Type Heterogeneity under Private Information: A Public-Good Analysis of Vaccination Behavior

Erwin Amann*[1] and Manar Alyousuf[1]

[1]Department of Business and Economics, University of Duisburg–Essen, Germany,
erwin.amann@uni-due.de manar.alyousuf@uni-due.de

January 20, 2026

Abstract

We study vaccination as a public-good problem with epidemiological externalities and multi-dimensional private heterogeneity. A unit-mass continuum of individuals chooses whether to vaccinate. Vaccination entails a privately observed cost; remaining unvaccinated exposes the individual to an infection loss that depends on the aggregate non-vaccination rate via a simple incidence map. Motivated by phenomena highlighted during the COVID-19 pandemic, in particular the presence of medically vulnerable individuals who cannot safely risk either infection or vaccination and therefore require substantially enhanced indirect protection, we develop a framework that isolates and quantifies the size of a strategically vulnerable subgroup.

We first analyze two benchmark models with one-dimensional heterogeneity (in infection loss and in vaccination cost, respectively) and derive closed-form expressions for equilibrium vaccination outcomes, utilitarian planner allocations, and uniform Pigouvian instruments that implement the planner under private information. For these benchmark models we also show how imperfect vaccine efficacy can be incorporated in a reduced-form way, by rescaling the incidence term, without sacrificing tractability.

Our main contribution is an analytically tractable model with joint (dual) heterogeneity in vaccination costs and infection losses. In this two-dimensional type space, equilibrium is

*Corresponding author: erwin.amann@uni-due.de

characterized by a slope-indexed indifference locus, and the non-vaccination rate exhibits a sharp two-regime structure depending on where this locus intersects the unit square. In the high-transmission regime, the planner's problem reduces to the choice of a single intercept parameter; we obtain closed-form expressions for the optimal intercept and for the residual mass of non-vaccinators, interpreted as an endogenously determined "at-risk" subgroup trapped in a strategic dilemma. This provides a transparent public-good framework for the design of future policy instruments aimed at protecting medically vulnerable individuals, and, to the best of our knowledge, constitutes the first vaccination game to treat vaccination disutility and infection loss as independent continuous private types while retaining full analytical solvability.

Keywords: Public goods; Vaccination games; Externalities; Population games; Pigouvian instruments; Heterogeneity.

JEL: C72, D64, H41, I18.

1 Introduction

Vaccination is a canonical public-good problem. Individuals face idiosyncratic private costs associated with both vaccination and infection, while the infection risk they bear depends on others' immunisation choices through herd immunity and reduced transmission.[1] Each additional vaccination lowers population risk, but this external benefit is not fully internalised in private decisions. As in standard models of public goods, individually optimal vaccination choices generate an outcome that is inefficiently low relative to a utilitarian benchmark.

The COVID-19 pandemic has made these issues particularly salient. It revealed substantial heterogeneity in perceived and actual vaccine risks, in infection losses, and in attitudes towards preventive measures, and it brought to the forefront a clinically vulnerable subgroup: individuals whose health conditions do not allow them to safely risk either infection or vaccination and who therefore depend critically on protection provided by others' immunization decisions. Understanding how such a group emerges in equilibrium, and how its size can be reduced by policy instruments, is central for the design of robust vaccination policies.

A central insight, already present in the public-good formulation of Brito et al. (1991), is that

[1] Vaccination costs may comprise, for example, monetary and time outlays, expected side effects, and perceived health risks associated with vaccination.

vaccination chosen by self-interested individuals falls short of the social optimum even with homogeneous population and perfectly effective vaccines. Subsequent work combines strategic vaccination decisions with epidemic dynamics and social learning (e.g., Bauch et al., 2003; Bauch and Earn, 2004), introduces preference and network heterogeneity (Manfredi et al., 2009; Tassier et al., 2015), and considers imperfect vaccines and interdependent security (e.g., Heal and Kunreuther, 2005). On the policy side, studies compare mandates to voluntary regimes (e.g., Browne, 2016) or analyze subsidy schemes under private information (e.g., Yamin and Gavious, 2013). A robust message emerges: in vaccination games, individually rational behavior does not deliver collectively efficient coverage.

This paper develops a static vaccination game that emphasizes multi-dimensional private heterogeneity and simple, implementable policy instruments, with a particular eye to the protection of medically vulnerable individuals. We model a unit mass of individuals who choose whether to vaccinate. Vaccination entails a privately observed cost; remaining unvaccinated exposes an individual to an infection loss that depends on the aggregate non-vaccination rate through a transparent linear incidence map. Each individual is described by a two-dimensional type capturing vaccination-cost and infection-loss parameters. Within this framework we study three related models and compare the outcome of individual behavior to a utilitarian planner's allocation under information constraints. The analysis isolates, and quantifies in closed form, a subset of types who remain in a strategic dilemma, namely those for whom neither vaccination nor non-vaccination strictly dominates, and provides a formal counterpart to the clinically vulnerable group highlighted by COVID-19.

The contribution is threefold. First, we analyze two benchmark models with one-dimensional heterogeneity—heterogeneous infection losses with homogeneous vaccination cost, and heterogeneous vaccination costs with homogeneous infection loss. In both cases, behavior is characterized by a scalar cutoff, and under linear incidence and uniform types we obtain closed-form expressions for the equilibrium vaccination outcome, the utilitarian planner's allocation, and the net Pigouvian transfer that implements the planner through a uniform subsidy–tax scheme. For these benchmark models we also show how imperfect vaccine efficacy can be incorporated in a reduced-form way, by rescaling the epidemiological incidence term, without sacrificing tractability.

Second, we study a model with joint (dual) heterogeneity in vaccination costs and infection losses. Types are two-dimensional and uniformly distributed, and the expected net gain from vaccinating is linear in the type vector. The non-vaccination region in type space is defined by a linear indifference locus whose slope aggregates epidemiological primitives, economic primitives, and the aggregate non-vaccination rate. We characterize equilibrium non-vaccination in terms of this slope and show that the equilibrium exhibits a sharp two-regime structure, depending on whether the indifference locus intersects the top or the right edge of the unit square. This geometry yields closed-form expressions for the equilibrium non-vaccination rate in each regime. To the best of our knowledge, this is the first vaccination game to treat vaccination disutility and infection loss as independent continuous private types while preserving full analytical tractability.

Third, in the specification with two-dimensional type heterogeneity, we analyze a utilitarian planner who internalizes the infection externality and chooses, in reduced form, among linear cutoff frontiers in type space. The planner's allocation can be represented by an indifference locus with slope pinned down by epidemiological and cost primitives and an intercept parameter that captures the net effect of the policy and behavioral environment. In the high-transmission regime, the planner's problem reduces to the choice of this single intercept. We show that the infeasibility cutoff in the infection-loss dimension is determined solely by primitives, whereas the intercept governs the size and composition of the residual non-vaccinating set. We derive in closed form the mass of individuals who remain in a strategic dilemma, namely those whose type makes both actions risky, and interpret this mass as an endogenously determined "at-risk" subgroup. This yields explicit welfare expressions under multi-dimensional private information and provides a solid analytical basis for assessing and designing future policy measures aimed at protecting medically vulnerable individuals who cannot safely rely on either vaccination or infection.

Section 2 introduces the population game, the incidence structure, and the planner's problem. Section 3 presents the two benchmark models with one-dimensional heterogeneity and characterizes equilibrium outcomes, planner allocations, and Pigouvian instruments, including the reduced-form treatment of imperfect efficacy. Section 4 contains the main analysis of joint heterogeneity, the slope-indexed equilibrium structure, and the planner's problem within the class of linear cutoff frontiers in type space, and derives closed-form expressions for the optimal intercept and the size of the at-risk group. Section 5 discusses policy implications and extensions, and Section 6 concludes. Technical derivations for the joint-heterogeneity model are collected in Appendix B.

2 The model

We consider a continuum population game with a unit mass of nonatomic individuals. Each individual chooses an action $a \in \{V, N\}$, interpreted as vaccinating (V) or not vaccinating (N). Let $x \in [0,1]$ denote vaccination coverage and $\pi = 1 - x$ the share of non-vaccinators.

Vaccination entails a private disutility capturing vaccine side effects, whereas infection entails a loss reflecting the consequences of infection.[2] These primitives are scaled by privately observed types.

Types, payoffs, and incidence. Each individual is endowed with a privately observed type

$$(\lambda, \theta) \in [0,1]^2,$$

where λ indexes idiosyncratic disutility from vaccination and θ idiosyncratic loss in the event of infection. Throughout the main text we assume that (λ, θ) is i.i.d. across individuals and uniformly distributed on $[0,1]^2$; in the benchmark environments below, heterogeneity is restricted to a single dimension.

Vaccinating a type-λ individual entails cost λc_v, with $c_v > 0$. Infection of a type-θ individual generates loss θc_i, with $c_i > 0$.

Let the infection probability of a non-vaccinator be $p(\pi)$, where $p : [0,1] \to [0,1]$ is continuous, strictly increasing, and satisfies $p(0) = 0$. In the body of the paper we adopt the linear incidence specification

$$p(\pi) = \alpha\pi, \tag{1}$$

where $\alpha \in [0,1]$ is a transmission parameter. This delivers transparent expressions while capturing the dependence of individual infection risk on aggregate non-vaccination.[3] We focus on the non-

[2] The parameters c_v and c_i can be interpreted in reduced form as aggregating all relevant components of vaccination cost and infection loss (monetary and time costs, health and productivity effects, and perceived risks). While the baseline formulation is phrased in terms of physiological damage, the analysis applies equally to broader economic and behavioral interpretations.

[3] More generally one may consider a convex incidence function $p(\pi) = \alpha\pi^\eta$ with $\eta \geq 1$. This preserves the cutoff structure of equilibrium strategies and the qualitative comparative statics (higher c_v lowers coverage; higher α or c_i raises coverage). The square-root and rational formulas obtained under the linear case are replaced by power or implicit solutions; see Appendix A.

degenerate case with $\alpha > 0$, $c_i > 0$, and $c_v > 0$. Interior cutoff solutions additionally require c_v not to be too large relative to αc_i; otherwise both equilibrium and planner solutions lie at the boundary with no vaccination.

In the baseline, vaccination is fully effective: vaccinated individuals face zero infection probability. Sections 3.1.1 and 3.2.1 introduce imperfect efficacy in reduced form.

Given a population state π and type (λ, θ), expected utilities in the baseline are

$$\mathbb{E}[u(V) \mid \lambda] = -\, c_v \lambda, \qquad \mathbb{E}[u(N) \mid \theta] = -\, \alpha \pi c_i \theta. \tag{2}$$

An individual vaccinates if and only if

$$c_v \lambda \leq \alpha \pi c_i \theta. \tag{3}$$

Thus equilibrium strategies are threshold strategies in type space: for any given π, the vaccinating set is described by a linear inequality in (λ, θ).

3 Benchmark models with one-dimensional heterogeneity

We begin with two benchmark models that isolate heterogeneity in a single dimension. These environments connect directly to classic public-good models of vaccination and deliver closed-form solutions for the individual equilibrium and the planner allocation, as well as the associated Pigouvian instruments.

3.1 Model 1: Type-dependent c_i, homogeneous c_v

Types differ only in infection loss θ, while vaccination cost is homogeneous. Vaccination is perfectly effective.

Environment. Types satisfy $\theta \sim \text{Unif}[0, 1]$. Vaccination costs $c_v > 0$ are homogeneous, while infection of type θ causes loss θc_i. Given π, expected utilities are

$$\mathbb{E}[u(V)] = -c_v, \qquad \mathbb{E}[u(N \mid \theta, \pi)] = -\alpha \pi c_i \theta.$$

Equilibrium. A type θ vaccinates iff $c_v \le \alpha\pi c_i\theta$. In a cutoff equilibrium, types with $\theta \ge \bar{\theta}$ vaccinate and $\pi = \bar{\theta}$. Imposing indifference for $\theta = \bar{\theta}$ yields

$$c_v = \alpha c_i \bar{\theta}^2,$$

so the interior equilibrium cutoff is

$$\bar{\theta}^* = \sqrt{\frac{c_v}{\alpha c_i}}, \qquad \pi^* = \bar{\theta}^*, \quad x^* = 1 - \bar{\theta}^*, \tag{4}$$

provided $c_v \le \alpha c_i$; for $c_v > \alpha c_i$, the unique equilibrium is $\bar{\theta}^* = 1$ (no vaccination).

Planner. A utilitarian planner chooses $\bar{\theta} \in [0, 1]$ to maximize

$$W(\bar{\theta}) = -\int_{\bar{\theta}}^{1} c_v \, d\theta \; - \; \int_{0}^{\bar{\theta}} \alpha\bar{\theta}c_i\theta \, d\theta = -(1-\bar{\theta})c_v - \frac{1}{2}\alpha c_i \bar{\theta}^3. \tag{5}$$

The first-order condition,

$$\frac{dW}{d\bar{\theta}} = c_v - \frac{3}{2}\alpha c_i \bar{\theta}^2 = 0,$$

yields the optimal cutoff

$$\bar{\theta}^{**} = \sqrt{\frac{2}{3}\frac{c_v}{\alpha c_i}}, \qquad \pi^{**} = \bar{\theta}^{**}, \quad x^{**} = 1 - \bar{\theta}^{**}, \tag{6}$$

for $c_v \le \frac{3}{2}\alpha c_i$; otherwise the optimum is at the boundary. Since $\bar{\theta}^{**} = \sqrt{\frac{2}{3}}\,\bar{\theta}^* < \bar{\theta}^*$, the individual equilibrium exhibits under-vaccination, $x^{**} > x^*$.

Subsidization and taxation. Because infection-loss types are privately observed, the social planner observes only the distribution of θ but not individual realizations. Hence the planner cannot determine which individuals have "high" or "low" infection losses and thus cannot directly target vaccination by type. The planner can, however, influence behavior by subsidizing vaccinators and/or taxing non-vaccinators. Let S denote a subsidy to vaccinators and T a tax on non-vaccinators. Since only $S - T$ matters for behavior, we normalize $T = 0$ and interpret S as the net subsidy.[4]

The cutoff type satisfies

$$-c_v + S = -\alpha\pi c_i\theta \quad \Leftrightarrow \quad c_v - S = \alpha\pi c_i\theta.$$

Evaluating this condition at $(\theta, \pi) = (\bar{\theta}^{**}, \pi^{**})$ with $\pi^{**} = \bar{\theta}^{**}$ gives

$$S^* = c_v - \alpha c_i(\bar{\theta}^{**})^2.$$

Using $\alpha c_i(\bar{\theta}^{**})^2 = \frac{2}{3}c_v$ from the planner's first-order condition, the Pigouvian subsidy is

$$S^* = \frac{1}{3}c_v, \tag{7}$$

which implements the planner's second-best under private information in this benchmark environment.

3.1.1 Imperfect efficacy

Let $\delta \in [0, 1]$ denote the vaccine failure probability: a vaccinated individual remains susceptible with probability δ. With vaccination coverage $x = 1 - \pi$, the mass of susceptibles is $\pi + \delta(1 - \pi)$, so a susceptible individual faces infection probability $\alpha[\pi + \delta(1 - \pi)]$. Expected utilities are

$$\mathbb{E}[u(V \mid \delta, \theta, \pi)] = -c_v - \delta\alpha\pi c_i\theta - \alpha(1-\pi)\delta^2 c_i\theta, \qquad \mathbb{E}[u(N \mid \delta, \theta, \pi)] = -\alpha\pi c_i\theta - \delta\alpha(1-\pi)c_i\theta. \tag{8}$$

In a cutoff equilibrium, $\pi = \bar{\theta}$ and the marginal type satisfies $\mathbb{E}[u(V \mid \delta, \bar{\theta}, \bar{\theta})] = \mathbb{E}[u(N \mid \delta, \bar{\theta}, \bar{\theta})]$. This yields a quadratic in $\bar{\theta}$,

$$\alpha c_i(1 - \delta)^2\bar{\theta}^2 + \alpha c_i\delta(1 - \delta)\bar{\theta} - c_v = 0, \tag{9}$$

[4]Subsidies are financed by non-distortionary lump-sum means, so transfers do not enter the planner's welfare objective.

whose economically relevant root is

$$\bar{\theta}^*(\delta) = \frac{\sqrt{\alpha c_i \delta^2 + 4c_v} - \sqrt{\alpha c_i}\,\delta}{2\sqrt{\alpha c_i}\,(1-\delta)}. \tag{10}$$

This converges to the benchmark cutoff as $\delta \to 0$ and yields the unique interior equilibrium whenever $\bar{\theta}^*(\delta) \in (0,1)$; otherwise the equilibrium is at the boundary $\bar{\theta}^*(\delta) = 1$ (no vaccination). At $\delta = 1$, vaccination is strictly dominated and $\bar{\theta}^*(1) = 1$.

The planner's problem yields a cubic welfare function in $\bar{\theta}$

$$\begin{aligned} W(\bar{\theta},\delta) &= \int_{\bar{\theta}}^{1} \left[-c_v - \delta\bar{\theta}\alpha c_i\theta - \alpha(1-\bar{\theta})\delta^2 c_i\theta\right] d\theta \;+\; \int_{0}^{\bar{\theta}} \left[-\bar{\theta}\alpha c_i\theta - \delta(1-\bar{\theta})\alpha c_i\theta\right] d\theta. \\ &= -\frac{1}{2}\alpha c_i(1-\delta)^2\bar{\theta}^3 - \frac{1}{2}\alpha c_i\delta(1-\delta)\bar{\theta}^2 + \Big(-\frac{1}{2}\alpha c_i\delta(1-\delta) + c_v\Big)\bar{\theta} - \Big(\frac{1}{2}\alpha c_i\delta^2 + c_v\Big). \end{aligned} \tag{11}$$

The first-order condition $\partial W/\partial\bar{\theta} = 0$ yields a quadratic equation in $\bar{\theta}$,

$$\frac{\partial W}{\partial \bar{\theta}} = -\frac{3}{2}\alpha c_i(1-\delta)^2\bar{\theta}^2 - \alpha c_i\delta(1-\delta)\bar{\theta} + \Big(-\frac{1}{2}\alpha c_i\delta(1-\delta) + c_v\Big) = 0. \tag{12}$$

The optimal cutoff is

$$\bar{\theta}^{**}(\delta) = \frac{\sqrt{4\alpha c_i\delta^2 - 3\alpha c_i\delta + 6c_v} - \sqrt{\alpha c_i}\,\delta}{3\sqrt{\alpha c_i}\,(1-\delta)}, \tag{13}$$

with $\lim_{\delta\to 0}\bar{\theta}^{**}(\delta) = \bar{\theta}^{**}$. For parameter values admitting interior solutions, $\bar{\theta}^{**}(\delta) < \bar{\theta}^*(\delta)$, so under-vaccination persists under imperfect efficacy. As $\delta \to 1$, both the equilibrium and planner's solution converge to no vaccination and the inefficiency vanishes.

3.2 Model 2: Type-dependent c_v, homogeneous c_i

Types differ only in idiosyncratic side-effect risk λ, while infection loss is homogeneous. Vaccination is perfectly effective.

Environment. Types satisfy $\lambda \sim \text{Unif}[0,1]$. Vaccination of type λ entails cost λc_v, while each infection causes loss c_i. Expected utilities are

$$\mathbb{E}[u(V \mid \lambda)] = -\lambda c_v, \qquad \mathbb{E}[u(N \mid \pi)] = -\alpha\pi c_i.$$

Equilibrium. A type λ vaccinates iff $\lambda c_v \leq \alpha\pi c_i$. In a cutoff equilibrium, types with $\lambda \leq \bar{\lambda}$ vaccinate and $\pi = 1 - \bar{\lambda}$. Indifference implies

$$\bar{\lambda} c_v = \alpha(1 - \bar{\lambda}) c_i,$$

so the equilibrium cutoff is

$$\bar{\lambda}^* = \frac{\alpha c_i}{c_v + \alpha c_i}, \qquad x^* = \bar{\lambda}^*, \quad \pi^* = 1 - \bar{\lambda}^* = \frac{c_v}{c_v + \alpha c_i}. \tag{14}$$

Equilibrium strategies are monotone in λ: individuals with sufficiently low side-effect risk vaccinate; those with high risk do not. If $\alpha = 0$ or $c_i = 0$, then $\bar{\lambda}^* = 0$ and no one vaccinates.

Planner. Given a cutoff $\bar{\lambda}$, the mass of vaccinators is $\bar{\lambda}$, each incurring side-effect cost λc_v; the mass of non-vaccinators is $\pi = 1 - \bar{\lambda}$, each facing expected infection loss $\alpha\pi c_i$. Aggregate welfare is

$$W(\bar{\lambda}) = -\frac{c_v}{2}\bar{\lambda}^2 - \alpha c_i (1 - \bar{\lambda})^2. \tag{15}$$

The first-order condition,

$$\frac{dW}{d\bar{\lambda}} = -c_v \bar{\lambda} + 2\alpha c_i (1 - \bar{\lambda}) = 0,$$

yields the planner's cutoff

$$\bar{\lambda}^{**} = \frac{2\alpha c_i}{c_v + 2\alpha c_i}, \qquad \pi^{**} = 1 - \bar{\lambda}^{**} = \frac{c_v}{c_v + 2\alpha c_i}. \tag{16}$$

Whenever $0 < \bar{\lambda}^{**} < 1$, we have $\bar{\lambda}^{**} > \bar{\lambda}^*$ and $x^{**} > x^*$: voluntary coverage is inefficiently low.

Subsidization and taxation. With privately observed λ, the planner again can use non-discriminatory transfers. Let S denote a subsidy to vaccinators and T a tax on non-vaccinators, with $T = 0$ without loss of generality.[5]

Under a uniform subsidy S, a type λ receives $-\lambda c_v + S$ if vaccinated and $-\alpha\pi c_i$ if not, so the

[5]Subsidies are financed by non-distortionary lump-sum means, so transfers do not enter the planner's welfare objective.

cutoff type satisfies

$$-\lambda c_v + S = -\alpha\pi c_i \quad\Leftrightarrow\quad \lambda c_v - S = \alpha\pi c_i.$$

Evaluating at $(\bar{\lambda}^{**}, \pi^{**})$ yields

$$S^* = \bar{\lambda}^{**} c_v - \alpha\pi^{**} c_i = \frac{\alpha c_i c_v}{c_v + 2\alpha c_i}, \tag{17}$$

which implements the planner's cutoff under private information in this benchmark environment.

3.2.1 Imperfect efficacy

Let $\delta \in [0,1]$ denote the vaccine failure probability. With coverage $\bar{\lambda}$, the non-vaccination rate is $\pi = 1 - \bar{\lambda}$, and the mass of susceptibles is $\pi + \delta(1-\pi)$. A susceptible individual faces infection probability $\alpha[\pi + \delta(1-\pi)]$. Expected utilities are

$$\mathbb{E}\big[u(V) \mid \lambda, \delta, \pi\big] = -c_v\lambda - \delta\alpha\pi c_i - \alpha(1-\pi)\delta^2 c_i, \qquad \mathbb{E}\big[u(N) \mid \delta, \pi\big] = -\alpha\pi c_i - \delta\alpha(1-\pi)c_i. \tag{18}$$

In a cutoff equilibrium, $\pi = 1 - \bar{\lambda}$ and the marginal type satisfies $\mathbb{E}[u(V) \mid \bar{\lambda}, \delta, \pi] = \mathbb{E}[u(N) \mid \delta, \pi]$, which simplifies to

$$\bar{\lambda}\big[c_v + \alpha c_i(1-\delta)^2\big] = \alpha c_i(1-\delta).$$

The equilibrium cutoff is

$$\bar{\lambda}^*(\delta) = \frac{\alpha c_i(1-\delta)}{c_v + \alpha c_i(1-\delta)^2}, \tag{19}$$

with non-vaccination rate

$$\pi^*(\delta) = 1 - \bar{\lambda}^*(\delta) = \frac{c_v + \alpha c_i\big[(1-\delta)^2 - (1-\delta)\big]}{c_v + \alpha c_i(1-\delta)^2}. \tag{20}$$

We have

$$\lim_{\delta\to 0} \bar{\lambda}^*(\delta) = \frac{\alpha c_i}{c_v + \alpha c_i}, \qquad \lim_{\delta\to 1} \bar{\lambda}^*(\delta) = 0,$$

so the equilibrium converges to the benchmark as $\delta \to 0$ and collapses to zero vaccination when the vaccine is useless.

Aggregate welfare can be written as

$$\begin{aligned} W(\bar{\lambda}, \delta) &= \int_0^{\bar{\lambda}} \left[-c_v \lambda - \delta\alpha(1-\bar{\lambda})c_i - \alpha\bar{\lambda}\delta^2 c_i \right] d\lambda \ + \ \int_{\bar{\lambda}}^1 \left[-\alpha(1-\bar{\lambda})c_i - \delta\alpha\bar{\lambda}c_i \right] d\lambda \\ &= \left[-\alpha c_i(1-\delta)^2 - \frac{c_v}{2} \right] \bar{\lambda}^2 + 2\alpha c_i(1-\delta)\bar{\lambda} - \alpha c_i, \end{aligned} \tag{21}$$

so the planner's first-order condition yields

$$\bar{\lambda}^{**}(\delta) = \frac{2\alpha c_i(1-\delta)}{2\alpha c_i(1-\delta)^2 + c_v}, \tag{22}$$

with

$$\lim_{\delta\to 0} \bar{\lambda}^{**}(\delta) = \frac{2\alpha c_i}{2\alpha c_i + c_v}, \qquad \lim_{\delta\to 1} \bar{\lambda}^{**}(\delta) = 0.$$

For any $\delta \in [0, 1)$ such that both solutions are interior, the planner vaccinates strictly more types:

$$\bar{\lambda}^{**}(\delta) - \bar{\lambda}^{*}(\delta) = \frac{\alpha c_i c_v (1-\delta)}{\left[c_v + \alpha c_i(1-\delta)^2\right]\left[c_v + 2\alpha c_i(1-\delta)^2\right]} \ > \ 0. \tag{23}$$

The gap is strictly positive for any imperfect but non-degenerate vaccine and converges to zero only as $\delta \to 1$.

Connection between Models 1 and 2. Comparing Model 1 (heterogeneity in infection loss) and Model 2 (heterogeneity in side-effect risk) illustrates that the qualitative externality is robust to the source of heterogeneity: voluntary coverage is inefficiently low, and the planner's threshold exceeds the decentralized threshold in the relevant type dimension. Imperfect efficacy weakens the epidemiological leverage of each vaccination but does not overturn this ranking for any imperfect yet non-degenerate vaccine. In both models, the inefficiency can be mitigated by simple linear Pigouvian instruments (subsidies to vaccination or taxes on non-vaccination) that shift the cutoff in the relevant type dimension while preserving the underlying threshold structure. Taken together, the two benchmark models are best viewed as two sides of the same coin.

4 Model 3: Joint heterogeneity, equilibrium geometry, and the at-risk group

We now turn to our main theoretical environment, in which both vaccination costs and infection losses are heterogeneous and privately observed. Vaccination remains fully effective in this model; imperfect efficacy is confined to the benchmark models.

4.1 Joint heterogeneity and equilibrium geometry

Setup. Types are two-dimensional and i.i.d. uniform: each individual observes $(\lambda, \theta) \in [0,1]^2$, where vaccination imposes cost λc_v and infection entails loss θc_i. The non-vaccination rate is π, infection risk is $p(\pi) = \alpha\pi$, and vaccination is fully effective.

Given (2), expected utilities are

$$\mathbb{E}[u(V \mid \lambda)] = -c_v\lambda, \qquad \mathbb{E}[u(N \mid \theta, \pi)] = -\alpha\pi c_i\theta.$$

Vaccination is (weakly) optimal if and only if

$$\lambda \le k\,\theta, \qquad k := \frac{\alpha\pi c_i}{c_v}. \tag{24}$$

Thus the allocation is governed by the slope k of the indifference ray in (θ, λ)-space, which aggregates epidemiological parameters, cost scales, and the endogenous non-vaccination rate.

Let $\mathcal{N}(k) = \{(\lambda, \theta) \in [0,1]^2 : \lambda > k\theta\}$ denote the non-vaccination region. With uniform types, its Lebesgue measure equals the non-vaccination rate, $\pi = \text{area}(\mathcal{N}(k))$. Equilibrium values of π and k are jointly determined by the consistency conditions

$$\pi = \text{area}(\mathcal{N}(k)) \quad \text{and} \quad k = \frac{\alpha\pi c_i}{c_v}.$$

Proposition 1 (Slope-indexed equilibrium structure)**.** *Under joint heterogeneity with uniformly distributed* (λ, θ) *and linear incidence* $p(\pi) = \alpha\pi$*, the equilibrium non-vaccination rate* π^* *is given* *by*

$$\pi^* = \begin{cases} \dfrac{2c_v}{2c_v + \alpha c_i} \in (0,1), & k^* < 1, \quad \textit{with } \lambda_1^* = \dfrac{2\alpha c_i}{2c_v + \alpha c_i} < 1, \\ \sqrt{\dfrac{c_v}{2\alpha c_i}} \in (0,1), & k^* > 1, \quad \textit{with } \theta_1^* = \sqrt{\dfrac{2c_v}{\alpha c_i}} < 1, \end{cases} \tag{25}$$

where $k^* = \frac{\alpha \pi^* c_i}{c_v}$ *is the equilibrium slope. The low-slope regime* $k^* < 1$ *occurs if and only if* $\alpha < \frac{2c_v}{c_i}$*, while the high-slope regime* $k^* > 1$ *occurs if and only if* $\alpha > \frac{2c_v}{c_i}$*.*

The proof is a straightforward geometric calculation in $[0,1]^2$ and is provided in Appendix B. In the low-slope regime the ray $\lambda = k^*\theta$ intersects the right edge of the unit square at $(\theta, \lambda) = (1, \lambda_1^*)$, and the behavioral margin is primarily along the infection-loss dimension θ. In the high-slope regime it intersects the top edge at $(\theta, \lambda) = (\theta_1^*, 1)$, and the behavioral margin is primarily along the vaccination-cost dimension λ.

Economically, the condition $\alpha < \frac{2c_v}{c_i}$ characterizes environments with moderate transmission intensity relative to the vaccination–infection cost ratio, a configuration that is naturally associated with seasonal influenza–type pathogens. In such settings, a large share of individuals is near the vaccination margin because their perceived infection consequences θc_i are relatively modest. Small changes in perceived or actual infection losses—for instance, revised assessments of influenza severity in specific risk groups—shift the indifference locus primarily along the θ-dimension and thereby reclassify a sizable mass of individuals between vaccination and non-vaccination.

By contrast, $\alpha > \frac{2c_v}{c_i}$ corresponds to high-transmission environments, such as those observed during major SARS-CoV-2 (COVID-19) waves. Here, the probability of infection is sufficiently large that, for a broad range of θ, infection consequences are already substantial. The predominant source of heterogeneity at the margin is then the disutility from vaccination, λc_v. In the model, this is reflected in a steep indifference frontier: small reductions in perceived or effective vaccination disutility—through improved safety communication, reduced access costs, or mitigation of side-effect concerns—translate into substantial movements along the λ-dimension and can mobilize a large mass of individuals whose infection losses are already high.

In what follows, we restrict attention to the case in which $\lambda/\theta > 1$. Addressing the opposite case is straightforward and does not alter the robustness of the conclusions.

4.2 Planner's optimization and the at-risk group

In the joint-heterogeneity environment, a utilitarian social planner internalizes the infection externality and chooses a measurable partition of the type space $[0,1]^2$ into vaccinating and non-vaccinating regions. Because payoffs are linear in types and types are uniformly distributed, we restrict attention to allocations that can be represented by a linear cutoff frontier in (θ, λ)-space, as illustrated in Figure 1.

In the high-slope regime characterized in Proposition 1, the equilibrium indifference locus intersects the top edge of the unit square at $(\theta_1, 1)$. We have shown that the corresponding threshold type

$$\theta_1^{**}$$

is pinned down by primitives and is independent of the intercept of the cutoff line. In particular, combining the expression for the intersection point

$$\theta_1 = \frac{1-\lambda_0}{k} = \frac{(1-\lambda_0)c_v}{\alpha\pi c_i}$$

with the induced non-vaccination rate

$$\pi(\lambda_0) = \frac{(1-\lambda_0)\theta_1}{2}$$

implies

$$\theta_1 = \frac{(1-\lambda_0)c_v}{\alpha c_i \cdot \frac{(1-\lambda_0)\theta_1}{2}} = \frac{2c_v}{\alpha c_i\, \theta_1},$$

and therefore

$$\theta_1^2 = \frac{2c_v}{\alpha c_i} \quad \Rightarrow \quad \theta_1^{**} = \sqrt{\frac{2c_v}{\alpha c_i}}.$$

It is convenient to define the scale parameter

$$z := \theta_1^{**} = \sqrt{\frac{2c_v}{\alpha c_i}} \in (0,1), \tag{26}$$

which summarizes the joint impact of the cost scales c_v, c_i and the transmission parameter α.

Given z, the planner's remaining degree of freedom is entirely captured by the intercept $\lambda_0 \in$

$[0,1]$ of the cutoff frontier on the λ-axis. At $\theta = 0$, infection loss vanishes, so individuals with types $(\lambda, 0)$ derive no private benefit from vaccination; for these types vaccination is purely costly. The intercept λ_0 therefore has a natural interpretation: it is the maximal vaccination-cost realization at $\theta = 0$ that the planner succeeds in mobilizing to vaccinate, for example through pecuniary or non-pecuniary instruments such as appeals to altruism, professional norms, or mandates.

Under this restricted class of allocations, the vaccinating region consists of

$$\big\{(\lambda,\theta) : \theta \in [0,z],\ \lambda \le \lambda(\theta)\big\} \cup \big\{(\lambda,\theta) : \theta \in [z,1],\ \lambda \in [0,1]\big\},$$

while the non-vaccinating region is its complement in $[0,1]^2$. The cutoff frontier is given by the affine function

$$\lambda(\theta) = \lambda_0 + (1-\lambda_0)\,\frac{\theta}{z}, \tag{27}$$

which connects the points $(0, \lambda_0)$ and $(z, 1)$. Note that for any choice of $\lambda_0 \in [0,1]$,

$$\lambda(z) = \lambda_0 + (1-\lambda_0)\frac{z}{z} = 1,$$

so all admissible frontiers pivot around the fixed point $(\theta_1^{**}, 1) = (z, 1)$. From a geometric perspective, varying λ_0 rotates the boundary clockwise or counterclockwise around $(z, 1)$ and thereby adjusts the set of low-θ, high-λ types that are mobilized to vaccinate.

Under this frontier, the mass of non-vaccinating types is

$$\pi(\lambda_0) = \int_0^z \big[1 - \lambda(\theta)\big]\, d\theta = \frac{(1-\lambda_0)z}{2}, \tag{28}$$

i.e. the area of the triangle above the cutoff line and below the top edge over $\theta \in [0, z]$. The planner's problem in the high-slope regime thus reduces to choosing λ_0 so as to trade off (i) the additional vaccination costs incurred by mobilizing increasingly high-λ types with θ near zero, against (ii) the reduction in aggregate infection losses due to the induced decline in $\pi(\lambda_0)$.

Aggregate welfare as a function of λ_0 can be written as

$$W(\lambda_0) = \int_0^z \int_{\lambda(\theta)}^1 \big(-\pi(\lambda_0)\,\alpha c_i \theta\big)\, d\lambda\, d\theta + \int_0^z \int_0^{\lambda(\theta)} \big(-c_v \lambda\big)\, d\lambda\, d\theta + \int_z^1 \int_0^1 \big(-c_v \lambda\big)\, d\lambda\, d\theta, \tag{29}$$

with $\lambda(\theta)$ given by (27) and z by (26). Evaluating the integrals and differentiating with respect to λ_0 yields a unique maximizer $\lambda_0^{**}(z) \in (0,1)$.

Proposition 2 (Constrained-optimal intercept and at-risk mass)**.** *In the high-slope regime, and within the class of allocations generated by linear cutoff frontiers passing through* $(\theta_1^{**}, 1)$*, the utilitarian planner's problem reduces to choosing the intercept* $\lambda_0 \in [0,1]$*. The unique optimal intercept is*

$$\lambda_0^{**}(z) = \frac{1 - \frac{z}{6}}{1 + \frac{z}{3}} \in (0,1), \tag{30}$$

where $z = \sqrt{2c_v/(\alpha c_i)}$*. The corresponding non-vaccination rate—the mass of individuals in the residual "at-risk" group, for whom neither vaccination nor remaining unvaccinated is uniformly safe in health terms—is*

$$\pi^{**}(z) = \frac{z^2}{4\big(1 + \frac{z}{3}\big)} \in (0,1), \tag{31}$$

which is strictly increasing and strictly convex in $z \in (0,1]$*.*[6] *The residual mass* $\pi^{**}(z)$ *corresponds to the triangular at-risk region in* (θ, λ)*-space illustrated in Figure 1.*

Proof. See Appendix B. □

Proposition 2 shows that, within this natural class of linear cutoff allocations, the planner rotates the frontier around $(\theta_1^{**}, 1)$ as far as is welfare-improving but cannot drive the non-vaccination rate to zero. The residual mass $\pi^{**}(z)$ constitutes an endogenous at-risk subgroup: individuals with relatively low infection losses $\theta \le z$ but sufficiently high vaccination costs $\lambda \ge \lambda_0^{**} + (1-\lambda_0^{**})\theta/z$ that mobilization beyond λ_0^{**} is not socially justified. Its size depends on the scale parameter z, which is increasing in c_v and decreasing in αc_i. Hence higher vaccination costs or lower infection losses lead to a larger residual mass of individuals who remain unvaccinated even under the planner's optimal use of the available (non-pecuniary) mobilization margin.

This at-risk subgroup is particularly relevant in epidemiological environments such as COVID-19, in which medically fragile individuals (for example, those with severe comorbidities or immunosuppression) may face substantial physiological risks both from infection and from vaccination. The

[6] Formally, one can microfound the intercept λ_0, for example via a uniform subsidy scheme that mobilizes individuals to vaccinate and thereby implements the planner's second-best allocation. We do not commit to a specific micro-founded instrument, however, and interpret λ_0 more broadly as a reduced-form index of non-targeted policy and social forces that shift the vaccination margin.

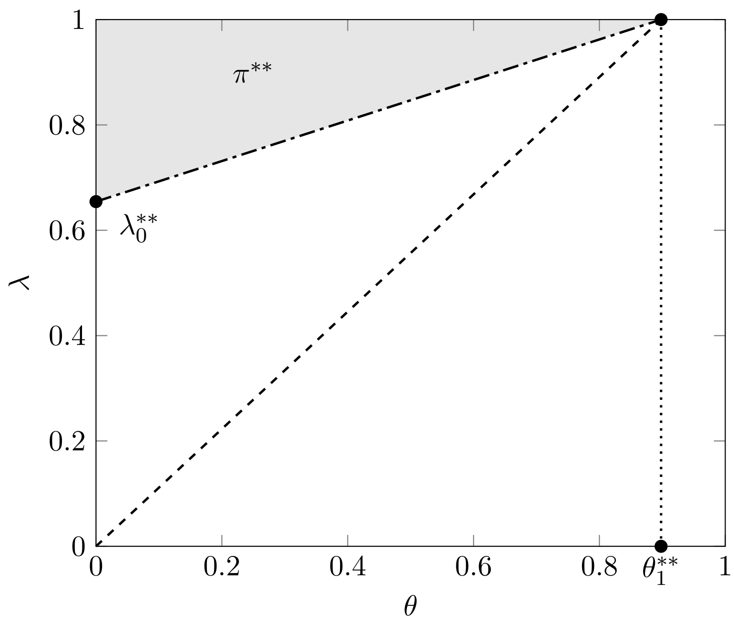

Figure 1: At-risk population in the type space $(\theta, \lambda) \in [0,1]^2$. The policy boundary $\lambda = \lambda_0^{**} + k\,\theta$ (dash–dotted) runs from $(0, \lambda_0^{**})$ to $(\theta_1^{**}, 1)$; the infeasibility threshold at $\theta = \theta_1^{**}$ is shown as a vertical dotted line; and the reference line from $(0,0)$ to $(\theta_1^{**}, 1)$ is dashed. Parameters: $c_v = 0.15$, $c_i = 0.60$, $\alpha = 0.62$. Implied values: $\frac{\alpha c_i}{c_v} = 2.479\,98$, $\theta_1^{**} = \sqrt{2c_v/\alpha c_i} \approx 0.898\,01$, $\lambda_0^{**} \approx 0.6544$, and $\pi^{**} \approx 0.155\,17$. $\alpha = 0.62$ is the estimated child-to-child transmission probability of SARS-CoV-2; see van Boven et al. (2024).

model provides a closed-form expression for the measure of such strategically vulnerable types and a clear link between this measure and the underlying epidemiological and cost parameters, thereby offering a tractable foundation for the design and calibration of protective policies targeted at these groups.

5 Policy implications

The comparative statics of $\pi^{**}(z)$ yield a simple organizing principle for policy. The size of the at-risk subgroup increases with the scale parameter $z = \sqrt{2c_v/(\alpha c_i)}$, that is, with higher vaccination costs or lower infection losses, and it does so at an increasing rate. In high-transmission environments (large α, as in severe COVID-19 waves), even a planner who is able to mobilize a substantial set of low-θ, high-λ types cannot eliminate non-vaccination entirely; policy should therefore prioritize instruments that approximate a high intercept λ_0, by encouraging or requiring vaccination among individuals who face little direct infection loss but whose behavior is epidemiologically pivotal, such as younger health-care workers in critical infrastructures (hospitals, outpatient clinics, and nursing or elderly-care homes).[7] For medically fragile individuals, whose health status makes both infection and vaccination potentially risky, the model provides an explicit expression for the measure of such strategically vulnerable types and a clear link between this measure and the underlying epidemiological and cost parameters.

6 Conclusion

We have developed a public-good model of vaccination with epidemiological externalities and type heterogeneity under private information. A linear incidence map and uniform type distributions yield tractable benchmark models with one-dimensional heterogeneity, in which unregulated vaccination is inefficiently low relative to the utilitarian planner. In these benchmarks, simple uniform Pigouvian instruments, namely a subsidy to vaccination and/or a tax on non-vaccination, implement the planner's allocation under private information, and welfare can be compared in closed form between the individual equilibrium and the planner's cutoff. Imperfect vaccine efficacy can be incorporated in a reduced-form way without sacrificing tractability, and it systematically widens the wedge between privately chosen and socially optimal coverage.

Our main contribution is the analysis of joint heterogeneity in vaccination costs and infection losses. In this environment the equilibrium non-vaccination rate is governed by the slope of a linear indifference locus in type space and exhibits a sharp two-regime structure, depending on whether the decision frontier intersects the right or the top edge of the type space. In the high-transmission regime, the relevant frontier pivots around a primitive-determined infection-loss threshold, and a utilitarian planner is effectively restricted to selecting the intercept of this frontier on the vaccination-cost axis. We derive closed-form expressions for the optimal intercept and for the size of the resulting residual non-vaccinating mass, which we interpret as an endogenously determined at-risk subgroup. These expressions highlight how vaccination costs and infection harms jointly shape policy leverage and the residual vulnerability of the population.

The framework provides an analytically transparent basis for thinking about recent and future

[7] A detailed analysis of concrete policy instruments and their implementation is beyond the scope of this paper and is left to future research.

vaccine-preventable epidemics, including COVID-19, in which medically vulnerable individuals often cannot freely accept either the risk of infection or the risk of adverse vaccine reactions. By deriving the size of this endogenously determined at-risk group within a fully specified game-theoretic environment, the model offers a solid foundation for constructing and calibrating future policy environments and protective arrangements targeted at such populations.

Several extensions suggest themselves. First, allowing for incomplete information and belief heterogeneity about infection risk or vaccine safety would introduce strategic uncertainty and potentially richer equilibrium phenomena. Second, embedding the static interaction in a dynamic epidemic model could connect the static at-risk subgroup to long-run prevalence and disease burden. Third, analyzing richer policy instrument sets, such as targeted subsidies, mandates, or screening mechanisms, would help to address informational constraints regarding types and to approximate the planner's second-best allocation. We hope that the simple structure and closed-form results presented here can serve as a useful benchmark for such extensions.

Appendix A Robustness to convex incidence

Throughout the main text we assume a linear incidence function

$$p(\pi) = \alpha\pi. \tag{32}$$

This appendix shows that the main structural and comparative-static results are robust to replacing (32) by a convex incidence function

$$p(\pi) = \alpha\pi^{\eta}, \qquad \eta \geq 1, \tag{33}$$

at the cost of losing some of the simple closed forms.

Model 1 (heterogeneous c_i, common c_v). With incidence (33), an unvaccinated type θ faces expected loss $-\alpha\pi^{\eta}c_i\theta$, while vaccination yields $-c_v$. Given π, type θ vaccinates iff $c_v \leq \alpha\pi^{\eta}c_i\theta$, so best responses are cutoff strategies in θ. Under $\theta \sim U[0,1]$, the non-vaccination rate equals the cutoff, $\pi = \bar{\theta}$, and the equilibrium cutoff solves

$$c_v = \alpha c_i \bar{\theta}^{\eta+1} \quad \Rightarrow \quad \bar{\theta}^*(\eta) = \left(\frac{c_v}{\alpha c_i}\right)^{\frac{1}{\eta+1}}.$$

Thus the cutoff structure is preserved and the qualitative comparative statics are unchanged: $\bar{\theta}^*(\eta)$ is increasing in c_v and decreasing in α and c_i, so equilibrium coverage $x^*(\eta) = 1-\bar{\theta}^*(\eta)$ is decreasing in c_v and increasing in α and c_i. For $\eta = 1$ this reduces to the square-root expression in the main text.

Model 2 (heterogeneous c_v, common c_i). With incidence (33), an unvaccinated individual's expected loss is $-\alpha\pi^{\eta}c_i$, while a type-λ vaccinator incurs $-\lambda c_v$. Given π, type λ vaccinates iff $\lambda c_v \leq \alpha\pi^{\eta}c_i$, so best responses are cutoff strategies in λ. Under $\lambda \sim U[0,1]$, the vaccination rate equals the cutoff, $x = \bar{\lambda}$, so $\pi = 1 - \bar{\lambda}$ and the equilibrium cutoff solves

$$\bar{\lambda}c_v = \alpha c_i(1-\bar{\lambda})^{\eta}.$$

Defining $f(\bar{\lambda}) = \bar{\lambda}c_v - \alpha c_i(1-\bar{\lambda})^\eta$, we have $f(0) = -\alpha c_i < 0$, $f(1) = c_v > 0$, and $f'(\bar{\lambda}) = c_v + \alpha c_i \eta(1-\bar{\lambda})^{\eta-1} > 0$, so there exists a unique $\bar{\lambda}^*(\eta) \in (0,1)$. Implicit differentiation of $f(\bar{\lambda}^*(\eta)) = 0$ yields $\partial\bar{\lambda}^*/\partial c_v < 0$ and $\partial\bar{\lambda}^*/\partial\alpha > 0$, $\partial\bar{\lambda}^*/\partial c_i > 0$, so equilibrium coverage $x^*(\eta) = \bar{\lambda}^*(\eta)$ is decreasing in c_v and increasing in α and c_i. For $\eta = 1$, we recover the rational closed-form expression in the main text; for general $\eta \geq 1$, the cutoff is characterized implicitly by the above equation.

In sum, replacing the linear incidence (32) by a convex power function (33) preserves (i) threshold strategies in the relevant type dimension and (ii) the sign of the key comparative statics, while replacing the square-root and rational formulas by power or implicit solutions.

Appendix B Joint heterogeneity in vaccination and infection costs: derivations

This appendix collects derivations for the joint-heterogeneity environment summarized in Section 4.

Appendix B.1 Individual choice and the indifference frontier

Types are two-dimensional and i.i.d., uniformly distributed. Each individual observes $(\lambda, \theta) \in [0,1]^2$, with vaccination cost λc_v and infection loss θc_i. Let the unvaccinated share be $\pi \in [0,1]$ and exposure be $p(\pi) = \alpha\pi$ for $\alpha \in (0,1]$. Vaccination is fully effective. Expected utilities are

$$\mathbb{E}[u(V \mid \lambda)] = -c_v\lambda, \qquad \mathbb{E}[u(N \mid \theta)] = -\alpha\pi c_i\theta. \tag{34}$$

An individual vaccinates if and only if

$$\lambda \leq k\theta, \qquad k := \frac{\pi\alpha c_i}{c_v}. \tag{35}$$

The equilibrium is governed by the slope of the indifference ray k, which aggregates the epidemiological force of infection (α), the economic primitives (c_i, c_v), and the endogenous non-vaccination share (π). For $(\lambda, \theta) \sim \text{Unif}[0,1]^2$, the non-vaccinator set

$$\mathcal{N}(k) = \{(\lambda, \theta) \in [0,1]^2 : \ \lambda > k\theta\}$$

has Lebesgue measure $\pi(k)$, and coverage is $x(k) = 1 - \pi(k)$. Solving delivers an individual equilibrium with a slope-based regime change at $k = 1$.

Appendix B.2 Slope-indexed equilibria

We distinguish two behavioral regimes, each yielding an interior equilibrium under specific parameter constraints.

Low-slope equilibrium $k < 1$. The indifference ray hits the right boundary at $(\theta, \lambda) = (1, k)$, and the non-vaccinating region is the rectangle $\{(\lambda, \theta) : \lambda \in [k, 1], \theta \in [0, 1]\}$. Hence

$$\pi(k) = 1 - k.$$

The equilibrium condition $k = \alpha\pi c_i/c_v$ therefore reads

$$k = \frac{\alpha c_i}{c_v}(1 - k),$$

which yields

$$k^* = \frac{\alpha c_i}{2c_v + \alpha c_i}, \qquad \pi^* = 1 - k^* = \frac{2c_v}{2c_v + \alpha c_i}.$$

This regime occurs if and only if $k^* < 1$, i.e.,

$$\alpha < \frac{2c_v}{c_i}.$$

High-slope equilibrium $k > 1$. Now the indifference ray hits the top boundary at $(\theta_1, 1)$, where $\theta_1 = 1/k$. The non-vaccinating region is the triangle above the ray and below the top edge for $\theta \in [0, \theta_1]$, so

$$\pi(k) = \frac{\theta_1}{2} = \frac{1}{2k}.$$

The equilibrium condition $k = \alpha\pi c_i/c_v$ becomes

$$k = \frac{\alpha c_i}{c_v} \cdot \frac{1}{2k},$$

so that

$$k^2 = \frac{\alpha c_i}{2c_v}, \qquad \pi^* = \frac{1}{2k^*} = \sqrt{\frac{c_v}{2\alpha c_i}}, \qquad \theta_1^* = \frac{1}{k^*} = \sqrt{\frac{2c_v}{\alpha c_i}}.$$

This regime occurs if and only if $k^* > 1$, i.e.,

$$\alpha > \frac{2c_v}{c_i}.$$

Appendix B.3 Endogenous unvaccinated mass and the scale parameter z

With a policy-induced intercept λ_0 on the λ-axis, the indifference locus is

$$\lambda(\theta) = \lambda_0 + k\theta.$$

In the high-slope regime, the line intersects the top boundary at $(\theta_1, 1)$, where

$$\theta_1 = \frac{1-\lambda_0}{k} = \frac{(1-\lambda_0)c_v}{\alpha \pi c_i}.$$

The induced unvaccinated mass is

$$\pi = \int_0^{\theta_1} (1-\lambda(\theta))\, d\theta = \int_0^{\theta_1} (1-\lambda_0-k\theta)\, d\theta = (1-\lambda_0)\theta_1 - \frac{k\theta_1^2}{2} = \frac{(1-\lambda_0)\theta_1}{2}. \tag{36}$$

Substituting for θ_1 and k and solving yields the infeasibility cutoff

$$\theta_1^{**} = \sqrt{\frac{2c_v}{\alpha c_i}} =: z,$$

and the compact representation

$$\pi(\lambda_0) = \frac{(1-\lambda_0)z}{2}.$$

The scale parameter z summarizes the joint impact of epidemiological and cost primitives.

Appendix B.4 Planner's optimization

The planner chooses $\lambda_0 \in [0, 1)$ to maximize aggregate welfare:

$$\begin{aligned} W(\lambda_0) &= \int_0^z \int_{\lambda(\theta)}^1 (-\pi\alpha c_i \theta)\, d\lambda d\theta \\ &+ \int_0^z \int_0^{\lambda(\theta)} (-c_v \lambda) d\lambda d\theta \\ &+ \int_z^1 \int_0^1 (-c_v \lambda)\, d\lambda d\theta \\ &= (-\pi\alpha c_i) \int_0^z (1-\lambda_0) \left(1 - \frac{\theta}{z}\right) d\theta \\ &+ \left(-\frac{c_v}{2}\right) \int_0^z \left(\frac{\theta}{2} + \lambda_0 \cdot \frac{z-\theta}{z}\right)^2 d\theta \\ &+ \int_z^1 \left(-\frac{c_v}{2}\right) d\theta, \end{aligned} \tag{37}$$

with $\lambda(\theta) = \lambda_0 + (1-\lambda_0)\frac{\theta}{z}$ and $z = \sqrt{2c_v/(\alpha c_i)}$. Using $\pi(\lambda_0) = (1-\lambda_0)z/2$ and collecting terms yields the first-order condition $W'(\lambda_0) = 0$, whose unique solution in $(0, 1)$ can be written

$$\lambda_0^{**} = \frac{1 - \frac{z}{6}}{1 + \frac{z}{3}} \in (0, 1).$$

Substituting λ_0^{**} into $\pi(\lambda_0)$ gives

$$\pi^{**}(z) = \frac{z^2}{4\left(1 + \frac{z}{3}\right)},$$

as stated in Proposition 2.

References

Bauch, C. T. and Earn, D. J. (2004). Vaccination and the theory of games. *Proceedings of the National Academy of Sciences*, 101(36):13391–13394.

Bauch, C. T., Galvani, A. P., and Earn, D. J. (2003). Group interest versus self-interest in smallpox vaccination policy. *Proceedings of the National Academy of Sciences*, 100(18):10564–10567.

Brito, D. L., Sheshinski, E., and Intriligator, M. D. (1991). Externalities and compulsary vaccinations. *Journal of Public Economics*, 45(1):69–90.

Browne, K. (2016). The measles and free riders: California's mandatory vaccination law. *Cambridge Quarterly of Healthcare Ethics*, 25(3):472–478.

Heal, G. and Kunreuther, H. (2005). Vaccination games with peer effects in a heterogeneous hospital worker population. *Columbia Business School and Wharton Risk Management and Decision Processes Center, Working Paper.*

Manfredi, P., Della Posta, P., d'Onofrio, A., Salinelli, E., Centrone, F., Meo, C., and Poletti, P. (2009). Optimal vaccination choice, vaccination games, and rational exemption: an appraisal. *Vaccine*, 28(1):98–109.

Tassier, T., Polgreen, P., and Segre, A. (2015). Vaccination games with peer effects in a heterogeneous hospital worker population. *Administrative Sciences*, 5(1):2–26.

van Boven, M., van Dorp, C. H., Westerhof, I., Jaddoe, V., Heuvelman, V., Duijts, L., Fourie, E., Sluiter-Post, J., van Houten, M. A., Badoux, P., Euser, S., Herpers, B., Eggink, D., de Hoog, M., Boom, T., Wildenbeest, J., Bont, L., Rozhnova, G., Bonten, M. J., Kretzschmar, M. E., and Bruijning-Verhagen, P. (2024). Estimation of introduction and transmission rates of sars-cov-2 in a prospective household study. *PLOS Computational Biology*, 20(1):1–20.

Yamin, D. and Gavious, A. (2013). Incentives' effect in influenza vaccination policy. *Management Science*, 59(12):2667–2686.