x}^e = \tilde{x}$, these facts hold for $\bar{x}^e$ as well, verifying Lemma 1. *Q.E.D.*

Proof of Lemma 2

We prove the lemma first for concave ϕ and then for strictly sigmoid ϕ.

Concave ϕ. Substituting (A9) into (A7) yields

$$\alpha'_\phi(x) = \frac{1}{x^2}\int_0^x \omega\phi''(\omega)d\omega, \tag{A23}$$

implying $\alpha'_\phi(x) < 0$ for all $x > 0$ because $\phi''(x) < 0$ for all $x > 0$. It is immediate that $\alpha_\phi(x)$ is strictly quasiconcave since strictly decreasing functions are strictly quasiconcave. Since $(0, \bar{x}^e)$ is empty, it is vacuously true that $\alpha_\phi(x)$ is strictly increasing on $(0, \bar{x}^e)$. Since $\alpha_\phi(x)$ is strictly decreasing for $x > 0$ and the definition $\alpha_\phi(0) = \phi'(0)$ is its right limit, we have that $\alpha_\phi(x)$ is maximized over $x \geq 0$ at $0 = \bar{x}^e$.

Strictly sigmoid ϕ. As an intermediate step in the analysis, the proof of the previous lemma showed that $\alpha_\phi(x)$ is strictly quasiconcave for $x > 0$, strictly increasing for $x \in (0, \bar{x}^e)$, strictly decreasing for $x > \bar{x}^e$, reaching a maximum for $\bar{x}^e$.

Connection to Karlin and Rubin (1956). Equations (A7)–(A9) imply

$$\alpha'_\phi(x) = \int_0^x \frac{\omega}{x^2}\phi''(\omega)d\omega \tag{A24}$$

$$= \int_0^x p(x \mid \omega)h(\omega)dF(\omega), \tag{A25}$$

where (A25) follows from translating into the notation in Karlin and Rubin (1956) by setting $p(x \mid \omega) \equiv \omega/x^2$, $h(\omega) \equiv \phi''(\omega)$ and $dF(\omega) \equiv d\omega$. It is straightforward to show $p(x \mid \omega)$ possesses monotone likelihood ratio. Furthermore, $h(\omega)$ changes sign once, since ϕ is sigmoid, implying $\phi''(x)$ changes sign from positive to negative once. Hence, the authors' Lemma 1 applies, providing that the integral in (A25) changes sign at most once. Therefore, $\alpha'_\phi(x)$ changes sign from positive to negative at most once, proving that $\alpha_\phi(x)$ is quasiconcave. *Q.E.D.*

Proof of Proposition 1

Nontrivial Solution. Before setting up the KKT conditions, the problem can be simplified by showing that optimum must be nontrivial: $n^*, x^* > 0$. If either $n = 0$ or $x = 0$, the planner's objective (7) equals $B(0) = 0$. A feasible alternative that does not violate the constraints sets $n = N$ and $x = Q/N$, generating $B\big(N\phi(Q/N)\big) > 0$ for the planner, a strict improvement over the trivial solution, proving $n^*, x^* > 0$.

KKT Conditions. Letting λ and μ be the Lagrange multiplier on constraints (3) and (4), respectively, the Lagrangian associated with the planner's maximization problem assuming costless vaccination is

$$\mathcal{L} = B\big(n\phi(x)\big) + \lambda(Q - nx) + \mu(N - n). \tag{A26}$$

Given the optimum is nontrivial as proved above, nonnegativity constraints (5) do not bind and thus can be omitted from the problem. The remaining KKT conditions are

$$\frac{\partial \mathcal{L}}{\partial n} = B'\big(n^*\phi(x^*)\big)\phi(x^*) - \lambda^* x^* - \mu^* = 0 \tag{A27}$$

$$\frac{\partial \mathcal{L}}{\partial x} = B'\big(n^*\phi(x^*)\big)n^*\phi'(x^*) - \lambda^* n^* = 0 \tag{A28}$$

$$Q \geq n^* x^* \tag{A29}$$

$$N \geq n^* \tag{A30}$$

$$\lambda^* \geq 0 \tag{A31}$$

$$\mu^* \geq 0 \tag{A32}$$

$$\lambda^*(Q - n^* x^*) = 0 \tag{A33}$$

$$\mu^*(N - n^*) = 0. \tag{A34}$$

Necessity of KKT Conditions. We will prove that a nontrivial solution to the planner's problem satisfies linear independence constraint qualification (LICQ). According to textbook results (see, e.g., Solodov, 2011), if the problem's constraints satisfy LICQ, the KKT conditions are necessary for an optimum.

LICQ requires the gradients of all equality constraints and all active inequality constraints to be linearly independent. Note that the planner's problem involves only inequality constraints, not equality constraints. The nonnegativity constraint (5) is inactive when $x^* > 0$ and $n^* > 0$. The remaining constraints (3) and (4) are the only potentially active ones in a nontrivial solution.

The gradients of constraints (3) and (4) are linearly independent if and only if the Jacobian matrix containing these gradients,

$$J = \begin{pmatrix} n & x \\ 0 & 1 \end{pmatrix}, \tag{A35}$$

has full rank. Its determinant is $|J| = n$, which is nonzero in a nontrivial optimum with $n^* > 0$, proving J has full rank in that case.

Derivation of Optimum. Suppose for the sake of contradiction that $\lambda^* = 0$. Then (A28) implies either $B'\big(n^*\phi(x^*)\big) = 0$ or $\phi'(x^*) = 0$, contradicting $B'(v) > 0$ for all $v > 0$ and $\phi'(x) > 0$ for all $x > 0$. Thus $\lambda^* > 0$, implying $n^*x^* = Q$ by (A33).

Condition A32 is consistent with two subcases we will proceed to analyze separately: $\mu^* = 0$ and $\mu^* > 0$. First suppose $\mu^* = 0$. Substituting $\mu^* = 0$ into (A27) and rearranging yields

$$\frac{\phi(x^*)}{x^*} = \frac{\lambda^*}{B'\big(n^*\phi(x^*)\big)}. \tag{A36}$$

Rearranging (A28) yields

$$\phi'(x^*) = \frac{\lambda^*}{B'\big(n^*\phi(x^*)\big)}. \tag{A37}$$

Equations (A36) and (A37) together imply $\phi'(x^*) = \phi(x^*)/x^*$, implying $\eta_\phi(x^*) = 1$. A nontrivial solution requires $x^* > 0$, implying $x^* = \bar{x}^e$. Further, $n^* = Q/\bar{x}^e$ since (A33) holds with equality.

Second, suppose $\mu^* > 0$. Then $n^* = N$ by (A34). Further, $x^* = Q/N$ since (A33) holds with equality.

One can see that the solution for $\mu^* = 0$ corresponds to the solution of the relaxed problem omitting the population constraint (4) and thus must be the solution to the original problem as long as it does not violate the omitted constraint, i.e., as long as $n^* = Q/\bar{x}^e \leq N$, or equivalently $\bar{x}^e \geq Q/N$. Otherwise the solution for $\mu^* > 0$ obtains, with $x^* = Q/N$. Combining these solutions yields $x^* = \max\{\bar{x}^e, Q/N\}$. Since $n^* = Q/x^*$, we have $n^* = \min\{Q/\bar{x}^e, N\}$, as stated in the proposition. *Q.E.D.*

Proof of Lemma 3

We need to prove both directions of the biconditional statement.

($\Longrightarrow$) Suppose the planner's solution is nontrivial. We will find $x > 0$ satisfying

$$B'(0)\phi(x) > g(x). \tag{A38}$$

Since the planner's objective (2) equals zero when evaluated at either $n = 0$ or $x = 0$, for the planner's solution (n^*, x^*) to be nontrivial, it must generate a positive value for (2). Thus,

$$0 < B\big(n^*\phi(x^*)\big) - n^*g(x^*) \tag{A39}$$

$$= \int_0^{n^*} \big[B'\big(n\phi(x^*)\big)\phi(x^*) - g(x^*)\big]\, dn \tag{A40}$$

$$\leq \int_0^{n^*} \big[B'(0)\phi(x^*) - g(x^*)\big]\, dn. \tag{A41}$$

Equation (A40) follows from the Fundamental Theorem of Calculus. Inequality (A41) follows from $B''(v) < 0$, implying $B'(v)$ is maximized over $v \geq 0$ by $v = 0$. Since the integral in (A41) is positive, its integrand must be positive. Thus, (A38) is satisfied by the nontrivial optimal dosage $x^* > 0$.

($\Longleftarrow$) Suppose (A38) is satisfied for some positive dosage, call it $\tilde{x} > 0$. We will construct a vaccination policy generating a positive value for the planner's objective (2). Since the optimal policy must do weakly better than our constructed policy, it must also generate positive planner surplus and thus cannot be trivial.

Consider a policy vaccinating a small number of people $\tilde{\varepsilon} > 0$ with a dose of potency $\tilde{x}$. Substituting these values into (2) yields planner surplus

$$B\big(\tilde{\varepsilon}\phi(\tilde{x})\big) - \tilde{\varepsilon}g(\tilde{x}) = \tilde{\varepsilon}g(\tilde{x})\left[\frac{B\big(\tilde{\varepsilon}\phi(\tilde{x})\big)}{\tilde{\varepsilon}g(\tilde{x})} - 1\right]. \tag{A42}$$

By l'Hôpital's Rule, the limit of the bracketed factor for small $\tilde{\varepsilon}$ is

$$\lim_{\tilde{\varepsilon}\downarrow 0}\left[\frac{B\big(\tilde{\varepsilon}\phi(\tilde{x})\big)}{\tilde{\varepsilon}g(\tilde{x})} - 1\right] = B'(0)\frac{\phi(\tilde{x})}{g(\tilde{x})} - 1, \tag{A43}$$

which is positive since $\tilde{x}$ satisfies (A38). Since the limit of the bracketed factor in (A42) is positive, by continuity it is positive for $\tilde{\varepsilon}$ in a neighborhood above 0, implying the entire expression in (A42) is positive for $\tilde{\varepsilon}$ in that neighborhood. *Q.E.D.*

Proof of Lemma 4

First statement of lemma. We need to prove that X^0 is a nonempty, finite set, which will allow us to conclude it has a unique largest element $\bar{x}^0$. We will further argue that $\bar{x}^0 > 0$.

We begin by using a sandwiching argument to establish

$$\lim_{x\uparrow\infty}\left[\frac{\phi(x)}{g(x)}\right] = 0. \tag{A44}$$

For all $x > 0$,

$$0 < \frac{\phi(x)}{g(x)} < \frac{1}{g(x)} < \frac{1}{px}. \tag{A45}$$

The first inequality follows from $\phi(x), g(x) > 0$ and the second inequality from $\phi(x) < 1$. The last inequality holds since $s(x) \geq 0$ and $k > 0$, which imply $g(x) \equiv s(x) + px + k > px$. Taking limits in (A45),

$$0 \leq \lim_{x\uparrow\infty}\left[\frac{\phi(x)}{g(x)}\right] \leq \lim_{x\uparrow\infty}\left(\frac{1}{px}\right) = 0, \tag{A46}$$

where the final equality follows from $p > 0$. This establishes (A44).

Pick any $x' \in (0,\infty)$, and set $L' \equiv \phi(x')/g(x')$. By (A44), there exists $x'' \in (0,\infty)$ such that $\phi(x)/g(x) < L'$ for all $x > x''$.

One can invoke the Extreme Value Theorem to argue that the function $\phi(x)/g(x)$ has a maximizer on $[0,x'']$. To see this, note that the preceding interval is compact. Note further that $\phi(x)/g(x)$ is real-analytic on $[0,\infty)$ since it is a ratio of real-analytic functions with a nonzero denominator, as $g(x) \geq k > 0$ for all $x \geq 0$. Since real-analytic functions are continuous, $\phi(x)/g(x)$ is continuous on $[0,\infty)$ and thus continuous on $[0,x'']$. A continuous function on a compact interval has a maximum by the Extreme Value Theorem.

As can be proved, a real-analytic function on a compact interval has a finite set of maximizers as long as it is not a constant function on that interval. Of course, $\phi(x)/g(x)$ is not constant over $[0,x'']$ since $\phi(0)/g(0) = 0/k = 0$ but $\phi(x)/g(x) > 0$ for all $x > 0$, so $\phi(x)/g(x)$ must initially increase. Thus, $\phi(x)/g(x)$ has a finite set of maximizers on $[0,x'']$, and this finite set has a unique largest element, $\bar{x}^0$. We have

$$x^0 = \max\left\{\underset{x\in[0,x'']}{\operatorname{argmax}}\left[\frac{\phi(x)}{g(x)}\right]\right\} \tag{A47}$$

$$= \max\left\{\underset{x\geq 0}{\operatorname{argmax}}\left[\frac{\phi(x)}{g(x)}\right]\right\}. \tag{A48}$$

The last step holds since $\phi(x')/g(x') = L' > \phi(x)/g(x)$ for all $x > x''$, implying

$$\underset{x\in[0,x'']}{\operatorname{argmax}}\left[\frac{\phi(x)}{g(x)}\right] = \underset{x\geq 0}{\operatorname{argmax}}\left[\frac{\phi(x)}{g(x)}\right]. \tag{A49}$$

We showed in the previous paragraph that $\phi(0)/g(0) = 0 < \phi(x)/g(x)$ for all $x > 0$. Hence, all elements of X^0 are positive, implying $\bar{x}^0 > 0$.

Second statement of lemma. To prove the second statement of the lemma, we will show that $\bar{x}^0 \geq \bar{x}^e$ implies $\underline{x}^0 = \bar{x}^0$ and thus $X^0 = \{\bar{x}^0\}$.

Suppose there exists $x^0 \in X^0$ such that $x^0 < \bar{x}^e$. Then $1 < \eta_\phi(x^0) = \eta_g(x^0)$, where the inequality follows from Lemma 1 and the equality from (11). We will show $\eta_g(x) > 1$ for all $x \geq x^0$. Differentiating $\eta_g(x)$ yields

$$\eta_g'(x) = \frac{xg''(x)}{g(x)} + \frac{1}{x}\eta_g(x)[1-\eta_g(x)]. \tag{A50}$$

For any x such that $\eta_g(x) = 1$, (A50) implies $\eta_g'(x) = xg''(x)/g(x) > 0$. Hence, $\eta_g(x)$ is upward sloping anywhere it crosses a line of height 1, implying it can only cross that line once. Therefore, $\eta_g(x^0) > 1$ implies $\eta_g(x) > 1$ for all $x \geq x^0$. But $\eta_\phi(x) \leq 1$ for all $x \geq \bar{x}^e$, implying $\eta_\phi(x) \neq \eta_g(x)$ for all $x \geq \bar{x}^e$, implying no $x \geq \bar{x}^e$ can maximize (8) by (11), implying $\bar{x}^0 < \bar{x}^e$.

Suppose for the remainder of the proof that $\bar{x}^0 \geq \bar{x}^e$. As just shown, we must have, $x^0 \geq x^e$ for all $x^0 \in X^0$, implying $\underline{x}^0 \geq x^e$. Thus $\underline{x}^0 > x^i$ since $\bar{x}^e > x^i$ by Lemma 1. If $\phi(x)$ is strictly sigmoid, then $\phi''(x) < 0$ for all $x > x^i$, implying $\phi''(x) < 0$ for all $x \geq \underline{x}^0$. If $\phi(x)$ is concave, then $\phi''(x) < 0$ for all $x > 0$. Thus, for each of the two possible sigmoid shapes, $\phi''(x) < 0$ for all $x \geq \underline{x}^0$.

Differentiating $\eta_\phi(x)$ and rearranging, we can write

$$\eta_\phi'(x) = \frac{x\phi''(x)}{\phi(x)} + H\big(x, \eta_\phi(x)\big) \tag{A51}$$

where

$$H(x,\eta) \equiv \frac{1}{x}\eta(1-\eta). \tag{A52}$$

Given $\phi''(x) < 0$ for all $x \geq \underline{x}^0$ as shown above,

$$\eta_\phi'(x) < H\big(x, \eta_\phi(x)\big). \tag{A53}$$

Let $\eta(x)$ be a solution to the differential equation $\eta'(x) = H\big(x, \eta(x)\big)$. Inequality (A53) establishes $\eta_\phi(x)$ as an *under-function* relative to $\eta(x)$ on $[\underline{x}^0, \infty)$ according to Definition 1.2.1 from Lakshmikantham and Leela (1969).

Similarly, we can write

$$\eta_g'(x) = \frac{xg''(x)}{g(x)} + H\big(x, \eta_g(x)\big). \tag{A54}$$

Given $g''(x) > 0$ for all $x > 0$,

$$\eta_g'(x) > H\big(x, \eta_g(x)\big), \tag{A55}$$

establishing $\eta_g(x)$ as an *over-function* relative to $\eta(x)$ on $[\underline{x}^0, \infty)$ according Definition 1.2.1 from Lakshmikantham and Leela (1969). By Theorem 1.3.2 of Lakshmikantham and Leela (1969), since $\eta_\phi(x)$ is an under-function and $\eta_g(x)$ an over-function relative to the same function on $[\underline{x}^0, \infty)$ and $\eta_\phi(\underline{x}^0) = \eta_g(\underline{x}^0)$, $\eta_\phi(x) < \eta_g(x)$ for all $x > \underline{x}^0$. Thus, no $x > \underline{x}^0$ can satisfy (11), a necessary condition for a solution x^0, implying $\bar{x}^0 \not\geq \underline{x}^0$, implying $\underline{x}^0 = \bar{x}^0$. *Q.E.D.*

Proof of Lemma 5

We will prove the conditions for a nontrivial planner optimum in this and Lemma 3 are equivalent. Suppose there exists $x > 0$ such that $B'(0)\phi(x) > g(x)$, the necessary and sufficient condition for a nontrivial planner optimum in Lemma 3. Rearranging,

$$1 < \frac{B'(0)\phi(x)}{g(x)} \leq \frac{B'(0)\phi(\bar{x}^0)}{g(\bar{x}^0)} = \rho^0(0). \tag{A56}$$

The second step follows from the fact that x^0 maximizes $\phi(x)/g(x)$ by (10). The last step follows from the definition of $\rho^0(n)$ in (12). Equation (A56) implies $\rho^0(0) > 1$, the necessary and sufficient condition for a nontrivial planner optimum stated in this lemma.

Suppose $\rho^0(0) > 1$. Then $B'(0)\phi(x^0) > g(x^0)$ by (12), demonstrating the existence of some $x > 0$, namely x^0, satisfying $B'(0)\phi(x) > g(x)$. Thus, the condition for a nontrivial planner optimum stated in this lemma implies the condition stated Lemma 3, completing the proof that the conditions are equivalent. *Q.E.D.*

Proof of Lemma 6

Since $B''(v) < 0$ for all $v \in [0, \infty)$, $B'(v)$ is strictly decreasing for all $v \in [0, \infty)$, implying $B'(v)$ has an inverse. It remains to show that $g(x^0)/\phi(x^0)$ is in the range of $B'(v)$. Suppose the optimum is nontrivial. Then $\rho^0(0) > 1$ by Lemma 5, implying $B'(0) > g(x^0)/\phi(x^0)$. Since $\lim_{v \uparrow \infty} B'(v) = 0$, there exists $v' > 0$ such that $B'(v') < g(x^0)/\phi(x^0)$. Since B' is continuous and decreasing, by the Intermediate Value Theorem, there exists $v \in (0, v')$ such that $B'(v) = g(x^0)/\phi(x^0)$. Thus, the inverse $(B')^{-1}\big(g(x^0)/\phi(x^0)\big)$ exists in $(0, v') \subset (0, \infty)$. *Q.E.D.*

Proof of Lemma 7

The objective function maximized in (15) is

$$B\big(N\phi(x)\big) - Ng(x). \tag{A57}$$

We will show (A57) is negative at extreme values of x. Its left-hand limit is

$$\lim_{x\downarrow 0}\left[B\big(N\phi(x)\big)-Ng(x)\right]=B\big(N\phi(0)\big)-Ng(0)=-Nk. \tag{A58}$$

The second equality follows from $B\big(N\phi(0)\big)=B(0)=0$ since $\phi(0)=0$ and from $g(0)=s(0)+p\cdot 0+k=k$ since $s(0)=0$. To derive the right-hand limit of (A57), note

$$B\big(N\phi(x)\big)-Ng(x)<B(N)-Npx \tag{A59}$$

since $\phi(x)<1$ and $g(x)=s(x)+px+k>px$ when $k>0$. Taking limits, (A59) implies

$$\lim_{x\uparrow\infty}\left[B\big(N\phi(x)\big)-Ng(x)\right]\leq\lim_{x\uparrow\infty}[B(N)-Npx] \tag{A60}$$

$$=B(N)-Np\lim_{x\uparrow\infty}x \tag{A61}$$

$$=-\infty. \tag{A62}$$

Therefore, both right- and left-hand limits of (A57) are negative.

We next show that (A57) is positive for some interior $x\in(0,\infty)$ when $\rho^0(N)\geq 1$. By (12), this last inequality is equivalent to

$$B'\big(N\phi(x^0)\big)\phi(x^0)\geq g(x^0). \tag{A63}$$

We have

$$B\big(N\phi(x^0)\big)-Ng(x^0)=\int_0^N\left[B'\big(n\phi(x^0)\big)-g(x^0)\right]dn \tag{A64}$$

$$>\int_0^N\left[B'\big(N\phi(x^0)\big)-g(x^0)\right]dn. \tag{A65}$$

The first line follows from the Fundamental Theorem of Calculus. The second line follows from the concavity of B, which implies $B'\big(n\phi(x^0)\big)>B'\big(N\phi(x^0)\big)$ for all $n\in(0,N)$. The right-hand side of (A65) is positive because its integrand is positive by (A63). Thus, (A57) is positive for some interior x, namely x^0.

Equation (A57) is real-analytic since all terms are. Since (A57) is real-analytic, it is continuous, implying it a nontrivial interior optimum because it is negative at its extremes and positive for some interior point. The set of positive maximizers (A57) is thus nonempty. The set of positive maximizers is finite because (A57) is real-analytic. This finite set has a unique maximum value. *Q.E.D.*

Proof of Lemma 8

Throughout the proof, maintain the assumption that the planner's optimum is nontrivial, equivalent to $\rho^0(0)>1$ by Lemma 5. The objective function maximized in (18) is

$$B\big(Q\alpha_\phi(x)\big)-Q\alpha_g(x). \tag{A66}$$

We will prove that the set of maximizers of (A66), $X^s(Q)$, is nonempty and finite set, allowing us to conclude it has a unique largest element $\bar{x}^s(Q)$. We will further argue that $\bar{x}^s(Q)>0$ and that (A66) is positive when evaluated at $\bar{x}^s(Q)$.

We proceed by showing (A66) is negative at extreme values of x. Its left-hand limit is

$$\lim_{x\downarrow 0}\left[B\big(Q\alpha_\phi(x)\big)-Q\alpha_g(x)\right]=B\big(Q\phi'(0)\big)-Q\lim_{x\downarrow 0}\alpha_g(x)=-\infty. \tag{A67}$$

The first step follows from (A1). The second step follows from $\lim_{x\downarrow 0}\alpha_g(x)=g(0)/0=k/0=\infty$. The right-hand limit of (A66) is

$$\lim_{x\uparrow\infty}\left[B\big(Q\alpha_\phi(x)\big)-Q\alpha_g(x)\right]=B(0)-Q\lim_{x\uparrow\infty}\alpha_g(x)\leq -Qp. \tag{A68}$$

The first step follows from (A3). The second step follows from $B(0)=0$ and

$$\alpha_g(x)=\frac{s(x)+px+k}{x}>\frac{px}{x}=p \tag{A69}$$

since $s(x),k>0$. Thus, $\lim_{x\uparrow\infty}\alpha_g(x)\geq p$.

We next show that (A57) is positive for some interior $x\in(0,\infty)$ when $\rho^0(0)>1$. By (12), this last inequality is equivalent to

$$B'(0)\phi(\bar{x}^0)>g(\bar{x}^0). \tag{A70}$$

Since all functions are real-analytic, they are continuous. By continuity, there exists $\tilde{v}>0$ such that

$$B'(v)\phi(x^0)>g(x^0). \tag{A71}$$

for all $v\in[0,\tilde{v}]$. Take

$$\tilde{Q}=\frac{\tilde{v}}{\alpha_\phi(\bar{x}^0)}. \tag{A72}$$

Then, for all $Q\in(0,\tilde{Q}]$,

$$B'\big(Q\alpha_\phi(\bar{x}^0)\big)\phi(\bar{x}^0)-g(\bar{x}^0)>0. \tag{A73}$$

Consider a policy of distributing doses of size $\bar{x}^0$ to $Q/\bar{x}^0$ individuals, for $Q\in(0,\tilde{Q}]$. The social surplus from this policy equals

$$B\big(Q\alpha_\phi(\bar{x}^0)\big)-Q\alpha_g(\bar{x}^0)=\int_0^{Q/\bar{x}^0}\left[B'\big(n\phi(\bar{x}^0)\big)-g(\bar{x}^0)\right]dn \tag{A74}$$

$$>\int_0^{Q/\bar{x}^0}\left[B'\big(Q\alpha_\phi(\bar{x}^0)\big)-g(\bar{x}^0)\right]dn. \tag{A75}$$

Equation (A75) follows from the Fundamental Theorem of Calculus. Condition (A76) follows since $B''<0$, implying $B'\big(n\phi(\bar{x}^0)\big)>B'\big((Q/\bar{x}^0)\phi(\bar{x}^0)\big)$ for all $n\in(0,Q/\bar{x}^0)$. The integral in (A76) is positive since its integrand is by (A73). The proposed policy thus generates positive social surplus.

For $Q\in(0,\tilde{Q}]$, we have shown that (A66) is negative outside of some compact interval $[x',x'']$ where $0<x'<x''$ and positive for some values inside, implying that the maximum over $[x',x'']$ is also the maximum over $[0,\infty)$. Since all functions in (A66) are real-analytic, the expression is real-analytic, and thus continuous. A continuous function on a compact interval has a maximum by the Extreme Value Theorem. Since the objective is real-analytic, the set of maximizers is finite

and has a unique maximum, denoted $\bar{x}^s(Q)$. Since $\bar{x}^s(Q)$ maximizes (A66), it generates a greater value than $\bar{x}^0$:

$$B\big(Q\alpha_\phi(\bar{x}^s(Q))\big) - Q\alpha_g(\bar{x}^s(Q)) \geq B\big(Q\alpha_\phi(\bar{x}^0)\big) - Q\alpha_g(\bar{x}^0) > 0, \tag{A76}$$

where the last inequality holds because (A75) is positive. Since $\bar{x}^s(Q) \in [x', x'']$, $\bar{x}^s(Q) \geq x' > 0$. *Q.E.D.*

Proof of Proposition 2

Differentiating objective function (18) with respect to x and dividing out inessential constants yields the first-order necessary condition

$$\frac{d\bar{x}^s(Q)}{dQ} = \frac{-B''\big(Q\alpha_\phi(\bar{x}^s(Q))\,\alpha'_\phi(\bar{x}^s(Q))\,\alpha_\phi(\bar{x}^s(Q))\big)}{B''\big(Q\alpha_\phi(\bar{x}^s(Q))\,\alpha'_\phi(\bar{x}^s(Q))^2 + B'\big(Q\alpha_\phi(\bar{x}^s(Q))\,\alpha''_\phi(\bar{x}^s(Q))\big) - \alpha''_g(\bar{x}^s(Q))\big)}. \tag{A77}$$

A necessary condition for an optimum is that the objective function be locally concave at the optimum, implying the denominator is negative, since it is the derivative of the first-order condition with respect to x, so the second derivative of the objective function with respect to x. The numerator of (A77) has a negative factor $B''\big(Q\alpha_\phi(\bar{x}^s(Q))\big)$ and a positive factor $\alpha_\phi\big(\bar{x}^s(Q)\big)$. Thus,

$$\operatorname{sgn}\left(\frac{d\bar{x}^s(Q)}{dQ}\right) = -\operatorname{sgn}\Big(\alpha'_\phi\big(\bar{x}^s(Q)\big)\Big) = \operatorname{sgn}\big(\bar{x}^s(Q) - \bar{x}^e\big), \tag{A78}$$

where the second step follows from Lemma 2.

To complete the proof, we will show

$$\operatorname{sgn}\big(\bar{x}^s(Q) - \bar{x}^e\big) = \operatorname{sgn}\big(1 - \eta_g(\bar{x}^e)\big). \tag{A79}$$

We will verify (A79) holds in the case $\eta_g(\bar{x}^e) < 1$. Verifying (A79) holds in the other cases $\eta_g(\bar{x}^e) > 1$ and $\eta_g(\bar{x}^e) = 1$ requires similar arguments, which are omitted.

Suppose for the remainder of the proof that $\eta_g(\bar{x}^e) < 1$. We will first argue that $\eta_g(x) < 1$ for all $x \in [0, \bar{x}^e]$ as well. To see this, note that for all $x > 0$ such that $\eta_g(x) = 1$,

$$\eta'_g(x) = \frac{xg''(x)}{g(x)} + H(x,1) = \frac{xg''(x)}{g(x)} \geq 0, \tag{A80}$$

where the first step follows from (A54), the second step from (A52), and the third step from the weak convexity of g. Condition (A80) implies that $\eta_g(x)$ cannot be downward sloping anywhere it is unit elastic. Therefore, after hitting the unit-elastic threshold, $\eta_g(x)$ can never be less than unit elastic. Formally, if $\eta_g(\tilde{x}) \geq 1$ for some $\tilde{x} > 0$, then $\eta_g(x) \geq 1$ for all $x \geq \tilde{x}$. The contrapositive of this statement must also be true. The contrapositive can be written as follows: if $\eta_g(\tilde{x}) < 1$ for some $\tilde{x} > 0$, then $\eta_g(x) < 1$ for all $x \in [0, \tilde{x}]$. Taking $\tilde{x} = \bar{x}^e$, we have $\eta_g(x) < 1$ for all $x \in [0, \bar{x}^e]$.

Take any $x \in [0, \bar{x}^e]$. By Lemma 1, $\eta_\phi(x) \geq 1$. Therefore,

$$\operatorname{sgn}\big(\eta_\phi(x) - 1\big) > \operatorname{sgn}\big(\eta_g(x) - 1\big). \tag{A81}$$

Equation (A19) implies $\operatorname{sgn}\big(\eta_\phi(x)-1\big)=\operatorname{sgn}\big(\alpha'_\phi(x)\big)$. A similar calculation can be used to show $\operatorname{sgn}\big(\eta_g(x)-1\big)=\operatorname{sgn}\big(\alpha'_g(x)\big)$. Subsituting these equalities into (A81) yields

$$\operatorname{sgn}\big(\alpha'_\phi(x)\big)>\operatorname{sgn}\big(\alpha'_g(x)\big). \tag{A82}$$

A necessary condition for x to satisfy (20) is that $\operatorname{sgn}\big(\alpha'_\phi(x)\big)=\operatorname{sgn}\big(\alpha'_g(x)\big)$ since $B'>0$. By (A82), therefore, x cannot satisfy (20). Hence, $x\neq x^s(Q)$. Since x is an arbitrary element of $[0,\bar{x}^e]$, we have $x^s(Q)\notin[0,\bar{x}^e]$, implying $\bar{x}^s(Q)\notin[0,\bar{x}^e]$, implying $\bar{x}^s(Q)>\bar{x}^e$ implying $\operatorname{sgn}\big(\bar{x}^s(Q)-\bar{x}^e\big)=1=\operatorname{sgn}\big(1-\eta_g(\bar{x}^e)\big)$. This verifies (A79) holds in the case $\eta_g(\bar{x}^e)<1$. The arguments that (A79) holds in other cases is similar and omitted. *Q.E.D.*

Proof of Proposition 6

KKT Conditions. Including a cost $g(x)$ of vaccinating individuals, the Lagrangian associated with the planner's maximization problem becomes

$$\mathcal{L}=B\big(n\phi(x)\big)-ng(x)+\lambda(Q-nx)+\mu(N-n). \tag{A83}$$

We will maintain condition $\rho^0(0)>1$ throughout the proof. Then $n^*,x^*>0$ by Lemma 5. Thus, nonnegativity constraint (5) can be omitted from the problem. The remaining KKT conditions are

$$\frac{\partial\mathcal{L}}{\partial n}=B'\big(n^*\phi(x^*)\big)\phi(x^*)-g(x^*)-\lambda^*x^*-\mu^*=0 \tag{A84}$$

$$\frac{\partial\mathcal{L}}{\partial x}=B'\big(n^*\phi(x^*)\big)n^*\phi'(x^*)-n^*g'(x^*)-\lambda^*n^*=0 \tag{A85}$$

along with (A29)–(A34), unchanged from Proposition 1.

Necessity of KKT Conditions. The constraints are identical to those in Proposition 1. Thus, the proof there that the KKT conditions are necessary for an optimum, which relied on the constraints satisfying LICQ, is identical here.

Optimum When $\rho^0(N)>1$. We next move to deriving the optimum in each of the two cases delineated in the statement of the proposition. First, consider the case $\rho^0(N)>1$. Then

$$1<\rho^0(N)=B'\big(N\phi(\bar{x}^0)\big)\frac{\phi(\bar{x}^0)}{g(\bar{x}^0)}, \tag{A86}$$

where the equality follows from the definition (12).

Suppose for the sake of contradiction that neither constraint (3) or (4) binds. Then $\lambda^*=\mu^*=0$. Substituting these values into (A84) yields $B'\big(n^*\phi(x^*)\big)\phi(x^*)=g(x^*)$. Rearranging,

$$1=B'\big(n^*\phi(x^*)\big)\frac{\phi(x^*)}{g(x^*)} \tag{A87}$$

$$=B'\big(n^0\phi(x^0)\big)\frac{\phi(x^0)}{g(x^0)} \tag{A88}$$

$$=B'\big(n^0\phi(x^0)\big)\frac{\phi(\bar{x}^0)}{g(\bar{x}^0)} \tag{A89}$$

Equation (A88) follows since (n^*,x^*) must be one of the solutions (n^0,x^0) to the unconstrained problem (8). Equation (A89) follows because x^0 and $\bar{x}^0$ both maximize $\phi(x)/g(x)$ and thus must generate the same value of this ratio:

$$\frac{\phi(x^0)}{g(x^0)} = \frac{\phi(\bar{x}^0)}{g(\bar{x}^0)} \tag{A90}$$

Together, (A86)–(A89) imply

$$B'\big(N\phi(\bar{x}^0)\big) > B'\big(n^0\phi(x^0)\big), \tag{A91}$$

implying $N\phi(\bar{x}^0) < n^0\phi(x^0)$ since $B'' < 0$, implying $n^0 > N$ since $\phi'(x) > 0$ and since $\phi(\bar{x}^0) \geq \phi(x^0)$ as $\bar{x}^0$ is as large as any other maximizing dosage by (10). But $n^0 > N$ violates KKT condition (A30), a contradition proving either (3) or (4) binds.

For Q sufficiently high, constraint (3) does not bind. Thus, (4) must bind, implying $n^* = N$. The problem reduces to maximizing objective (2) ignoring constraint (3) fixing n at a certain value (here $n = N$). This is the problem solved in the universal-coverage benchmark, implying $x^* = x^u$ for some $x^u \in X^u$. The solution $(n^*,x^*) = (N,x^u)$ obtains as long as it satisfies (3), requiring $Nx^u = n^*x^* \leq Q$. This constraint is relaxed to the extent possible by choosing the smallest of the solutions, i.e., $x^* = \underline{x}^u$.

If $N\underline{x}^u > Q$, then the solution ignoring (3) violates (3), implying (3) binds. The resulting solution is determined by whether (4) also binds. If (4) does not bind, the problem reduces to maximizing objective (2) ignoring constraint (4), fully utilizing supply Q. This is exactly the problem solved in the short-supply benchmark, implying $x^* = x^s(Q)$ for some $x^s(Q) \in X^s(Q)$. Since (3) binds, $n^*x^* = Q$, implying $n^* = Q/x^* = Q/x^s(Q)$. This solution obtains as long as it satisfies (4), requiring $n^* = Q/x^s(Q) \leq N$. This constraint is relaxed to the extent possible by choosing the largest of the solutions, i.e., $x^* = \bar{x}^s(Q)$.

If $Q/x^s(Q) > N$, then both (3) and (4) bind. We are left with a system of two equations in two unknowns: $n^* = N$ and $n^*x^* = Q$, implying $x^* = Q/n^* = Q/N$.

Optimum When $\rho^0(N) \leq 1$. Next, suppose $\rho^0(N) \leq 1$. Repeating (A86)–(A89) reversing the inequality yields

$$n^0 \leq N. \tag{A92}$$

The unconstrained solution (n^0,x^0) thus satisfies the population constraint (4). It is the optimum if it also satisfies supply constraint (3), requiring $n^0x^0 \leq Q$.

We will show that the largest solution $\bar{x}^0$ relaxes that constraint to the extent possible. Suppose there exist two solutions, x_1^0 and x_2^0 such that $x_1^0 < x_2^0$. By Lemma 4 we must have $x_1^0, x_2^0 < \bar{x}^e$, since if $x^0 \geq \bar{x}^e$ for any $x^0 \in X^0$, then X^0 is a singleton and cannot contain two unequal solutions. By Lemma 2, $\alpha'_\phi(x) > 0$ for all $x < \bar{x}^e$, implying $\alpha_\phi(x_2^0) > \alpha_\phi(x_1^0)$, which after rearranging implies

$$\phi(x_2^0) > \frac{\phi(x_1^0)x_2^0}{x_1^0}. \tag{A93}$$

Equation (14) implies

$$n_1^0\phi(x_1^0) = (B')^{-1}\left(\frac{g(x_1^0)}{\phi(x_1^0)}\right) = (B')^{-1}\left(\frac{g(\bar{x}^0)}{\phi(\bar{x}^0)}\right), \tag{A94}$$

where the second equality holds because all solutions in X^0 maximize $g(x)/\phi(x)$, so produce the same value for that ratio. Similarly,

$$n_2^0\phi(x_2^0) = (B')^{-1}\left(\frac{g(\bar{x}^0)}{\phi(\bar{x}^0)}\right). \tag{A95}$$

Thus,

$$n_1^0\phi(x_1^0) = n_2^0\phi(x_2^0) > \frac{n_2^0 x_2^0 \phi(x_1^0)}{x_1^0}, \tag{A96}$$

where the first step follows from equating the left-hand sides of (A94)–(A95) and the second step follows from substituting from (A93). After rearranging, (A96) implies $n_1^0 x_1^0 > n_2^0 x_2^0$, implying that the higher solution relaxes the omitted capacity constraint more than the lower solution, implying $\bar{x}^0$ relaxes the omitted capacity constraint to the extent possible.

Therefore, the pair $(\bar{n}^0, \bar{x}^0)$ is a solution to the planner's problem unless $\bar{n}^0\bar{x}^0 > Q$. In that case, the capacity constraint binds. Then, the remaining problem is the same as the short-supply benchmark. We saw above that the largest solution $\bar{x}^s(Q)$ relaxes omitted constraints to the extent possible. Hence, the solution when $\bar{n}^0\bar{x}^0 > Q$ is $x^* = \bar{x}^s(Q)$ and $n^* = Q/\bar{x}^s(Q)$. *Q.E.D.*

Proof of Proposition 7

The steps in the proof of Proposition 1 establishing the existence of nontrivial solution, characterizing the KKT conditions, and establishing the necessity of KKT conditions were independent of whether efficacy is increasing or peaked, so continue to apply here. The proof picks up here with the derivation of the optimum.

Suppose for the sake of contradiction that $\lambda^* = \mu^* = 0$. Then KKT condition (A27) implies $B'\big(n^*\phi(x^*)\big)\phi(x^*) = 0$, implying $\phi(x^*) = 0$ since $B'(v) > 0$ for all $v \geq 0$. But $\phi(x^*) = 0$ yields social surplus $B\big(n^*\phi(x^*)\big) = 0$, no better than the trivial solution, a contradiction to the existence of a nontrivial solution. Thus, one of λ^* or μ^* or both must be positive. Hence, the optimum must be one of the following three candidates analyzed next: (a) $\lambda^* > 0$, $\mu^* > 0$; (b) $\lambda^* > 0$, $\mu^* = 0$; (c) $\lambda^* = 0$, $\mu^* > 0$.

We analyzed candidates (a) and (b) in the proof of Proposition 1, showing candidate (a) entails $n^* = N$ and $x^* = Q/N$, and candidate (b) entails $x^* = x^e$ and $n^* = Q/x^e$. To characterize the candidate (c) solution, substituting $\lambda^* = 0$ into (A28) yields $B'\big(n^*\phi(x^*)\big)n^*\phi'(x^*) = 0$, implying $\phi(x^*) = 0$ since $B'(v) > 0$ for all $v \geq 0$ and $n^* > 0$ in a nontrivial optimum. But $\phi(x^*) = 0$ implies $x^* = x^p$. Substituting $\mu^* > 0$ into (A34) implies $n^* = N$.

It remains to piece together these candidates to construct the planner optimum. Suppose $Q \geq Nx^p$. Then we can show candidate (b) is infeasible, specifying an infeasible number of vaccinations $n^* = Q/x^e > Nx^p/x^e > N$ since $x^p > x^e$. To see this last inequality, note that $\phi'(x) \leq 0$ for all $x \geq x^p$, while $\phi'(x^e) = \phi(x^e)/x^e > 0$, implying $x^e \ngeq x^p$, implying $x^e < x^p$. One of the two remaining candidates, (a) or (c), must provide the optimum. Candidate (c) omits a constraint relative to case (a) so must provide a weakly better solution unless it violates the omitted capacity constraint. But the maintained condition $Q \geq Nx^p$ ensures candidate (c) respects the omitted capacity constraint.

Suppose $Q \leq Nx^e$. Then we can show candidate (c) is infeasible, requiring an infeasible capacity: $n^*x^* = Nx^p > Nx^e \geq Q$ since $x^p > x^e$ as shown in the previous paragraph. One of the two remaining candidates, (a) or (b), must provide the optimum. Candidate (a) omits a constraint

relative to case (b) so must provide a weakly better solution unless it violates the omitted population constraint. The number of vaccinations specified by candidate (a) is $n^* = Q/x^e \leq N$, where the inequality follows from the maintained condition $Q \leq Nx^e$. Hence candidate (a) respects the omitted population constraint and provides the optimum.

The remaining possibility is $Q \in (Nx^e, Nx^p)$. Candidate (b) is infeasible since it specifies the number of vaccinations $n^* = Q/x^e > Nx^e/x^e = N$, where the inequality follows from $Q > Nx^e$. Candidate (c) is infeasible since it requires capacity $n^*x^* = Nx^p > Q$. This leaves candidate (a) as the only feasible candidate, which therefore provides the optimum. *Q.E.D.*

Proof of Proposition 8

The proof of Proposition 6 only uses the assumption of everywhere-increasing ϕ in one step: the step showing (A91) implies $n^0 > N$. Here, we establish that step without assuming $\phi'(x) > 0$ for all $x > 0$, substituting an argument that works for peaked efficacy.

The key is to show $\phi(x^0) \leq \phi(\bar{x}^0)$ for all $x^0 \in X^0$ when efficacy is peaked. Equation (A90) holds since x^0 and $\bar{x}^0$ both maximize $\phi(x)/g(x)$. Since $g'(x) > 0$ and $\bar{x}^0 \geq x^0$, $g(\bar{x}^0) \geq g(x^0)$. Substituting this inequality into (A90) and rearranging yields $\phi(x^0) \leq \phi(\bar{x}^0)$.

We that argument in hand, we have (A91) implies $N\phi(\bar{x}^0) < n^0\phi(x^0)$ since $B'' < 0$, implying $n^0 > N$ since $\phi(\bar{x}^0) \geq \phi(x^0)$, as shown in the previous paragraph. But $n^0 > N$ violates KKT condition (A30), a contradition proving either (3) or (4) binds. The remainder of the proof of Proposition 6 applies since it does not rest on the assumption of increasing ϕ. *Q.E.D.*

Appendix A2. Calibrating Dose-Response Functions

This appendix provides details behind the calibration of the concave and sigmoid dose-response functions shown in Figure 8.

The text provided three conditions pinning down the dose-response curves:

$$\phi(1) = 0.95 \tag{A97}$$
$$\phi(0.5) = 0.8 \tag{A98}$$
$$\phi'(1) = 0. \tag{A99}$$

For the concave curve, we adopt the general functional form

$$\phi(x) = \alpha_1 x(1 - \alpha_2 x^{\alpha_3}). \tag{A100}$$

Numerically solving the system of equations (A97)–(A99) for the three unknown parameter values yields $\alpha_1 = 26.3$, $\alpha_2 = 0.964$, and $\alpha_3 = 0.0375$ to three significant digits.

For the sigmoid curve, we adapt a two-parameter specification from Furness (2018) with the convenient properties that its domain and image are in $[0,1]$, its derivatives are zero at both endpoints, and the parameters allow the location of the point of inflection and the steepness there to be controlled. Adding a parameter α_1 that allows efficacy to be scaled below 1 at a full dose, we obtain the three-parameter specification

$$\phi(x) = \begin{cases} \alpha_1 x^{\zeta(\alpha_3)} \alpha_2^{1-\zeta(\alpha_3)} & x \leq \alpha_2 \\ \alpha_1 \left[1 - (1-x)^{\zeta(\alpha_3)} (1-\alpha_2)^{1-\zeta(\alpha_3)}\right] & x > \alpha_2, \end{cases} \tag{A101}$$

where $\zeta(\alpha_3) = 2/(1-\alpha_3)$. Numerically solving the system of equations (A97)–(A99) for the three unknown parameter values yields $\alpha_1 = 0.950$, $\alpha_2 = 0.210$, and $\alpha_3 = 0.434$ to three significant digits. With this parameterization, one can compute the unit-elastic point as $\bar{x}^e = 0.4$.

Appendix A3. Calibration for Slow-Growing Epidemic

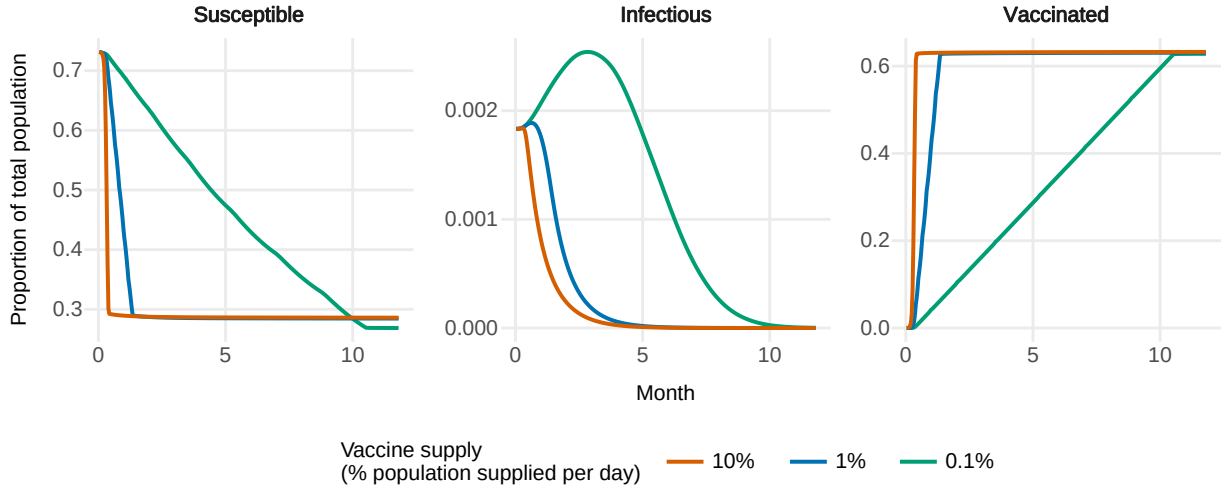


Figure A1: Evolution of Compartments in Calibrated Model in a Slow-Growing Epidemic. Reprises results from Figure 9, reducing $\mathscr{R}_0$ from 2 to 1.1. See that figure for additional notes.

Table A1: Calibration Results for Optimal Dosing and Averted Deaths in a Slow-Growing Epidemic

	Fraction of full dose X/x^p			Fraction of no-vaccine deaths averted $[D(0)-D(X)]/D(0)$		
Vaccine capacity (fraction of population supplied per day)	Full dose $X = x^p$	Uniform fractional dose $X = x^*$	Age-varying fractional dose $X = (x_a^*)$	Full dose $X = x^p$	Uniform fractional dose $X = x^*$	Age-varying fractional dose $X = (x_a^*)$
Sigmoid Dose Response						
0.01%	100%	41%	Fig. A3	6%	10%	10%
0.1%	100%	43%	Fig. A3	43%	57%	57%
1%	100%	60%	Fig. A3	85%	87%	87%
10%	100%	78%	Fig. A3	91%	91%	91%
Concave Dose Response						
0.01%	100%	3%	Fig. A3	6%	23%	23%
0.1%	100%	15%	Fig. A3	43%	65%	67%
1%	100%	55%	Fig. A3	85%	87%	87%
10%	100%	87%	Fig. A3	91%	91%	91%

Notes: Reprises results from Table A1, reducing $\mathscr{R}_0$ from 2 to 1.1. See that table for additional notes. The optimal age-varying dosage is not a single number but the vector plotted in Figure A3.

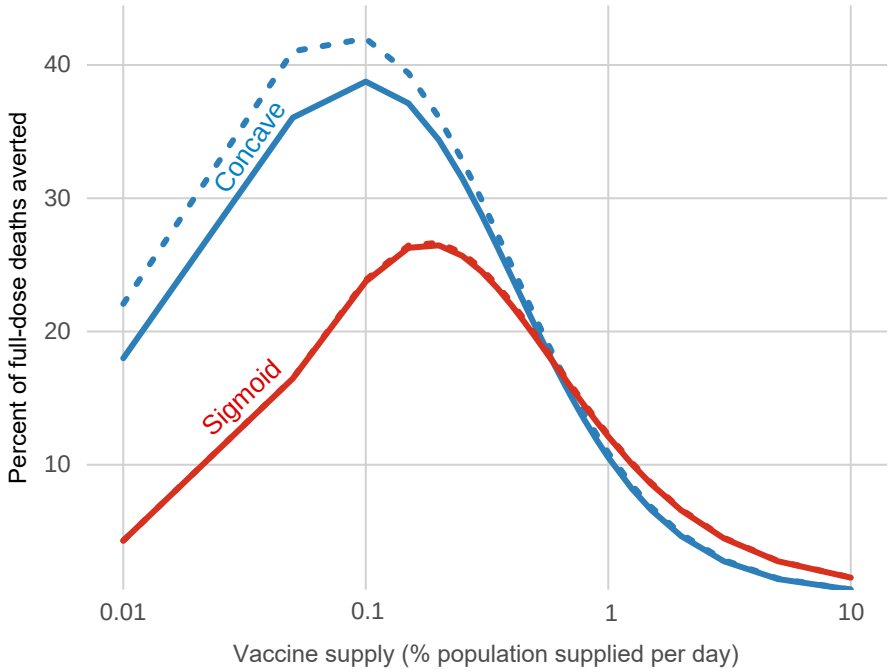


Figure A2: Deaths Averted Moving from Full to Fractional Dosing in a Slow-Growing Epidemic. Reprises results from Figure 10, reducing $\mathscr{R}_0$ from 2 to 1.1. See that figure for additional notes.

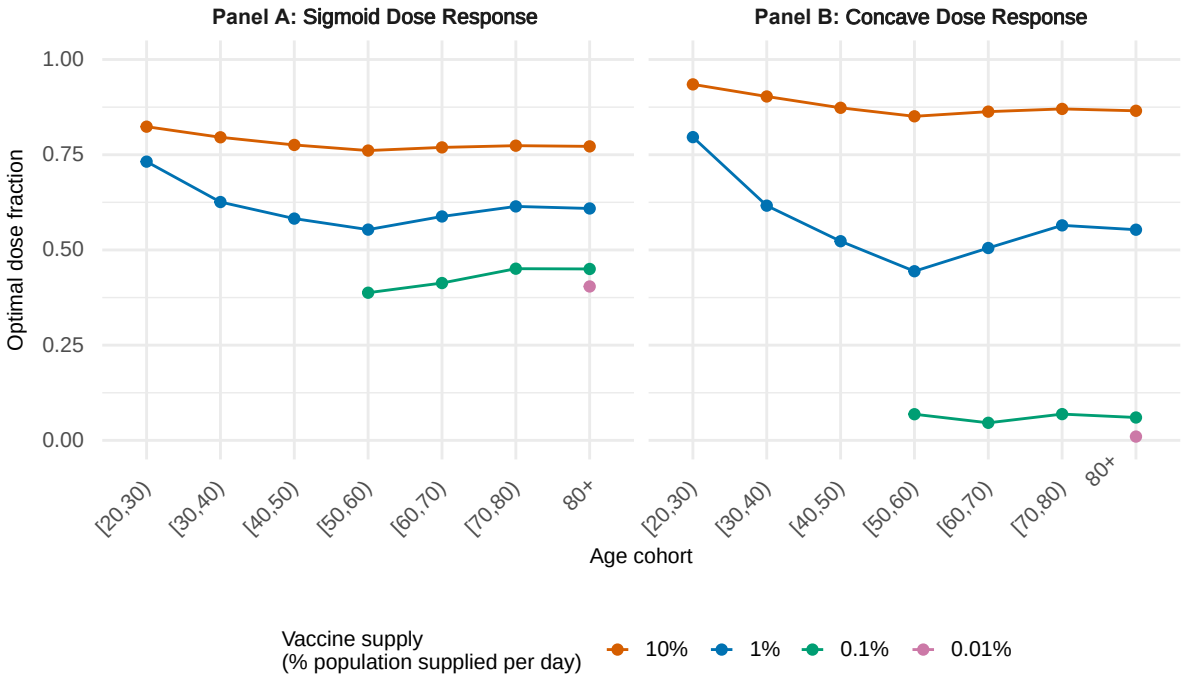


Figure A3: Optimal Age-Varying Fractional Doses in a Slow-Growing Epidemic. Reprises results from Figure 11, reducing $\mathscr{R}_0$ from 2 to 1.1. See that figure for additional notes.

References

Ahuja, Amrita, Susan Athey, Arthur Baker, Eric Budish, Juan Camilo Castillo, Rachel Glennerster, Scott D. Komienrs, Michael Kremer, Jean Lee, Canice Prendergast, Christopher M. Snyder, Alex Tabarrok, Brandon J. Tan, and Witold Więcek. (2021) "Preparing for a pandemic: Accelerating vaccine availability," *American Economic Association Papers and Proceedings* 111: 331–335.

Anderson, Roy. M. and Robert. M. May. (1992) *Infectious Diseases of Humans: Dynamics and Control*. New York: Oxford University Press.

An-Yeung, Byron B., Julie Zikherman, James L. Mueller, Judith F. Ashouri, Mehrdad Matloubian, Debra A. Cheng, Yiling Chen, Kevan M. Shokat, and Arthur Weiss. (2014) "A sharp T-cell antigen receptor signaling threshold for T-cell proliferation," *Proceedings of the National Academy of Sciences* 111 (35): E3679–E3688.

Atkeson, Andrew. (2021) "Behavior and the dynamics of epidemics," *Brookings Papers on Economic Activity* Spring 2021: 67–88.

Avery, Christopher. (2024) "The economics of social distancing and vaccination," *Review of Economic Design* 28: 781–812.

Avery, Christopher, Frederick Chen, and David McAdams. (2024) "Steady-state social distancing and vaccination," *American Economic Review: Insights* 6 (1): 1–19.

Azman, Andrew S., Francisco J. Luquero, Iza Ciglenecki, Rebecca F. Frais, David A. Sack, and Justin Lessler. (2015) "The impact of a one-dose versus two-dose oral cholera vaccine regimen in outbreak settings: A modeling study," *PLoS Medicine* 12 (8): e1001867.

Baden, Lindsey R., et al. (2020) "Efficacy and safety of the mRNA-1273 SARS-CoV-2 vaccine," *New England Journal of Medicine* 384 (5): 403–416.

Baylor, Norman W. and Valerie B. Marshall. (2013) "Regulation and testing of vaccines," chapter 73 in Stanley A. Plotkin, Walter A. Orenstein, and Paul A. Offit, eds., *Vaccines* (sixth edition): Elsevier.

Benest, John, Sophie Rhodes, Sara Afrough, Thomas Evans, and Richard White. (2020) "Response type and host species may be sufficient to predict dose-response curve shape for adenoviral vector vaccines," *Vaccines* 8: 155.

Billeskov, Rolf, Thomas Lindenstrøm, Joshua Woodworth, Cristina Vilaplana, Pere-Joan Cardona, Joseph P. Cassidy, Rasmus Mortensen, Else Marie Agger, and Peter Andersen. (2018) "High antigen dose is detrimental to post-exposure vaccine protection against tuberculosis," *Frontiers in Immunology* 8 (article 1973).

Borghans, José A. M., Leonie S. Taams, Marca H. M. Wauben, and Rob J. De Boer. (1999) "Competition for antigenic sites during T cell proliferation: A mathematical interpretation of *in vitro* data," *Proceedings of the National Academy of Sciences* 96: 10782–10797.

Castillo, Juan Camillo, Amrita Ahuja, Susan Athey, Arthur Baker, Eric Budish, Tasneem Chipty, Rachel Glennerster, Scott D. Kominers, Michael Kremer, Greg Larson, Jean Lee, Canice Prendergast, Christopher M. Snyder, Alex Tabarrok, Brandon J. Tan, and Witold Więcek. (2021) "Market design to accelerate COVID-19 vaccine supply," *Science* 371 (6534): 1107–1019.

Centers for Disease Control and Prevention. (2026) "Vaccine for monkeypox prevention in the United States," CDC report, March 10. Available at https://www.cdc.gov/monkeypox/hcp/vaccine-considerations/index.html.

Dalton, Alexandra F., et al. (2023) "Estimated effectiveness of JYNNEOS vaccine in preventing mpox: A multijurisdictional case-control study—United States, August 19, 2022–March 31, 2023," *Morbidity and Mortality Weekly Report* 70 (20): 553–558.

de Jonge, Jørge, Harry van Dijken, Femke de Jeij, Sanne Spijkers, Justin Mouthaan, Rineke de Jong, Paul Roholl, Eduardo Alfredo Adami, Milena Apetito Akamatsu, Paulo Lee Ho, Livia Brunner, Nicolas Collin, Martin Friede, José A. Ferreira, and Willem Luytjes. (2020) "H7N9 influenza split vaccine with SWE oil-in-water adjuvant greatly enhances cross-reactive humoral immunity and protection against severe pneumonia in ferrets," *NPJ Vaccines* 5: 38.

Dimitrov, Dobromir, Blythe Adamson, and Laura Matrajt. (2023) "Evaluation of mpox vaccine dose-sparing strategies," *PNAS Nexus* 2 (5): 1–4.

Dunn, David T., Richard Gilson, Sheena McCormack, and Laura E. McCoy. (2023) "Licenced doses of approved COVID-19 vaccines may not be optimal: A review of the early-phase, dose-finding trials," *Vaccine* 41 (31): 4472–4479.

Faksova, K., et al. (2024) "COVID-19 vaccines and adverse events of special interest: A multinational Global Vaccine Data Network (GVDN) cohort study of 99 million vaccinated individuals," *Vaccine* 42 (9): 2200–2211.

Farrington, C. P. (2003) "On vaccine efficacy and reproduction numbers," *Mathematical Biosciences* 185 (1): 89–109.

Furness, James. (2018) "Sigmoid functions in game design," [Internet]. Available at https://www.jfurness.uk/sigmoid-functions-in-game-design.

Gans, Joshua S. (2022) "The economic consequences of $\hat{\mathscr{R}} = 1$: Toward a workable behavioural epidemiological model of pandemics," *Review of Economic Analysis* 14 (1): 3–25.

Goodkin-Gold, Matthew, Michael Kremer, Christopher M. Snyder, and Heidi Williams. (2024) "Optimal vaccine subsidies for epidemic diseases," *Review of Economics and Statistics* 106 (4): 895–909.

Halloran, M. Elizabeth, Michael Haber, and Ira M. Longini, Jr. (1992) "Interpretation and estimation of vaccine efficacy under heterogeneity," *American Journal of Epidemiology* 136 (3): 328–343.

Imai, Natsuko, Thomas Rawson, Edward S. Knock, Raphael Sonabend, Yasin Elmaci, Pablo N. Perez-Guzman, Lilith K. Whittles, Divya Thekke Kanapram, Katy A. M. Gaythorpe, Wes Hinsley, Bimandra A. Djaafara, Haowei Wang, Keith Fraser, Richard G. FitzJohn, Alexandra B. Hogan, Patrick Doohan, Azra C. Ghani, Meil M. Ferguson, Marc Baguelin, and Anne Cori. (2023) "Quantifying the effect of delaying the second COVID-19 vaccine dose in England: a mathematical modelling study," *Lancet* 8 (3): e174–e183.

Itoh, Yasushi, and Ronald N. Germain. (1997) "Single cell analysis reveals regulated hierarchical T cell antigen receptor signaling thresholds and intraclonal heterogeneity for individual cytokine responses of CD4+ T cells," *Journal of Experimental Medicine* 186 (5): 757–766.

Juan-Giner, Aitana, Derick Kimathi, Kyra H. Grantz, Mainga Hamaluba, Patrick Kazooba, Patricia Njuguna, Gamou Fall, Moussa Dia, Ndeye S. Bob, Thomas P. Monath, Alan D. Barrett, Joachim Hombach, Edgar M. Mulogo, Immaculate Ampeire, Henry K. Karanja, Dan Nyehangane, Juliet Mwanga-Amumpaire, Derek A. T. Cummings, Philip Bejon, George M. Warimwe, and Rebecca F. Graisner, (2021) "Immunogenicity and safety of fractional doses of yellow fever vaccines: A randomised, double-blind, non-inferiority trial," *Lancet* 397 (10269): 119–127.

Kahneman, Daniel, Jack L. Knetsch, and Richard Thaler. (1986) "Fairness as a constraint on profit seeking: Entitlements in the market," *American Economic Review* 76 (4): 728–741.

Karlin, Samuel, and Herman Rubin. (1956) "The theory of decision procedures for distributions with monotone likelihood ratio," *Annals of Mathematical Statistics* 27 (2): 272–299.

Lakshmikantham, V. and S. Leela. (1969) *Differential and Integral Inequalities: Theory and Applications*. New York: Academic Press.

LaTorre, Davide, Tufail Malik, and Simone Marsiglio. (2020) "Optimal control of prevention and treatment in a basic macroeconomic–epidemiological model," *Mathematical Social Sciences* 108: 100–108.

Leung, Tiffany, Julia Eaton, and Laura Matrajt. (2022) "Optimizing one-dose and two-dose cholera vaccine allocation in outbreak settings: A modeling study," *PLoS Neglected Tropical Diseases* 16 (4): e0010358.

Manski, Charles F. (2020) "Forming COVID-19 policy under uncertainty," *Journal of Benefit Cost Analysis* 11 (3): 341–356.

Matrajt, Laura, Tom Britton, M. Elizabeth Halloran, and Ira M. Longini, Jr. (2015) "One versus two doses: What is the best use of vaccine in an influenza pandemic?" *Epidemics* 13: 17–27.

Matrajt, Laura, Julia Eaton, Tiffany Leung, Dobromir Dimitrov, Joshua T. Schffer, David A. Swan, and Holly Janes. (2021) "Optimizing vaccine allocation for COVID-19 vaccines shows the potential role of single-dose vaccination," *Nature Communications* 12: 3449.

McAdams, David. (2021) "The blossoming of economic epidemiology," *Annual Reviews of Economics* 13: 539–570.

Mossong, Joël, Niel Hens, Mark Jit, Philippe Beutels, Kari Auranen, Rafael Mikolajczyk, Marco Massari, Stefania Salmaso, Gianpaolo Scalia Tomba, Jacco Mallinga, Janneke Heijne, Malgorzata Sadkowska-Todys, Magdalena Rosinska, and W. John Edmunds. (2008) "Social contacts and mixing patterns relevant to the spread of infectious diseases," *PLOS Medicine* 5 (3): e74.

Okayasu, Hiromasa, Carolyn Sein, Diana Chang Blanc, Alejandro Ramirez Gonzalez, Darin Zehrung, Courtney Jarrahian, Grace Macklin, and Roland W. Sutter. (2017) "Intradermal administration of fractional doses of inactivated poliovirus vaccine: A dose-sparing option for polio immunization," *Journal of Infectious Disease* 216 (supplement 1): S161–S167.

Our World in Data. (2024) "Coronavirus (COVID-19) Vaccinations," global dataset compiled from official national reports. Available at https://ourworldindata.org/covid-vaccinations.

Pan-American Health Organization. (2018) "Brazil launches world's largest campaign with fractional-dose yellow fever vaccine," PAHO report, January 25. Available at https://www.paho.org/hq/index.php?option=com_content&view=article&id=14065:brazil-launches-worlds-largest-campaign-with-fractional-dose-yellow-fever-vaccine&Itemid=135&lang=en.

Philipson, Tomas. (2000) "Economic epidemiology and infectious diseases," chapter 33 in Anthony J. Culyer and Joseph P. Newhouse, eds., *Handbook of Health Economics, Volume 1*: Elsevier.

Philipson, Tomas. (2016) "Economic epidemiology," in *The New Palgrave Dictionary of Economics*: Palgrave Macmillan.

Polack, Fernando. P., et al. (2020) "Safety and efficacy of the BNT162b2 mRNA Covid-19 vaccine," *New England Journal of Medicine* 383 (27): 2603–2615.

Pompano, Rebecca R., Jianjun Chen, Emily A. Verbus, Huifang Han, Arthur Fridman, Tessie McNeely, Joel H. Collier, and Anita S. Chong. (2014) "Titrating T cell epitopes within self-assembled vaccines optimizes CD4+ helper T cell and antibody outputs," *Advanced Healthcare Materials* 3 (11): 1898–1908.

Rhodes, Sophie, Neal Smith, Thomas Evans, and Richard White. (2022) "Identifying COVID-19 optimal vaccine dose using mathematical immunostimulation/immunodynamic modelling," *NPJ Vaccines* 40 (49): 7032–7041.

Rhodes, Sophie J., Andrea Zelmer, Gwenan M. Knight, Satria Arief Prabowo, Lisa Stockdale, Thomas G. Evans, Thomas Lindenstrøm, Richard G. White, and Helen Fletcher. (2016) "The TB vaccine H56 + IC31 dose-response curve is peaked not saturating: Data generation for new mathematical modelling methods to inform vaccine dose decisions," *Vaccine* 34 (50): 6285–6291.

Roozendall, Ramon, et al. (2021) "SARS-CoV-2 binding and neutralizing antibody levels after Ad26.COV2.S vaccination predict durable protection in rhesus macaques," *Nature Communications* 12: 5877.

Riley, Steven, Joseph T. Wu, and Gabriel M. Leung. (2007) "Optimizing the dose of pre-pandemic influenza vaccines to reduce the infection attack rate," *PLoS Medicine* 4 (6): e218.

Sakaguchi, Shimon, Tomoyuki Yamaguchi, Takashi Nomura, and Masahiro Ono. (2008) "Regulatory T cells and immune tolerance," *Cell* 133 (5): 775–787.

Solodov, Mikhail V. (2011) "Constraint qualification," in James J. Cochran, ed., *Encyclopedia of Operations Research and Management Science*. New York: Wiley.

Toxvaerd, Flavio and Robert Rowthorn. (2022) "On the management of population immunity," *Journal of Economic Theory* 204: 105501.

World Health Organization. (2020a) "WHO SAGE roadmap for prioritizing uses of COVID-19 vaccines in the context of limited supply (Version 1.1, 13 November 2020)," WHO report, September 14. Available at https://www.who.int/docs/default-source/immunization/sage/covid/sage-prioritization-roadmap-covid19-vaccines.pdf.

World Health Organization. (2020b) "WHO SAGE Values Framework for the allocation and prioritisation of COVID-19 vaccination," WHO report, September 14. Available at https://iris.who.int/bitstream/handle/10665/334299/WHO-2019-nCoV-SAGE_Framework-Allocation_and_prioritization-2020.1-eng.pdf?ua=1.

World Health Organization. (2022a) "Human papillomavirus vaccines: WHO position paper (2022 update)," *Weekly Epidemiological Record* 97 (50): 645–672.

World Health Organization. (2022b) "One-dose Human Papillomavirus (HPV) vaccine offers solid protection against cervical cancer.,' WHO report, April 11. Available at https://www.who.int/news/item/11-04-2022-one-dose-human-papillomavirus-(hpv)-vaccine-offers-solid-protection-against-cervical-cancer.

Więcek, Witold, Amrita Ahuja, Esha Chaudhuri, Michael Kremer, Alexandre Simoes Gomes, Christopher M. Snyder, Alex Tabarrok, and Brandon J. Tan. (2022) "Testing fractional doses of COVID-19 vaccines'," *Proceedings of the National Academy of Sciences* 119 (8): e2116932119.

Więcek, Witold, David Johnston, Tomas Dulka, Danny Toomey, and Enlli Lewis. (2023) "Vaccines at velocity: Evaluating potential lives saved by earlier vaccination in the COVID-19 pandemic," medRxiv reprint, June 30. Available at https://doi.org/10.1101/2023.06.16.23291442.

Williams, Jane, Chris Degeling, Joie McVernon, and Angus Dawson. (2021) "How should we conduct pandemic vaccination?" *Vaccine* 39 (6): 994–999.

Wu, Joseph T., Corey M. Peak, Babriel M. Leung, and Marc Lipsitch. (2016) "Fractional dosing of yellow fever vaccine to extend supply: A modelling study," *Lancet* 388: 2904–2911.

Verity, Robert, et al. (2020) "Estimates of the severity of coronavirus disease 2019: A model-based analysis," *Lancet Infectious Diseases* 20 (6): 669–677.