This page is intentionally blank. Proper e-companion title page, with INFORMS branding and exact metadata of the main paper, will be produced by the INFORMS office when the issue is being assembled.

Appendix A: Literature Review - Supplement

Several studies have been conducted in the clinical context to evaluate the immunogenic effects of fractional-dose vaccines corresponding to various diseases and infections. Table EC.1 provides a summary of these studies.

Disease	Study	Vaccine type	Summary of Results
Seasonal Influenza	Palache et al. (1993)	Dose-ranging ($1/10^{\text{th}}$ to six-fold of the usual dose)	Meta analysis: finds consistent result that 33% decrease in vaccine dosage from usual is not associated with a meaningful decrease in immune response
	Treanor et al. (2002)	Half-dose	Drop in immune response was statistically significant but not "substantial" in adults <50 years
	Belshe et al. (2004)	40% of usual dose, intradermal	Elicited "similarly vigorous immune response" as the full-dose intramuscular vaccine in adults <60 years
	Kenney et al. (2004)	$1/5^{\text{th}}$ of usual dose, intradermal	Elicited immune response that was "similar to or better than" the full-dose intramuscular vaccine
	Engler et al. (2008)	Half-dose	Immune response of healthy adults <50 years was not "substantially inferior" to the full-dose vaccine
	Vajo et al. (2017)	Dose-ranging (quarter-dose to full-dose)	All doses $\geq$40% of the full-dose vaccine showed no significant difference in immune response compared with the full-dose vaccine
H1N1-Pandemic Influenza	Greenberg et al. (2009)	Full and double-dose, one dose regimen	A single full-dose vaccine was able to elicit sufficient antibody response, even though two doses are usually required for novel virus strains
	Clark et al. (2009)	Half-dose, one dose regimen	Elicited immune response that is consistent with protection against the H1N1 pandemic strain in adults
	Plennevaux et al. (2010)	Half-dose, one dose regimen	Elicited immune response that exceeds FDA's requirement in adults >18 years
	Liang et al. (2010)	Half-dose, one dose regimen	Elicited nearly the same immune response as the full-dose vaccine in individuals >12 years
	Vajo et al. (2010)	40% dose, one-dose regimen	Elicited immune response that fulfilled all international licensing criteria in adults >18 years
H5N1-Avian Influenza	Langley et al. (2011)	Quarter-dose, two dose regimen	Elicited immune response that fulfilled US and European international licensing criteria in adults >18 years

Table EC.1 List of clinical studies that evaluated the immunogenicity of fractional-dose vaccines

Appendix B: Proof of Theoretical Results

B.1. Optimal Control Problem Setup

We first present some notation to simplify the presentation of the results. We define the vector $c = (c_1, c_2)$ to denote the vaccine dosages and $\epsilon = (\epsilon_1, \epsilon_2)$ for the vaccine efficacies. We use $x = (S, I, v)$ for the state variables, $x_0 = (S_0, I_0, v_0)$ for the initial conditions, and $u = (u_1, u_2)$ for the control variables. We rewrite (1) as $\dot{x} = f(x, u)$, where $f(x, u) = (-\beta SI - \epsilon^T u,\ \beta SI - \gamma I,\ c^T u)$. From (OCP), we use $F(x, u) = -\beta SI$ to denote the rate of the objective function. Next, we define $g(u) = (u_{max} - u_1 - u_2,\ u_1,\ u_2)$ for the control constraints and $h(x) = (\alpha t - v, S)$ for the state constraints. For the terminal constraints, we define $a(x(T)) = v_{max} - v(T)$ and $b(x(T)) = I_{min} - I(T)$.

Note that earlier we had not defined $I(T) = I_{min}$ explicitly as a constraint in (OCP). Recall our definition of the end time of an epidemic: $T = \inf\{t \mid I(t) = I_{min}\}$, where I_{min} is an arbitrarily chosen constant. Thus, we have an optimal control problem with free terminal time. This implicitly makes T a decision variable in (OCP) with $I(t) = I_{min}$ as an additional terminal constraint.

Using this notation, we can redefine (OCP) as

$$\begin{aligned} \max_{\pi} \ & \int_0^T F(x(t), u(t))dt \\ s.t.\ & \dot{x} = f(x, u),\ x(0) = x_0 \\ & g(u(t)) \geq 0,\ h(x(t)) \geq 0 \\ & a \geq 0,\ b \geq 0 \end{aligned} \tag{OCP}$$

The Hamiltonian for the system can be formulated using the objective function and the state equations in (1) as follows.

$$\begin{aligned} H(x, u, \lambda) &= F(x, u) + \lambda^T f(x, u) \\ &= -\beta SI + \lambda_S\ (-\beta SI - \epsilon_1 u_1 - \epsilon_2 u_2) + \lambda_I\ (\beta SI - \gamma I) + \lambda_v\ (c_1 u_1 + c_2 u_2) \end{aligned}$$

where $\lambda = (\lambda_S,\ \lambda_I,\ \lambda_v)$ are the adjoint variables corresponding to each of the state equations. We can rearrange the terms in the Hamiltonian to express it as

$$H(x,u,\lambda) = -\beta SI\ (1+\lambda_S-\lambda_I) - \gamma I\lambda_I + \omega_1 u_1 + \omega_2 u_2 \tag{H}$$

where $\omega_1 = \lambda_v c_1 - \lambda_S \epsilon_1$ and $\omega_2 = \lambda_v c_2 - \lambda_S \epsilon_2$.

The Lagrangian can be formulated by augmenting the Hamiltonian with the constraints of (OCP). We use the indirect method as outlined in Hartl et al. (1995).

$$\begin{aligned} L(x,u,\lambda,\mu,\eta) &= H(x,u,\lambda) + \mu^T g(u) + \eta^T \dot{h}(x) \\ &= -\beta SI\ (1+\lambda_S-\lambda_I) - \gamma I\lambda_I + \omega_1 u_1 + \omega_2 u_2 + \mu_0(u_{max} - u_1 - u_2) + \\ &\quad \mu_1 u_1 + \mu_2 u_2 + \eta_v(\alpha - c_1 u_1 - c_2 u_2) + \eta_S(-\beta SI - \epsilon_1 u_1 - \epsilon_2 u_2) \end{aligned} \tag{L}$$

where $\mu = (\mu_0, \mu_1, \mu_2)$ and $\eta = (\eta_v, \eta_S)$ denote the Lagrangian variables corresponding to the control and state constraints respectively. Note that λ, μ and η are functions of time. We omit the time-dependency, as for the state and control variables, for notation brevity.

B.2. Pontryagin's Maximum Principle

We use Pontryagin's Maximum Principle (PMP) to analyse the necessary conditions for optimality. The application of PMP to (OCP) is summarized in the following result.

THEOREM EC.1. *Let π^* be an optimal solution to* (OCP) *and let u^* and x^* denote the corresponding optimal control and state trajectories. Then, there exist piecewise continuously differentiable adjoint functions $\lambda^* : [0,T] \to \mathbb{R}^3$, piecewise continuous Lagrangian functions $\mu^* : [0,T] \to \mathbb{R}^3$ and $\eta^* : [0,T] \to \mathbb{R}^2$ and jump parameters $\zeta^*(\tau) : [0,T] \to \mathbb{R}^3$ at each τ where λ^* is discontinuous, such that the following conditions hold.*

1. *Ordinary Differential Equation Condition:*

$$\dot{x}^* = f(x^*, u^*),\ x^*(0) = x_0; \tag{EC.1}$$

2. *Adjoint Condition:*

$$\dot{\lambda}^* = -\nabla_x L(x^*, u^*, \lambda^*, \mu^*, \eta^*); \tag{EC.2}$$

3. *Transversality Condition:*

$$\begin{aligned} \lambda^*(T^-) &= p_a \nabla_x a(x^*(T)) + p_b \nabla_x b(x^*(T)) + p_h \nabla_x h(x^*(T)), \\ &\text{with } p_a \geq 0,\ p_a a(x^*(T)) = 0,\ p_h \geq 0,\ p_h^T h(x^*(T)) = 0; \end{aligned} \tag{EC.3}$$

where $p_a, p_b, p_h = (p_v, p_S)$ are constants.

$$H(x^*(T), u^*(T^-), \lambda^*(T^-)) + p_a \nabla_T a(x^*(T)) + p_b \nabla_T b(x^*(T)) + p_h \nabla_T h(x^*(T)) = 0 \tag{EC.4}$$

4. *Hamiltonian Maximization Condition:*

$$H(x^*(t), u^*(t), \lambda^*(t)) \geq H(x^*(t), u(t), \lambda^*(t)) \quad \forall\, t \in [0,T]; \tag{EC.5}$$

for each $u(t)$ satisfying for all $t \in [0,T]$

$$g(u(t)) \geq 0$$

$$\dot{h}(x^*(t), u(t)) \geq 0 \text{ whenever } h(x^*(t)) = 0$$

5. *Jump Condition:*

$$\begin{aligned} \lambda(\tau^-) &= \lambda(\tau^+) + \zeta(\tau) \nabla_x h(x^*(\tau)), \\ \zeta(\tau) &= 0 \text{ for every } \tau \text{ where } \dot{h}(x^*(\tau)) \neq 0, \\ H(x^*(\tau), u^*(\tau^-), \lambda^*(\tau^-)) &= H(x^*(\tau), u^*(\tau^+), \lambda^*(\tau^+)) - \zeta(\tau) \nabla_t h(x^*(\tau)); \end{aligned} \tag{EC.6}$$

6. *Lagrange Multiplier Condition:*

$$\left.\frac{\partial L}{\partial u}\right|_{u=u*} = 0, \quad \frac{dH}{dt} = \frac{dL}{dt} = \frac{\partial L}{\partial t}; \tag{EC.7}$$

7. *Complementary Slackness Condition:*

$$\begin{aligned} &\mu(t) \geq 0, \quad \mu(t)^T g(u^*(t)) = 0, \\ &\eta(t) \geq 0, \quad \eta(t)^T h(x^*(t)) = 0, \\ &\zeta(\tau) \geq 0, \quad \zeta(\tau)^T h(x^*(\tau)) = 0; \end{aligned} \tag{EC.8}$$

Let us analyze the necessary conditions laid out in Theorem EC.1 in further detail. From (EC.2), we get the following system of equations for the adjoint variables.

$$\begin{aligned} \dot{\lambda}_S &= -\frac{\partial L}{\partial S} = \beta I(1 + \lambda_S - \lambda_I + \eta_S) \\ \dot{\lambda}_I &= -\frac{\partial L}{\partial I} = \beta S(1 + \lambda_S - \lambda_I + \eta_S) + \gamma \lambda_I \\ \dot{\lambda}_v &= -\frac{\partial L}{\partial v} = 0 \end{aligned} \tag{EC.9}$$

where $\eta_S(t) = 0$ for all t such that $S(t) > 0$. From (EC.3), we get the following terminal conditions for the adjoint variables.

$$\lambda_S(T) = p_S, \quad \lambda_I(T) = p_I, \quad \lambda_v(T) = -p_v - p_a \tag{EC.10}$$

where $p_S, p_v, p_a \geq 0$, $p_S S(T) = 0$, $p_a(v(T) - v_{max}) = 0$, and $p_v(v(T) - \alpha T) = 0$. Note that the maximum value that $v(T)$ can take is v_{max} while T can be made arbitrarily large. This way we can ensure that $v(T) < \alpha T$ holds for any set of parameter conditions, thus ensuring $p_v = 0$. From (EC.4), we have

$$\begin{aligned} &H(x^*(T), u^*(T), \lambda^*(T)) + p_v \alpha = 0 \\ \implies &H(x^*(T), u^*(T), \lambda^*(T)) = 0 \end{aligned} \tag{EC.11}$$

(EC.5) along with the constraints in (OCP) lead to the following linear program that can be solved at every time $t \in [0, T]$ to obtain $u^*(t)$.

$$\begin{aligned} \max_{u_1, u_2} \quad & \omega_1 u_1 + \omega_2 u_2 \\ s.t. \quad & u_1 + u_2 \leq u_{max} \\ & u_1,\ u_2 \geq 0 \\ & c_1 u_1 + c_2 u_2 \leq \alpha \quad \text{whenever } v(t) = \alpha t \\ & u_1 = u_2 = 0 \quad \text{whenever } S(t) = 0 \end{aligned} \tag{LP}$$

Let us enumerate all possible corner points of the linear program (LP): $(0,0)$, $(u_{max}, 0)$, $(0, u_{max})$, $(\alpha/c_1, 0)$, $(0, \alpha/c_2)$, $\left(\frac{\alpha - c_2 u_{max}}{c_1 - c_2}, \frac{c_1 u_{max} - \alpha}{c_1 - c_2}\right)$. It is important to note that the optimal control takes the form of piecewise constant functions. We will use this later in our analysis.

B.3. Intermediate Results

Consider the jump conditions in (EC.6). The complimentary slackness conditions in (EC.8) ensure that the jump parameters $\zeta(\tau)$ are zero for every τ where the state constraints $h(x(\tau))$ satisfy with strict inequality. Therefore, to evaluate the values of $\zeta(\tau)$, we only need to consider the times where $h(x(\tau))$ become tight.

- $\zeta_S(\tau)$: We see from (1) that $S(t)$ is non-increasing. Let τ be the first time that we have $S(t) = 0$. Then, we must have $S(t) > 0 \ \forall\ t \in [0, \tau)$ and $S(t) = 0 \ \forall\ t \in [\tau, T]$. Consequently, $\dot{S}(\tau^-) = -\epsilon_1 u_1 - \epsilon_2 u_2 < 0$ while $\dot{S}(\tau^+) = 0$, implying that $S(t)$ is not differentiable at $t = \tau$. Then, by (EC.6), we must have $\zeta_S(\tau) = 0$.
- $\zeta_I(\tau)$: There are no state constraints in $I(t)$. Hence, we do not have discontinuities in $\lambda_I(t)$.

- $\zeta_v(\tau)$: Let τ be the time at which the constraint $\alpha t - v(t) \geq 0$ becomes tight. At $t = \tau^-$, the corner points of the feasible region of (LP) are $(0,0)$, $(0, u_{max})$ and $(u_{max}, 0)$. Clearly then, $\dot{v}(\tau^-)$ is constant and greater than α. However, $\dot{v}(\tau^+) \leq \alpha$ since the state constraint is now tight. Therefore, $v(t)$ is not smooth at τ which leads to $\zeta_v(\tau) = 0$. The same rationale can be applied to any time τ at which the constraint $\alpha t - v(t) \geq 0$ becomes loose.

These findings, along with the adjoint trajectories in (EC.9), lead to the following result.

LEMMA EC.1. *The adjoints $\lambda_S(t)$ and $\lambda_I(t)$ are continuous while $\lambda_v(t)$ is constant for all $t \in [0,T]$ such that $S(t) > 0$.*

Next, we note that the Lagrangian in (L) is not explicitly a function of t. Then, from (EC.7), we can conclude that $dH/dt = 0$. Combining this finding with (EC.11) and the values of the jump parameters, we obtain the following result.

LEMMA EC.2. *At optimality, the Hamiltonian $H(x(t), u(t), \lambda(t)) = 0$ for all $t \in [0,T]$ such that $S(t) > 0$.*

Using this result along with (H) and (EC.9), we get the following equations.

$$\begin{aligned} \dot{\lambda}_I I &= \omega_1 u_1 + \omega_2 u_2 \\ \dot{\lambda}_S S &= -\gamma I \lambda_I + \omega_1 u_1 + \omega_2 u_2 \end{aligned} \tag{EC.12}$$

Next we prove some monotonicity properties of the adjoint variables.

PROPOSITION EC.1. *$\lambda_I(t)$ is non-decreasing for all $t \in [0,T]$. $\lambda_S(t)$ is increasing for all $t \in [0,T]$ such that $\lambda_S(t) < 0$ and $S(t) > 0$.*

Proof of Proposition EC.1 Consider the linear program in (LP). $(u_1, u_2) = (0,0)$ is always a feasible solution. Given that it is a maximization problem, at optimality, the objective function $\omega_1 u_1 + \omega_2 u_2$ must take a value greater than or equal to 0. Then it follows from (EC.12) that $\dot{\lambda}_I > 0$ for all $t \in [0,T]$.

For λ_S, we use a contradiction argument. Say, for some time $t_s \in [0,T]$, we have $\dot{\lambda}_S(t_s) \leq 0$ and $\lambda_S(t_s) < 0$.

$$\dot{\lambda}_S(t_s) \leq 0 \implies \lambda_I(t_s) \geq 1 + \lambda_S(t_s) \quad \text{[Using (EC.9)]}$$

From (EC.10), we have that $\lambda_S(T) \geq 0$. Given that λ_S is continuous, there must exist $t_p > t_s$ such that $\lambda_S(t_p) \leq \lambda_S(t_s)$ and $\dot{\lambda}_S(t_p) > 0$.

$$\dot{\lambda}_S(t_p) > 0 \implies \lambda_I(t_p) < 1 + \lambda_S(t_p) \leq 1 + \lambda_S(t_s) \leq \lambda_I(t_s)$$

This violates the monotonicity of λ_I. By contradiction, we must have $\dot{\lambda}_S(t) > 0$ for all t such that $\lambda_S(t) < 0$. ■

Singularities are regions where the optimal control becomes indeterminate. Consider the linear program in (LP). Say, we have for some $t \in [0,T]$, $\omega_1(t) = 0$ and $\omega_2(t) < 0$. If we had $\omega_1(t^-) > 0$ and $\omega_1(t^+) < 0$, t would represent a time at which the control switches from one corner solution to another. On the other hand, if we had $\omega_1(t) = 0$ and $\omega_2(t) < 0$ over an interval, the optimal control would be indeterminate in that interval, thus representing a singularity. The following result rules out the possibility of singularities in the optimal control.

PROPOSITION EC.2. *The optimal control does not exhibit singularities.*

Proof of Proposition EC.2 We consider all possible conditions for singularities and rule them out on a case-by-case basis. Say, we have the following singularity condition at time t_s

- $\omega_1 > \omega_2$, $\omega_1 = 0$, $\dot{\omega}_1 = 0$.

$$\begin{aligned} &\omega_2(t_s) < 0 \implies u_2(t_s) = 0 \\ &\dot{\omega}_1(t_s) = 0 \implies \dot{\lambda}_S(t_s) = 0 \implies \lambda_I(t_s) = 0 \quad \text{[Since } \omega_1(t_s) = 0, u_2(t_s) = 0 \text{ and using (EC.12)]} \\ &\qquad \text{and } \lambda_I(t_s) = 1 + \lambda_S(t_s) \implies \lambda_S(t_s) = -1 \quad \text{[Using (EC.9)]} \end{aligned}$$

  Since $\lambda_S(T) \geq 0$, there must exist a time $t_p > t_s$ such that $\lambda_S(t_p) \leq -1$ and $\dot{\lambda}_S(t_p) > 0$. From (EC.9), this implies that $\lambda_I(t_p) < 0$. This violates the monotonicity condition of λ_I.

- $\omega_2 > \omega_1, \omega_2 = 0, \dot{\omega}_2 = 0$.

  This follows a similar line of reasoning as the preceding case.

- $\omega_1 = \omega_2 > 0, \dot{\omega}_1 = \dot{\omega}_2$.

$$\dot{\omega}_1(t_s) = \dot{\omega}_2(t_s) \implies \dot{\lambda}_S(t_s)\epsilon_1 = \dot{\lambda}_S(t_s)\epsilon_2$$

  Either $\epsilon_1 = \epsilon_2$ or $\dot{\lambda}_S(t_s) = 0$.

$$\omega_1(t_s) = \omega_2(t_s) \implies -\lambda_S(t_s)\epsilon_1 + \lambda_v c_1 = -\lambda_S(t_s)\epsilon_2 + \lambda_v c_2 \implies \lambda_S(t_s) = \frac{c_1 - c_2}{\epsilon_1 - \epsilon_2}\lambda_v$$

  Since λ_S is continuous in $[0, T]$, we cannot have $\epsilon_1 = \epsilon_2$. Consider $\dot{\lambda}_S(t_s) = 0$.

$$\omega_1(t_s) > 0 \implies \lambda_S(t_s) < \frac{c_1}{\epsilon_1}\lambda_v \leq 0.$$

  This contradicts Proposition EC.1.

- $\omega_1 = \omega_2 = 0, \dot{\omega}_1 = 0$ or $\dot{\omega}_2 = 0$.

$$\begin{aligned} \omega_1(t_s) = \omega_2(t_s) = 0 &\implies \lambda_S(t_s) = 0 \text{ and } \lambda_v = 0 \\ \dot{\omega}_1 = 0 &\implies \dot{\lambda}_S(t_s) = 0 \implies \lambda_I(t_s) = 1 \quad \text{[Using (EC.9)]} \\ &\text{and } \lambda_I(t_s) = 0 \quad \text{[Using (EC.12)]} \end{aligned}$$

  Clearly, this is a contradiction. A similar line of reasoning follows for the case where $\dot{\omega}_2 = 0$. ■

COROLLARY EC.1. *It follows from Proposition EC.1 and Proposition EC.2 that if, for some $t_s \in [0, T]$, we have $(u_1^*(t_s), u_2^*(t_s)) = (0, 0)$, we must have $(u_1^*(t), u_2^*(t)) = (0, 0)$ for all $t \in (t_s, T]$.*

In words, once vaccination stops under the optimal policy, it does not restart. This is because $\omega_i > 0$ implies $\lambda_S < 0$. Since λ_S is strictly increasing while $\lambda_S < 0$, ω_i is strictly decreasing while $\omega_i > 0$. Furthermore, the absence of singularities guarantees that ω_i does not remain at 0 over an interval.

PROPOSITION EC.3. *If there are more vaccines than can be used to vaccinate susceptibles under the optimal policy ($v(T) < v_{umax}$),* complete vaccination *must be achieved ($S(T) = 0$). Alternatively, if the optimal policy does not lead to a* complete vaccination *scenario ($S(T) > 0$), the stockpile must be depleted ($v(T) = v_{max}$).*

Proof of Proposition EC.3 We prove this by contradiction. Assume that $v(T) < v_{max}$ and $S(T) > 0$. By (EC.10), we have $\lambda_v = 0$ and $\lambda_S(T) = 0$. Recall that we can make T arbitrarily large. This means that vaccination must have stopped at some time before T. If this weren't the case, the piecewise constant controls would have led to either $v(t) = v_{max}$ or $S(t) = 0$ at some time $t < T$.

Let $t_s < T$ be the time at which vaccination stopped. Given that $(u_1^*(t_s), u_2^*(t_s)) = (0, 0)$, we must have $\omega_1(t_s) = 0, \omega_2(t_s) \leq 0$ or $\omega_1(t_s) \leq 0, \omega_2(t_s) = 0$.

$$\begin{aligned} &\omega_1(t_s) = -\lambda_S(t_s)\epsilon_1 \text{ and } \omega_2(t_s) = -\lambda_S(t_s)\epsilon_2 \quad \text{[Since } \lambda_v = 0\text{]} \\ &\implies \lambda_S(t_s) = 0 \text{ and } \dot{\lambda}_S(t_s) > 0 \quad \text{[Using Proposition } EC.2\text{]} \\ &\implies \lambda_I(t_s) < 0 \text{ and } \dot{\lambda}_I = 0 \; \forall \, t \in [t_s, T] \quad \text{[Using (EC.12)]} \\ &\implies \lambda_I(t) < 0 \; \forall \, t \in [0, T] \quad \text{[Using Proposition } EC.1\text{]} \\ &\implies \dot{\lambda}_S(t) > 0 \; \forall \, t \in [0, T] \quad \text{[Using (EC.12)]} \end{aligned}$$

Since $\lambda_S(T) = 0$, we must have $\lambda_S(t) < 0 \; \forall \, t \in [0, T)$. This is a contradiction to $\lambda_S(t_s) = 0$. ■

COROLLARY EC.2. *For all $t \in [0, T]$ such that $v(t) < v_{max}$ and $S(t) > 0$, at least one of the following must hold: $u_1^*(t) > 0$ or $u_2^*(t) > 0$.*

In words, vaccination continues until either the stockpile is depleted or there are no susceptibles remaining to be vaccinated.

B.4. Proof of Main Results

Proof of Proposition 1 We first note that $\lambda_S(t) < 0$ as long as vaccination continues. Now, let us analyze the two cases separately.

- $\mathcal{E} \leq 1 : \frac{c_1}{\epsilon_1} \leq \frac{c_2}{\epsilon_2} \implies \lambda_S c_1 \epsilon_2 \geq \lambda_S c_2 \epsilon_1 \implies c_2(-\lambda_S \epsilon_1 + \lambda_v c_1) \geq c_1(-\lambda_S \epsilon_2 + \lambda_v c_2) \implies \frac{\omega_1}{c_1} \geq \frac{\omega_2}{c_2}$
  Under this condition, the optimal solution to (LP) is $u_1^* = \min\left(u_{max}, \frac{\alpha}{c_1}\right)$ and $u_2^* = 0$. $t_1 = \min\left(\frac{v_{max}}{c_1 u^F}, \tau(\pi^F)\right)$ follows from Corollary EC.2.
- $\mathcal{E} \geq c_1/c_2 : \epsilon_1 \leq \epsilon_2 \implies -\lambda_S \epsilon_1 \leq -\lambda_S \epsilon_2 \implies -\lambda_S \epsilon_1 + \lambda_v c_1 \leq -\lambda_S \epsilon_2 + \lambda_v c_2 \implies \omega_1 \leq \omega_2$
  Under this condition, the optimal solution to (LP) is $u_1^* = 0$ and $u_2^* = \min\left(u_{max}, \frac{\alpha}{c_2}\right)$. $t_2 = \min\left(\frac{v_{max}}{c_2 u^R}, \tau(\pi^R)\right)$ follows from Corollary EC.2. ■

Proof of Theorem 1 We first note that, under $1 < \mathcal{E} < c_1/c_2$, we have the following order.

$$\frac{c_1 - c_2}{\epsilon_1 - \epsilon_2}\lambda_v \leq \frac{c_1}{\epsilon_1}\lambda_v \leq \frac{c_2}{\epsilon_2}\lambda_v \leq 0 \tag{EC.13}$$

When $\lambda_S(t) < \frac{c_1 - c_2}{\epsilon_1 - \epsilon_2}\lambda_v$, we have $\omega_1(t) > \omega_2(t)$. This signifies the early epidemic phase as described in §4.2. Here, the solution to (LP) is as follows.

$$(u_1^*(t), u_2^*(t)) = \begin{cases} (u_{max}, 0), & \text{if } \alpha \geq c_1 u_{max} \\ \left(\frac{\alpha - c_2 u_{max}}{c_1 - c_2}, \frac{c_1 u_{max} - \alpha}{c_1 - c_2}\right), & \text{if } \alpha \in (c_2 u_{max}, c_1 u_{max}) \\ \left(0, \frac{\alpha}{c_2}\right), & \alpha \leq c_2 u_{max} \end{cases}$$

When $\lambda_S(t) = \frac{c_1 - c_2}{\epsilon_1 - \epsilon_2}\lambda_v$, we have $\omega_1(t) = \omega_2(t)$. This represents the switching point t_1 between the early and mid epidemic phase.

For $\frac{c_1 - c_2}{\epsilon_1 - \epsilon_2}\lambda_v < \lambda_S(t) < \frac{c_2}{\epsilon_2}\lambda_v$, we have $\omega_1(t) < \omega_2(t)$ and $\omega_2(t) > 0$. This region is the mid epidemic phase and the optimal solution to (LP) is as follows.

$$(u_1^*(t), u_2^*(t)) = \begin{cases} (0, u_{max}), & \text{if } \alpha \geq c_1 u_{max} \\ (0, u_{max}), & \text{if } \alpha \in (c_2 u_{max}, c_1 u_{max}) \\ \left(0, \frac{\alpha}{c_2}\right), & \text{if } \alpha \leq c_2 u_{max} \end{cases}$$

The feasible polytopes for each case and the corresponding optimal solutions to (LP) are visually represented in Figure EC.1.

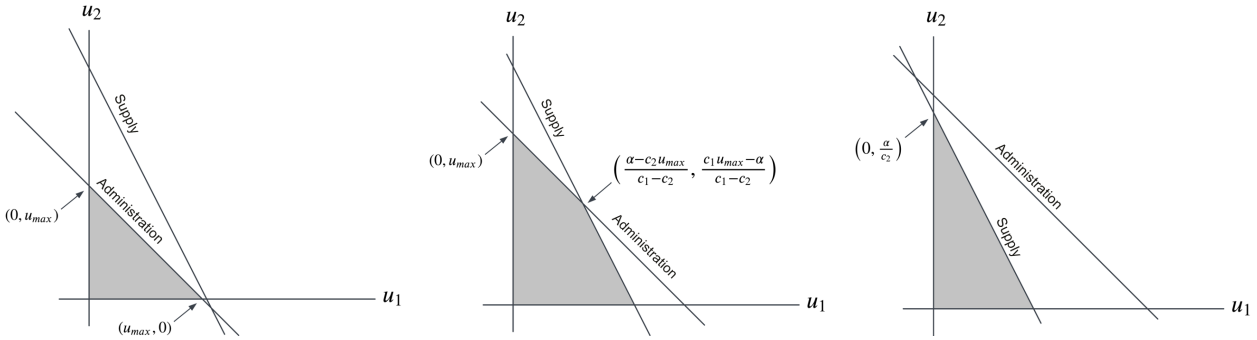

(a) $\alpha \geq c_1 u_{max}$ **(b)** $\alpha \in (c_2 u_{max}, c_1 u_{max})$ **(c)** $\alpha \leq c_2 u_{max}$

Figure EC.1 Visual reprentation of (LP)

Notes. The supply rate constraint is $c_1 u_1 + c_2 u_2 \leq \alpha$ and the administration rate constraint is $u_1 + u_2 \leq u_{max}$.

At $\lambda_S(t) = \frac{c_2}{\epsilon_2}\lambda_v$, we have $\omega_2(t) = 0$, denoted by time t_2 where the optimal control switches off. For $\lambda_S(t) > \frac{c_2}{\epsilon_2}\lambda_v$, we have $\omega_1(t), \omega_2(t) < 0$. The optimal solution in this region is $(u_1^*(t), u_2^*(t)) = (0, 0)$. This represents the end-epidemic phase where no vaccination takes place.

If the vaccine stockpile is depleted by the end of the epidemic ($v(T) = v_{max}$), we have $\lambda_v < 0$, resulting in $\frac{c_1 - c_2}{\epsilon_1 - \epsilon_2}\lambda_v < \frac{c_2}{\epsilon_2}\lambda_v$. This implies $t_1 < t_2$. As the optimal control does not result in *complete vaccination*, we have $\tau(\pi^*) = \infty$. In this case, the relationship between the time-thresholds t_1 and t_2 can be derived using the state equation for $v(t)$ as follows.

$$\int_0^T v(t)dt = v_{max} \implies c_1 \int_0^{t_1} u_1(t)dt + c_2 \int_0^{t_2} u_2(t)dt = v_{max}$$

If, however, the optimal control is such that *complete vaccination* is achieved ($S(T)=0$), we have $\tau(\pi^*)<\infty$. As the vaccine stockpile does not get used up entirely, we have $v(T)<v_{max}$, leading to $\lambda_v=0$. In this scenario, we obtain $\omega_1(t)>\omega_2(t)$ for as long as $\lambda_S(t)<0$ and $\omega_1(t),\omega_2(t)<0$ when $\lambda_S(t)>0$, thus implying that $t_1=t_2$. It follows from Corollary EC.2 that $\lambda_S(t)<0$ as long as $S(t)>0$, that is for all $t\in[0,\tau(\pi^*))$. This results in $t_1=t_2=\tau(\pi^*)$. ■

Proof of Proposition 2 Let us consider each case one-by-one.

- If $v_{max}\to 0$, $\pi^*=\pi^R$.

  Let $v_{max}=\Delta v\to 0$. If we use only the fractional-dose vaccine at the maximum rate, the stockpile will be depleted at time $\Delta t=\Delta v/(c_2u_{max})\to 0$. Any other convex combination of the full- and fractional-dose vaccines implemented at the maximum vaccination rate will deplete the stockpile at some time $t<\Delta t$. Let θ be the proportion of stockpile allocated to the full-dose vaccine. In what follows, we prove that for $\Delta v\to 0$, the optimal policy is to use the only the fractional-dose vaccine ($\theta=0$).

  Assume that we first use the full-dose vaccination at the maximum rate (until time $t_1=\theta\Delta v/(c_1u_{max})$). The remaining stockpile is used to vaccinate using the full-dose vaccine at the maximum rate (until time $t_2=\theta\Delta v/(c_1u_{max})+(1-\theta)\Delta v/(c_2u_{max})$). Once the stockpile is depleted, no vaccination takes place for $t\in[t_2,\Delta t]$. We use first-order approximation to first obtain the state trajectories at time t_1.

$$S(t_1)=S_0-\beta S_0I_0\frac{\theta\Delta v}{c_1u_{max}}-\theta\Delta v\frac{\epsilon_1}{c_1}$$
$$I(t_1)=I_0+\beta S_0I_0\frac{\theta\Delta v}{c_1u_{max}}-\gamma I_0\frac{\theta\Delta v}{c_1u_{max}}$$

  Ignoring the second order terms in Δv, we obtain the state trajectories at time t_2.

$$S(t_2)=S_0-\beta S_0I_0\frac{\Delta v}{u_{max}}\left(\frac{\theta}{c_1}+\frac{1-\theta}{c_2}\right)-\Delta v\left(\theta\frac{\epsilon_1}{c_1}+(1-\theta)\frac{\epsilon_2}{c_2}\right)$$
$$I(t_2)=I_0+\beta S_0I_0\frac{\Delta v}{u_{max}}\left(\frac{\theta}{c_1}+\frac{1-\theta}{c_2}\right)-\gamma I_0\frac{\Delta v}{u_{max}}\left(\frac{\theta}{c_1}+\frac{1-\theta}{c_2}\right)$$

  Next, we compute the state trajectories at time Δt following a similar method.

$$S(\Delta t)=S_0-\beta S_0I_0\Delta t-\Delta v\left(\theta\frac{\epsilon_1}{c_1}+(1-\theta)\frac{\epsilon_2}{c_2}\right)$$
$$I(\Delta t)=I_0+\beta S_0I_0\Delta t-\gamma I_0\Delta t$$

  Interestingly, for an infinitesimally small vaccine stockpile, the trajectory of I does not depend on the allocation θ. Since, $\epsilon_1/c_1<\epsilon_2/c_2$, $S(\Delta t)$ is minimized at $\theta=0$. As the vaccine stockpile is depleted by time Δt, the epidemic proceeds without vaccination for $t\in[\Delta t,T]$. The initial conditions for the epidemic in this duration are $S(\Delta t),I(\Delta t)$. It is rather straightforward to show that, for an epidemic with identical initial I, a lower initial S leads to a lower value for the total number of infections. Initially, we assumed that the full-dose vaccine is used first and the fractional-dose vaccine next. Altering the sequence of vaccination does not change the values of $S(\Delta t)$ and $I(\Delta t)$. Then, it follows that the optimal policy for an infinitesimally small vaccine stockpile ($v_{max}\to 0$) is to use only the fractional-dose vaccine ($\pi^*=\pi^R$) for $t\in[0,\Delta t)$.

- If $v_{max}\geq\frac{S_0c_1}{\epsilon_1}$, $\pi^*=\pi^F$.

  In this case, the vaccine stockpile is sufficient to effectively immunise all initial susceptibles S_0 using the full-dose vaccine. As the epidemic progresses, some susceptibles get infected and move to the recovered compartment while others get vaccinated. This means that, even at a maximum vaccination rate, any feasible vaccination policy will be able to effectively immunise fewer than S_0 individuals. As a result, we will have $v(T)<S_0c_1/\epsilon_1\implies v(T)<v_{max}\implies\lambda_v=0$. Additionally, *complete vaccination* will be achieved since vaccination will continue until $S(t)$ reaches 0. We have seen in the Proof of Theorem 1, that when $\lambda_v=0$, we have $\omega_1(t)>\omega_2(t)$ for as long as $\lambda_S(t)<0$ and $\omega_1(t),\omega_2(t)<0$ when $\lambda_S(t)>0$. This means that the optimal solution to (LP) will be $(u_1^*(t),u_2^*(t))=(u_{max},0)=\pi^F$ for all $t\in[0,\tau(\pi^*))$.

- If $u_{max} \to 0$, $\pi^* = \pi^F$.

  With $u_{max} \to 0$, the minimum time required to deplete the vaccine stockpile is $v_{max}/(c_1 u_{max}) \to \infty$. Since T is finite, we must then have $v(T) < v_{max}$. Using the same argument as in the preceding case, we obtain the optimal policy as π^F for all $t \in [0, T)$.
- If $u_{max} \to \infty$, $\pi^* = \pi^R$.

  An infinite vaccination rate signifies instantaneous vaccination. Effectively, the initial state S_0 is reduced by the number of effectively immunised individuals. Assume that a proportion θ of the stockpile is allocated to the full-dose vaccine. The initial state of the epidemic can be computed as $S(0^+) = S_0 - \theta v_{max} \frac{\epsilon_1}{c_1} - (1-\theta) v_{max} \frac{\epsilon_2}{c_2}$. Clearly, $S(0^+)$ is minimized for $\theta = 0$. Following a similar line of reasoning as in the case of $v_{max} \to 0$, we obtain the optimal policy as $\pi^* = \pi^R$ for $t \in [0, v_{max}/(c_2 u_{max}))$. ■

Proof of Proposition 3 When $\alpha \leq c_1 u_{max} \frac{\epsilon_2}{\epsilon_1}$, the maximum effective immunisation rate using the fractional-dose vaccine is $\epsilon_2 u_{max}$. Note that the full-dose vaccine cannot be administered at rate u_{max} as the supply is constrained at $\alpha < c_1 u_{max}$. The maximum effective immunisation rate using the full-dose vaccine can be computed as $\epsilon_1 \alpha / c_1 \leq \epsilon_2 u_{max}$. As a result, the fractional-dose 1D policy is optimal for any stockpile size.

Appendix C: Numerical Study - Supplement

C.1. Vaccine Efficacy Estimation

Vaccine efficacy is typically computed as the post-vaccination relative risk of infection between vaccinated and control groups during a clinical efficacy trial (Halloran et al. 1997). Let $Y \in \{0, 1\}$ be a Bernoulli indicator random variable associated with the event of an infection. The vaccine efficacy can be expressed as

$$\epsilon = 1 - \frac{\mathbb{E}[Y_1]}{\mathbb{E}[Y_0]}, \tag{EC.14}$$

where Y_1 and Y_0 denote infection events in the vaccinated and control groups, respectively. For the full-dose vaccine, the efficacy can be estimated from a sample average approximation of the realised infection events in each group. This efficacy estimation method does not apply to fractional-dose vaccines however, because for them only immunogenicity (as opposed to infection outcomes) data are available.

The immune response induced by a vaccine is typically measured through a serological assay (example - hemagglutination inhibition assay for influenza, enzyme-linked immunosorbent assay for Hepatitis A/B, microneutralisation assay for COVID-19), producing a response metric value denoted x (Plotkin 2010). Following Dunning (2006), we use an inverse-logit function to model the relationship between assay values and protection against infection. Let $\rho(x)$ denote the *clinical protection function*, or probability that an individual with assay value x is protected against the disease conditional on exposure. We thus model the clinical protection function as

$$\rho(x) = \frac{1}{1 + \exp(a + bx)}, \tag{EC.15}$$

where a and b are parameters specific to the disease.

Let λ denote the baseline probability of infection, defined as the probability of exposure to the disease irrespective of vaccination status over the entire period in which the subject is observed. λ may depend on the prevalence of infection in the population, contact rate of individuals, infectiousness of the disease and other factors independent of the immunogenic response. The probability that an individual with assay x develops an infection is thus given by

$$\mathbb{P}(Y = 1|x) = \lambda(1 - \rho(x)), \tag{EC.16}$$

and the probability of observing an infection outcome y conditional on x is

$$\mathbb{P}(Y = y|x) = \left(\frac{\lambda \exp(a + bx)}{1 + \exp(a + bx)}\right)^y \left(1 - \frac{\lambda \exp(a + bx)}{1 + \exp(a + bx)}\right)^{(1-y)}. \tag{EC.17}$$

Consider a vaccine for which an efficacy trial has been carried out. We can use the immunogenicity and infection outcome data of the full-dose vaccine to estimate the parameters (λ, a and b) of the above Clinical Protection (CP) model. Examples of studies that have carried out this exercise for different diseases are Storsaeter et al. (1998) for pertussis, Coudeville et al. (2010) for seasonal influenza and Khoury et al. (2021) for COVID-19. These studies also typically report the 95% credible intervals of each parameter estimate, representing a source of uncertainty in the efficacy estimation method.

Consider now a fractional dose of the same vaccine for which an immunogenicity trial was carried out but a vaccine efficacy trial was not. We use a Bayesian framework to estimate its efficacy. Specifically let $\mathbf{x_{i,i\in\{0,1\}}}$ denote the observed assay values in the immunogenicity trial, where 0 and 1 denote unvaccinated and vaccinated groups respectively. Let $\theta_{i,i\in\{0,1\}}$ denote the distributions of the corresponding population parameters. It is established in the clinical literature that observed assay values are approximately normally distributed (Nauta et al. 2009). We assume a Normal-Inverse-Chi-Square ($NI\text{-}\chi^2$) prior for the population parameters θ_i. Using the observed assay values, we obtain the posterior distribution for θ_i, which is also $NI\text{-}\chi^2$ distributed:

$$\mathbb{P}(\theta_i|\mathbf{x_i}) \propto \mathbb{P}(\mathbf{x_i}|\theta_i) \cdot \mathbb{P}(\theta_i) \tag{EC.18}$$

Clearly, larger clinical trials which report larger samples of x_i have the potential to produce more reliable estimates of θ_i.

Next, we draw the parameters of the CP model (λ, a and b) from their respective credible intervals. We also sample $\theta_{i,i\in\{0,1\}}$ from the posterior and compute the expected infection risk for both vaccinated and unvaccinated groups using the sampled parameters. This can be represented as $\mathbb{E}_{x\sim N(\theta_i)}[\mathbb{P}(Y=1|x;\lambda,a,b)]$, where the expectation is over assay values normally distributed with parameter θ_i. The efficacy corresponding to the sampled parameters can be estimated as the inverse of the relative infection risk between the vaccinated and unvaccinated groups, or

$$\hat{\epsilon} = 1 - \frac{\mathbb{E}_{x\sim N(\theta_1)}[\mathbb{P}(Y=1|x)]}{\mathbb{E}_{x\sim N(\theta_0)}[\mathbb{P}(Y=1|x)]} \tag{EC.19}$$

Resampling the parameters λ, a, b, and $\theta_{i,i\in\{0,1\}}$ from the respective distributions enables us to compute the distribution of $\hat{\epsilon}$. We truncate the resulting distribution at the efficacy of the full-dose vaccine because the fractional-dose vaccine can be expected to be less efficacious than the full-dose in realistic scenarios. We use the associated 95% credible intervals to evaluate the robustness of vaccination policies with respect to the estimated fractional dose vaccine efficacy in the following numerical analyses.

C.2. Data Estimation for Case Studies

C.2.1. Case Study 1: 2004-05 Seasonal Influenza in the United States The TIV typically comprises three different strains of the influenza virus with 15µg of active component per strain, which corresponds to the full-dose vaccine in our model. Engler et al. (2008) evaluated a half-dose TIV (7.5µg of active component per strain) and found that immune responses induced "were not substantially inferior to the full-dose vaccine"; we adopt this as the fractional-dose vaccine. The study reports the immunogenicity of the TIV against each of the three influenza strains separately. We use the immunogenicity against Influenza A (H3N2) in the age group 18-49 years to estimate the vaccine efficacy, as this was the dominant strain circulating in the United States in that influenza season (CDC 2005).

Next, we draw on the findings in Coudeville et al. (2010) for the parameters of the clinical protection model (we use the parameters of the ALL model, which includes all datasets and no covariates) corresponding to seasonal influenza: $a = 2.844$ (credible interval: $2.25 - 3.36$), $b = 1.299$ (credible interval: $1 - 1.69$). We then use the methodology outlined in Appendix C.1 with the immunogenicity data of the TIV and the parameter estimates described above to estimate the vaccine efficacies. This yields $\epsilon_1 = 0.434$ (credible interval: $0.191 - 0.683$) and $\epsilon_2 = 0.350$ (credible interval: $0.145 - 0.594$). For the purpose of this analysis, we treat $\epsilon_1 = 0.434$ as a known deterministic value. This assumption is based on the premise that, even though we estimate it here using immunogenicity data, ϵ_1 is typically ascertained with substantial accuracy through extensive efficacy trials.

This approach, however, presents a challenge to accurately quantifying the uncertainty surrounding ϵ_2: given the interdependence of ϵ_1 and ϵ_2 through the parameters of the CP model, it is inappropriate to directly apply the raw credible intervals of ϵ_2 while only

considering the mean value of ϵ_1. To address this issue, we numerically compute the 95% bounds for the ratio ϵ_2/ϵ_1. Multiplying these bounds by the mean value of ϵ_1 provides us with a refined set of bounds for ϵ_2, specifically $0.274 - 0.429$. Note that we have $1 < \mathcal{E} < c_1/c_2$ for all values of ϵ_2 within its 95% bounds.

The total stockpile available for vaccination in the US at the onset of the 2004-05 flu season was 65 million (full) doses (Seiguer 2005). With a population of 330 million and a single-dose regimen of the TIV, this results in the estimate $v_{max} = 0.22$. We also obtained data on weekly doses of vaccine administered in the US for influenza seasons 2018-19 through 2021-22 and observe that the maximum weekly vaccination rate was approximately 18 million doses in each season (CDC 2022). We use this as a proxy for the maximum vaccine administration capacity, resulting in the estimate $u_{max} = 0.008$. We have been unable to obtain reliable data on the supply rate of influenza vaccines. Considering alternative contextually relevant data, we note that a total of 11.2 billion doses of COVID-19 vaccines were distributed globally in 2021 (World Economic Forum 2022), resulting in a supply rate of $\alpha = 0.0019$. Given that the production and distribution of influenza vaccines starts well ahead of the flu season, we anticipate a significantly higher supply rate for influenza vaccines, and accordingly set our assumed baseline supply rate for influenza vaccines at $\alpha = 0.005$.

In the US, the influenza vaccination campaign typically commences around week 30 of each calendar year. By analyzing data on reported "influenza-like cases" during week 30 across a decade (2010-2019), we estimated an average initial prevalence $I_0 = 0.00007$. Given the early stage of the flu season in week 30, it is reasonable to assume that the recovered population is negligible at this point, so that $S_0 = 1 - I_0$ and the initial epidemic conditions were set to $S_0 = 0.99993$ and $I_0 = 0.00007$. We used the reproduction number $\mathcal{R}_0 = 1.28$ of seasonal influenza reported in the literature review Biggerstaff et al. (2014). Additionally, we used the average infectious period for influenza of 5 days reported in Suess et al. (2012), resulting in estimated rate of recovery $\gamma = 0.2$ and transmission rate $\beta = 0.256$.

For the SEIR model that we discuss in §C.3.3, the additional epidemic parameters were drawn on Hill et al. (2019) as follows: incubation period of 1.4 days and recovery period of 3.8 days. For the heterogeneous population model discussed in §C.3.4, we estimated the average number of contacts per day among population groups from the contact matrix in Mossong et al. (2008) as follows: 5.18 between adults, 3.02 between the elderly and 2.28 across individuals in the two groups. The proportion of adults was estimated at 0.52 based on the demographic distribution in high-income countries presented in Więcek et al. (2021). The probability of transmission on contact was computed at 0.0801 to match the desired $\mathcal{R}_0$.

C.2.2. Case Study 2: Nigeria's Response to COVID in 2021-22 During the phase 1 trial for the Pfizer-BioNTech COVID-19 vaccine, three dosages of two vaccine candidates, BNT162b1 and BNT162b2, were examined for safety and immunogenicity, specifically at 10, 20, and 30 µg dosage (Walsh et al. 2020). Finally, the 30µg dosage of BNT162b2 was advanced for the efficacy trial. This forms the basis for the full-dose vaccine in our model with $c_1 = 1$ and $\epsilon_1 = 0.95$ (Polack et al. 2020). For the fractional-dose vaccine, we consider the 10 µg dosage of BNT162b2, which implies $c_2 = 1/3$. We obtain its immunogenicity data from Walsh et al. (2020), which reports the phase 1 trial results. For the clinical protection model, we use the parameters estimated in Khoury et al. (2021) along with their 95% credible intervals. The estimated mean efficacy of the fractional-dose vaccine is $\epsilon_2 = 0.817$ (95% credible interval: 0.677-0.923). Note that the relative efficacy condition $1 < \mathcal{E} < c_1/c_2$ is satisfied for the entire range of ϵ_2.

While operational challenges affecting COVID-19 vaccinations were ubiquitous, these were particularly pronounced in developing countries like Nigeria where the acquisition of vaccine stockpiles largely depended on donations from other nations or through international non-profit initiatives like COVAX. To date, Nigeria has administered approximately 133 million doses of COVID-19 vaccines (Our World in Data 2024a). Considering a population of 213 million and that a full vaccination course entails two doses, this translates to a stockpile size of $v_{max} = 0.312$. Obtaining precise data on a country's maximum vaccination capacity is challenging. For our model, we use the peak daily per capita vaccination rate observed during the campaign as a proxy for this capacity, yielding $u_{max} = 0.00115$ (Our World in Data 2024a). Nigeria received their first shipment of COVID-19 vaccines on 02 March, 2021. By the same time next year, they had received a total of approximately 69 million doses (UNICEF 2024). Data on subsequent

shipments were hard to obtain. Therefore, we consider the information until this time to compute the vaccine supply rate. Assuming linear stockpile growth between 02 March, 2021 and 02 March, 2022, we obtain $\alpha = 0.000434$.

As vaccination started in the first week of March, 2021, the initial conditions of the state variables are based on the state of the epidemic in Nigeria at this juncture: $S_0 = 0.9926$ and $I_0 = 0.00004$ (Our World in Data 2024b). The basic reproduction number ($\mathcal{R}_0$) for COVID-19, drawn from the meta-analysis Ahammed et al. (2021), is estimated at $\mathcal{R}_0 = 2.69$. The literature presents a range of estimates for the average infectious period of COVID-19, typically between 6 and 14 days (Byrne et al. 2020, Arino and Portet 2020). We assume here an intermediate infectious period of 10 days, leading to an estimated recovery rate (γ) of 0.1 and a corresponding transmission rate (β) of 0.269.

C.3. Additional Results from the Case Studies

C.3.1. Case Study 2: Nigerias Response to COVID in 2021-22 In §5.2, we presented the optimal policy map for different combinations of vaccine stockpile size (v_{max}) and maximum administration rate (u_{max}) along with the improvement in infections averted that the optimal policy achieves over the full-dose 1D policy. Figure EC.2a illustrates the corresponding share of fractional-dose vaccine as specified by the optimal policy. We see that a *Mix* policy with an allocation of 82% of the available stockpile to the fractional-dose vaccine (see §4.2) would have been optimal and preferable to the full-dose 1D policy used historically in Nigeria.

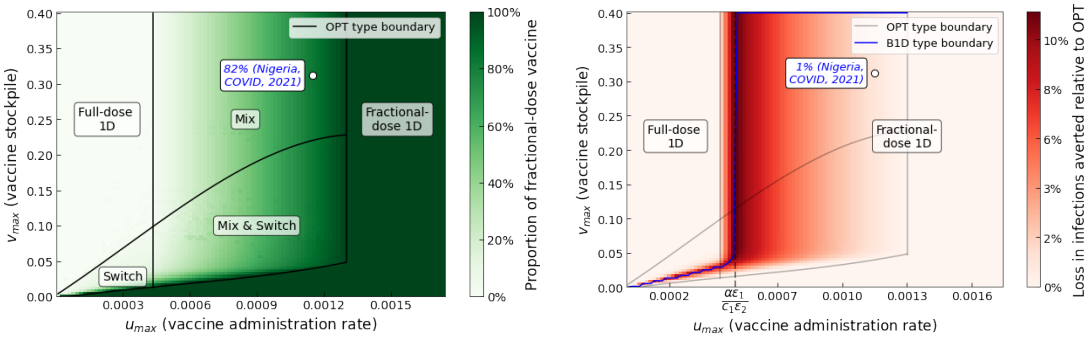

(a) Optimal policy (OPT) map and the corresponding share of fractional-dose vaccine

(b) Best 1D policy (B1D) map and the loss in infections averted relative to the optimal policy (OPT)

Figure EC.2 Optimal policy map and performance of the Best 1D policy (B1D)

Figure EC.2b plots the policy map for the best 1D policy. As in the preceding case study, we note that the vertical part of the boundary separating the parameter space into the two 1D policies (π^F and π^R) is almost exactly aligned with the value of the bound computed in Proposition 3. Figure EC.2b also suggests that, using the best 1D policy would be appropriate in the Nigerian context as the loss in infections averted with respect to the optimal policy would be only 1%. However, in some regions the best 1D policy can be suboptimal by as much as 12%, and therefore requires careful consideration before implementing it.

Figure EC.3a shows that the optimal vaccination policy would still be of a Mix type and would still bring about a substantial decrease of infections ($\delta(\pi^*) = 96\%$) relative to the full-dose 1D policy if the actual fractional-dose efficacy ϵ_2 assumed the 2.5^{th} percentile of its estimate distribution (0.677) instead of its mean (0.817). The optimal vaccination policy seems even more robust to the uncertainty in ϵ_2 as compared to the previous case study, as it outperforms the full-dose 1D policy in the worst-case efficacy scenario in the entire parameter space (u_{max}, v_{max}) considered.

Figures EC.2b and EC.3b suggest that the fractional-dose 1D policy's performance would be robust to both moderate changes around the estimated operational constraint parameters (u_{max}, v_{max}) and to the uncertainty affecting the fractional-dose efficacy ϵ_2. Beyond this particular case study however (where a recommendation to use the fractional-dose only policy would seem justified), Figure EC.3b indicates that, in contrast with the optimal policy, the best 1D policy is not always robust to the uncertainty on ϵ_2.

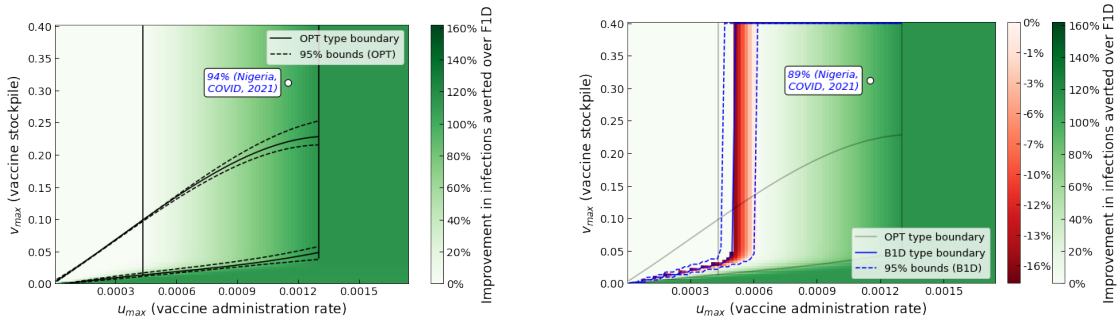

(a) 95% bounds of the optimal policy (OPT) map and improvement in infections averted by the baseline OPT policy over the full-dose 1D policy (F1D)

(b) 95% bounds of the best 1D policy (B1D) map and improvement (loss) in infections averted by the baseline B1D policy over F1D policy

Figure EC.3 Performance of optimal and best 1D policies under uncertain fractional-dose vaccine efficacy

Indeed, there exists values of the operational constraints (u_{max}, v_{max}) where that policy would underperform the conventional full-dose 1D policy by as much as 18% under the worst-case realization of ϵ_2 considered.

Figure EC.4 illustrates the sensitivity of the optimal policy to the vaccine supply rate, α, and the basic reproduction number, $\mathcal{R}_0$. Similar to the previous case study, a decreasing $\mathcal{R}_0$ supports increased use of fractional-dose vaccines, indicated by a shift in the optimal policy from *Mix* to *Mix & Switch*. In contrast to the US influenza scenario, the parameter conditions in Nigeria render the optimal policy highly sensitive to α: a small reduction in α from its estimated value would necessitate using only fractional-dose vaccines under the optimal policy.

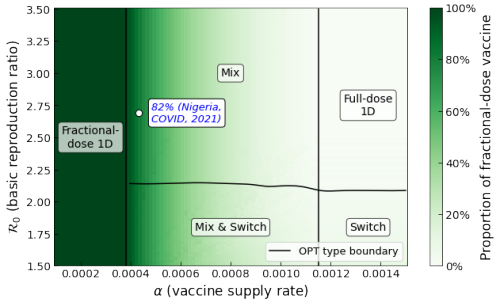

Figure EC.4 Optimal policy (OPT) map (against α and $\mathcal{R}_0$) and the corresponding share of fractional-dose vaccine

C.3.2. SEIRDH Model for the COVID-19 Analysis We consider an SEIRDH model with vaccination (Figure EC.5b) tailored to the Nigerian COVID-19 context. The SEIRDH model is an extension of the SEIR model, which has been widely used to describe COVID-19 dynamics. It accounts for three key structural features that were not captured by the base SIR model:

1. By introducing a compartment for exposed (E) individuals, it allows for an incubation period of the infection.
2. It introduces new compartments for vaccinated individuals - protected (or effectively immunised) (P) and not protected (N). This relaxes the unrealistic assumption of the SIR model in which ineffectively vaccinated individuals remain in the S compartment and have the opportunity to get repeatedly vaccinated (subject to vaccine stockpile availability) until they achieve protection.
3. It includes compartments for hospitalized (H) and deceased (D) individuals, which helps us understand the impact of the vaccination policies on epidemic outcomes other than infections.

In this framework, both susceptible (S) and recovered (R) individuals have the opportunity to get vaccinated but, unlike susceptibles, all recovered individuals are effectively immunised upon vaccination. We also assume that vaccines are administered to susceptible

and recovered individuals in proportion to the sizes of their respective compartments. The epidemic and vaccination dynamics are governed by the following set of differential equations.

$$\begin{aligned}
&\dot{S} = -\beta SI - (u_1+u_2)\tilde{S}, \quad \dot{E} = \beta(S+N)I - \sigma E, \quad \dot{I} = \sigma E - \gamma I \\
&\dot{R} = (1-p_D-p_H)\gamma I + (1-p_{HD})\eta H - (u_1+u_2)\tilde{R} \\
&\dot{P} = (\epsilon_1 u_1 + \epsilon_2 u_2)\tilde{S} - (u_1+u_2)\tilde{R}, \quad \dot{N} = ((1-\epsilon_1)u_1 + (1-\epsilon_2)u_2)\tilde{S} - \beta NI \\
&\dot{D} = p_D\gamma I + p_{HD}\eta H, \quad \dot{H} = p_H\gamma I - \eta H, \quad \dot{v} = c_1u_1 + c_2u_2
\end{aligned} \tag{EC.20}$$

where $\tilde{S} = S/(S+R)$ and $\tilde{R} = R/(S+R)$. The additional epidemic parameters are the reciprocals of the average incubation period ($\sigma = 1/5$, Więcek et al. (2021)) and the average hospitalization period ($\eta = 1/17.3$, Gauchon et al. (2021)); the proportion of infections that result in hospitalizations ($p_H = 0.08$, same) and deaths without hospitalization ($p_D = 0.002405$, same); and the proportion of hospitalizations that result in deaths ($p_{HD} = 0.0315$, same).

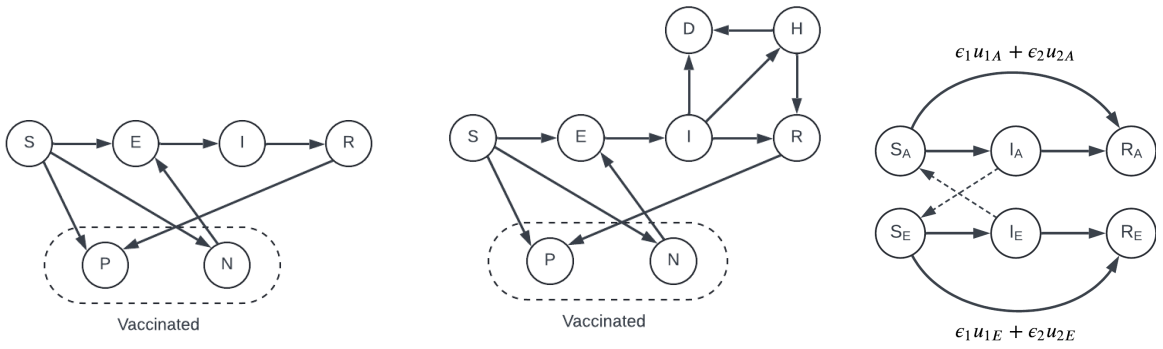

(a) SEIR model with vaccination **(b)** SEIRDH model with vaccination **(c)** Age-structured SIR Model

Figure EC.5 Model extensions

The total number of hospitalizations can be expressed as

$$\begin{aligned}
\int_0^T p_H\gamma I dt &= p_H\int_0^T \sigma E dt - p_H\int_0^T \dot{I} dt = p_H\int_0^T \beta(S+N)Idt - p_H\int_0^T \dot{E} dt - p_H(I_{min}-I_0) \\
&= p_H\int_0^T \beta(S+N)Idt - p_H(E(T)-E(0)) - p_H(I_{min}-I_0)
\end{aligned} \tag{EC.21}$$

The first term, $\int_0^T \beta(S+N)Idt$, is simply the total number of infections; $E(0)$, I_{min} and I_0 are constant parameters; and $E(T)$ (number of exposed individuals at the end of the epidemic) can be expected to take a very small value. This implies that minimizing hospitalizations is near-equivalent to minimizing infections. The total number of deaths can be expressed as

$$\begin{aligned}
D(T) &= \int_0^T p_D\gamma I dt + \int_0^T p_{HD}\eta H dt = p_D\int_0^T \gamma I dt + p_{HD}\int_0^T p_H\gamma I dt - p_{HD}\int_0^T \dot{H} dt \\
&= (p_D + p_{HD}p_H)\int_0^T \gamma I dt - p_{HD}H(T)
\end{aligned} \tag{EC.22}$$

We argued in (EC.21) that the first term, $\int_0^T \gamma I dt$, is approximately equal to the total number of infections. $H(T)$ (number of hospitalized individuals at the end of the epidemic) can be expected to take a very small value. Thus, we have that minimizing deaths is near-equivalent to minimizing infections.

Let us first introduce some terminology to simplify the presentation of the results. We denote the optimal policy and the best 1D (single-dosage) policy derived using the SIR model as the SIR-OPT and SIR-B1D policies respectively. For the SEIRDH model, we obtain the vaccination policy that minimizes infections within the class of policies that follow the specific policy structure laid out in Theorem 1. We refer to this policy as the SEIRDH-OPT policy. It is important to note that this policy may not be the optimal vaccination policy for the more sophisticated model. It is rather a heuristic policy that implements the structure laid out in Theorem

1 and only optimizes for the switching point, t_1, using the SEIRDH dynamics. The SEIRDH-B1D policy is very similar to the SIR-B1D policy: it is the vaccination policy that is restricted to use only one vaccine dosage and that minimizes infections given the dynamics of the corresponding SEIRDH model.

First, we aim to understand the sensitivity of our proposed vaccination policies to the epidemic modelling choice. Figure EC.6a maps the optimal and best 1D policy types obtained using each model (as described above). It also illustrates the absolute difference between the SIR-OPT and SEIRDH-OPT policies in the proportion of fractional-dose vaccine used for different (u_{max}, v_{max}) combinations. We find that the boundaries of the policy types are very close across the two models, indicating that the optimal or best 1D policy types are the same in most of the parameter space. Furthermore, the difference in the fractional-dose quantity between the two optimal policies is low (< 5 percentage points) in most of the parameter space. These findings suggests that the vaccination policies derived on the SIR model (SIR-OPT and SIR-B1D policies) may provide reasonable heuristics for epidemics governed by more complex dynamics.

A natural follow-up question is: how does the "true" performance of these SIR-based policies, when applied in real-world situations, diverge from that predicted by the SIR model? To answer this question, we take the SIR-OPT policy and compare its impact on epidemic outcomes under both SIR and SEIRDH dynamics. Figure EC.6b illustrates the absolute difference in improvement (in terms of infections averted) that the SIR-OPT policy achieves over the full-dose 1D policy across the two models. We find that this difference is around 3 percentage points for the specific Nigerian context, indicating that the optimal policy derived using the simple SIR model would still result in a 128% improvement in epidemic outcomes (instead of 131% as predicted by the SIR model) under the more complex dynamics of the SEIRDH model. The difference in improvement is also less than 6 percentage points for all of the studied (u_{max}, v_{max}) combinations. These results thus suggest that our theoretical findings about vaccination policies and their impact on epidemic outcomes may be robust to epidemic modeling choices.

Consideration of the SEIRDH model also allows us to estimate the impact of fractional-dose vaccines on epidemic outcomes beyond infections - specifically, hospitalizations and deaths. Our model predicts that the traditional full-dose 1D vaccination policy that was implemented saved approximately 650,000 hospitalizations and 40,000 deaths compared to a scenario with no vaccines. The same model predicts that the optimal *Mix* vaccination policy could have averted approximately 1.5 million hospitalizations and 95,000 deaths if implemented.

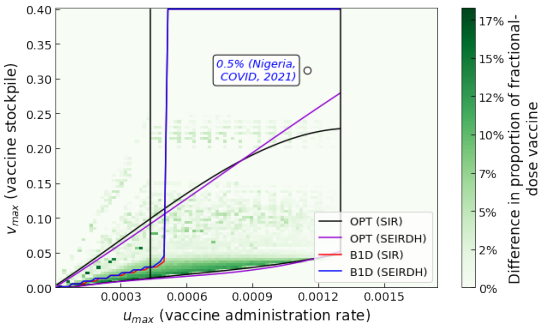

(a) Comparison of the optimal (OPT) and best 1D (B1D) policy map under the SIR and SEIRDH models

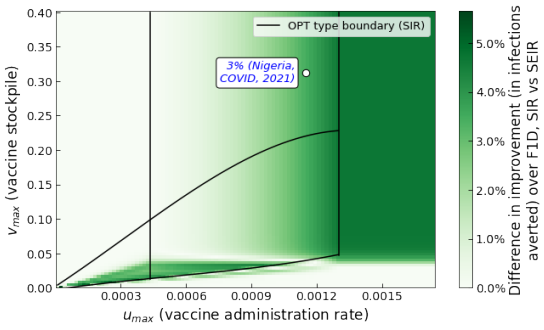

(b) Comparison of the improvement in infections averted by the optimal policy (OPT) over the F1D policy under the SIR and SEIRDH models

Figure EC.6 Robustness of the proposed vaccination policies and their performances against more complex epidemic models

C.3.3. SEIR Model for the Seasonal Influenza Analysis To evaluate the robustness of our findings against alternative epidemic dynamics, we consider an SEIR model with vaccination (Figure EC.5a) which allows for an incubation period of the infection. It also removes vaccinated individuals from the susceptibles compartment thereby preventing ineffectively vaccinated individuals from receiving additional vaccine doses. Our findings here are in line with what we observed in the Nigerian COVID context (see §C.3.2): the optimal policy maps are very close across the two epidemic models (Figure EC.7a) except for a substantial region in the parameter space where the Mix and Switch policy is optimal under the SIR model. This divergence is an artifact of the different vaccination dynamics across the two models. Because susceptibles can be repeatedly vaccinated in the SIR model, it may require a vaccine stockpile $v_{max} > 1$ to achieve complete vaccination. On the other hand, a stockpile of v_{max} close to unity would lead to complete vaccination under the SEIR model since individuals now have a much smaller chance of getting a repeat vaccination (an ineffectively vaccinated individual may still get an additional vaccine regimen if they got infected and moved to the R compartment, following which they would be effectively vaccinated with probability 1). Hence, policies which incorporate a *Switch*, which can be optimal only when complete vaccination is not achievable, are observable primarily for stockpiles of $v_{max} < 1$. Despite the difference in the optimal policy type, we note that the difference in the proportion of fractional-dose vaccine under the optimal policy is less than 3 percentage points in most of the parameter space studied.

In Figure EC.7b, we see that the performance (in terms of infections averted) of the optimal policy as predicted by the SIR model diverges from its performance under the SEIR model by less than 3 percentage points for all (u_{max}, v_{max}) combinations considered. More importantly, in the specific US context, the optimal policy would remain the same whether computed using the SIR model or the SEIR model. If this policy were to be applied when the epidemic follows the SEIR dynamics, the improvement of 64% over the traditional full-dose 1D policy that was predicted under the SIR model would continue to hold.

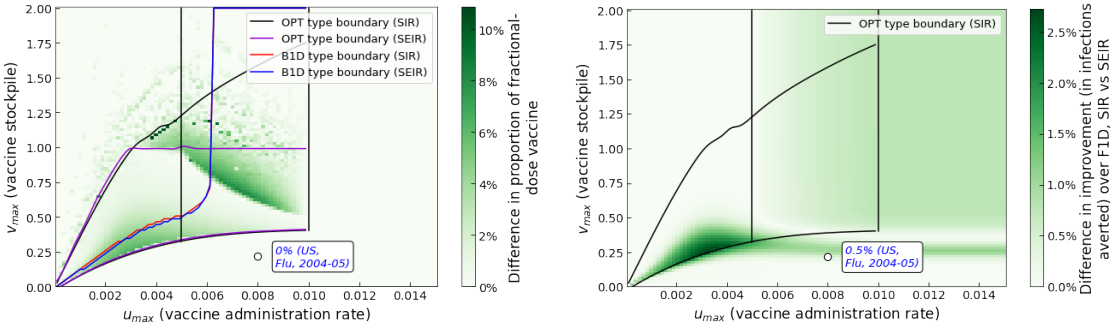

(a) Comparison of the optimal (OPT) and best 1D (B1D) policy map under the SIR and SEIR models

(b) Comparison of the improvement in infections averted by the optimal policy (OPT) over the F1D policy under the SIR and SEIR models

Figure EC.7 Robustness of the proposed vaccination policies and their performances against more complex epidemic models

C.3.4. Age-structured SIR Models Our main analysis did not account for differences in vulnerability and contact rates across different population groups, which may raise concerns about the effectiveness of fractional-dose vaccination policies in environments where the population is heterogeneous. To address this, we conduct a numerical investigation using an age-structured SIR model with two subgroups - young adults aged 20-49 (A) and the elderly aged 50+ (E) - drawn on the epidemic model in Więcek et al. (2021). Notably, our model is considerably more parsimonious than their 9-subgroup SEIR-type model, as our goal is to illustrate the effectiveness of fractional-dose vaccines when the homogeneity assumption of the basic SIR model is relaxed. In contrast, theirs is a full-scale simulation study aimed at generating reliable, COVID-19-specific policy prescriptions. We incorporate varying contact rates between subpopulations, estimated from the contact matrix in Mossong et al. (2008), and draw on Więcek

et al. (2021) for demographic distributions in high-income (for the US influenza case) and low-income (for the Nigeria COVID-19 case) countries. Consistent with Więcek's approach, we assume identical vaccine efficacies and recovery rates across population groups.

We estimated the average number of contacts per day among population groups from the contact matrix in Mossong et al. (2008) as follows: 5.18 between adults, 3.02 between the elderly and 2.28 across individuals in the two groups. The proportion of adults was estimated at 0.78 based on the demographic distribution in low-income countries presented in Więcek et al. (2021). The probability of transmission on contact was computed at 0.0625 to match the desired $\mathcal{R}_0$.

Determining the optimal vaccination policy is generally intractable in compartmental models that account for population heterogeneity (Brauer 2017). Without analytical insights into the qualitative nature of the optimal policy, numerical optimization becomes challenging due to the high dimensionality of the decision space. Specifically, one would need to determine the optimal rates of vaccination with the full and fractional-dose vaccines ($u_1(t)$ and $u_2(t)$, respectively) for each subgroup at every time t in the epidemic. Nevertheless, our findings for a homogeneous population, particularly the concept of the marginal value of vaccination (see §4.2), allow for the development of a reasonable class of heuristic policies, as outlined below.

Group A (young adults) has a higher contact rate, resulting in greater transmissibility and a correspondingly high marginal value of vaccination early in the epidemic. This suggests that the early-vaccination regime (full-dose, Mix, or fractional-dose) should initially focus on group A. As the size of group A decreases, the marginal value of vaccinating individuals in group E (elderly) may surpass that of group A, indicating that a reasonable heuristic policy would switch to vaccinating group E at some point, say, t_1'. Once the susceptible populations in both groups are sufficiently diminished, the policy might transition to the dose-sparing mid-vaccination regime at time t_2', using fractional-dose vaccines on group A, followed by a switch to group E at time t_3'. Vaccination continues until either the stockpile is depleted or complete vaccination is achieved. The intuition of vaccinating group A before group E aligns with recent research (example: Bubar et al. (2021)).

Even with the aforementioned assumptions on the policy structure, determining the switching points that minimize infections remains numerically challenging. In the homogeneous population scenario, specifying the optimal policy involved a simple line search over the switching point t_1. In contrast, the heterogeneous population model with the three switching points requires searching in a 3-dimensional space, which significantly increases computational complexity. Given the specific scope of our study, we restrict our analysis of the proposed heuristic policy to the baseline parameter conditions specific to each case study scenario.

We find that, in the US influenza context, the heuristic policy would involve using a fractional-dose 1D policy on group A until the stockpile is depleted. Notably, group E does not receive any vaccines due to their lower transmissibility and the severely limited stockpile. If the vaccine stockpile was sufficiently larger, the optimal policy would recommend achieving complete vaccination on group A followed by vaccinating group E using the fractional-dose 1D policy. On the other hand, in the Nigerian COVID-19 case, the heuristic policy adopts a *Mix* vaccination regime on group A followed by group E until complete vaccination is achieved in each group. As in the homogeneous population model, the *Mix* regime is an outcome of a limited vaccine supply rate. In this case, group E gets vaccinated, unlike in the influenza case because of a larger stockpile relative to the supply and administration capacities. The dose-sparing regime is not used since the early vaccination regime achieves complete vaccination in both population groups.

To assess the impact of fractional-dose vaccines in each case, we compare the infections averted by our proposed heuristic policy with a benchmark full-dose vaccination policy that we designed based on the approaches used in Więcek et al. (2021) and Bubar et al. (2021). The benchmark policy deploys the full-dose 1D regime on group A, followed by group E, with the switching point optimized to minimize infections. Our numerical analyses indicate that the heuristic policy would have averted 66% and 134% more infections than the benchmark policy in the influenza and COVID-19 case studies, respectively. Our analysis thus offers a straightforward and intuitive approach for adapting the optimal policies from this study to heterogeneous populations. It also demonstrates that fractional-dose vaccines, at least in some parameter scenarios, can significantly improve epidemic outcomes when applied in heterogeneous population settings.