Vaccination Equilibrium: Externality and Efficiency

Ching-to Albert Ma

Department of Economics, Boston University

ma@bu.edu

May 2024

Abstract

I study individual consumers' choices of getting vaccinated. A vaccine reduces the severity of an infectious illness, but may produce side-effects and other disutilities at the time of administration. Such private benefits and disutilities vary across consumers in the population. The infection probability depends negatively on the total mass of vaccinated consumers. This is an externality. One consumer's vaccination choice has negligible contribution to the total mass of vaccinated consumers. Consumers do not internalize the externality. I characterize a unique Vaccination Equilibrium, the sustainable vaccination mass resulting from individual decisions. I show how vaccine improvements in benefits, side-effects, and infection likelihood change the Vaccination Equilibrium. The first-best or efficient vaccination mass takes into account the externality and consumers' benefits and costs. The unique Vaccination Equilibrium is never first best. Taxes, subsidies, and mandates may change the Vaccination Equilibrium.

Keywords: Vaccination, equilibrium, externality, efficiency, free-riding, vaccine costs and benefits, infection

JEL: I10, H00

Acknowledgment: My thanks to Jacopo Bizzotto, Lester Chan, Randy Ellis, Bernhard Ganglmair, Sambuddha Ghosh, Hsueh-ling Huynh, Sanghoon Kim, Liisa Laine, Ting Liu, Michael Manove, Wendy Xu, and Pei-Cheng Yu for their comments and discussions. For their comments, I thank participants of the Third Workshop of the Advanced Pharmco & Health Economics Study and Research Center, in Genova, June 2023, Annual Conference of the Taiwan Society of Health Economics, in Taipei, December 2023, and the International Industrial Organization Conference, in Boston, May 2024. An editor and two referees provided insightful suggestions. Part of this research was undertaken during my visits at the University of New South Wales, Sydney, Australia, and at the National Taiwan University, Taipei, Taiwan; I am grateful for the host universities' hospitality and support.

1 Introduction

I study consumers' vaccination decisions. Individual choices together result in what I call a Vaccination Equilibrium, the sustainable percentage of consumers getting vaccinated. I compare the Vaccination Equilibrium mass with the first best, where an individual may be compelled to get vaccinated according to a rule that maximizes total welfare.

By taking a vaccine, a consumer benefits from illness severity reduction, but may have to incur a monetary cost, and suffer from personal disutility or side effects. These consumer private benefits and costs vary among individuals in the population. Besides private benefits and costs, consumers' vaccinations reduce the infection rate of a potentially contagious illness. This is an externality; a vaccine is also a public good.

A Vaccination Equilibrium results from individuals' private benefit-cost comparisons. Each consumer takes the population vaccination percentage as given, and decides whether to get vaccinated or not. In a Vaccination Equilibrium, the given population vaccination percentage must be the same as the total mass of consumers who choose to be vaccinated. It is a fixed point, the usual way an equilibrium is formally characterized. I show that there is a unique fixed point, a unique Vaccination Equilibrium.

I derive some comparative static results. A vaccine's improvement may increase benefit, decrease disutility, or reduce infection likelihood. I show how these three aspects of vaccine improvement may change the Vaccination Equilibrium. Higher benefit and reduced disutility will increase Vaccination Equilibrium; these are intuitive results from an individual's cost and benefit calculus. However, an improvement in infection likelihood will reduce Vaccination Equilibrium. This apparently counter-intuitive result has a natural resolution. Suppose that the illness infection likelihood is reduced. Now, an individual's vaccination benefit is realized with a lower probability, so would not like to suffer the disutility due to vaccination. Hence, each consumer has a reduced incentive to get vaccinated. As a result, the Vaccination Equilibrium decreases. This aspect of vaccine improvement is about the externality, not individual considerations.

I let social welfare be given by a utilitarian function. The first best is an allocation—which consumer is to be vaccinated—that maximizes social welfare. The first-best vaccination rate, of course, must balance consumers' benefits and costs, but there is an additional consideration. An increase in vaccination mass reduces the infection rate, so affects all consumers. The first best incorporates this externality, which is ignored by consumers. As a result, a Vaccination Equilibrium is never first best. Taxes and subsidies may change the Vaccination Equilibrium. Such Pigouvian monetary instruments simply alter private benefits and disutilities. With unrestricted use of taxes and subsidies, the Vaccination Equilibrium may be made to become first best. There is only the externality to correct, and a single tax-subsidy instrument is sufficient. A vaccine mandate, which means levying a penalty to consumers who refuse to get vaccinated, may also achieve the same.

Obviously, the recent pandemic has heightened attention to vaccines.[1] The US Center for Disease Control and Prevention and the Federal Drug Administration essentially regulate vaccine availability in the United States. Also, CDC publishes online vaccination schedules, guidelines for health care professionals, and advice for consumers.[2] The long list of available vaccines cover illnesses that are more or less contagious, and those that are seasonal.[3] However, the CDC and the FDA do not have the authority to compel individuals to be vaccinated. A Vaccination Equilibrium is one that is sustainable, from consumers' uncoordinated vaccination decisions. The model here is an interpretation of the current institution.

Vaccination Equilibrium offers a long-run perspective. The Susceptibility-Infection-Recovery model has been the classic framework for epidemiological studies of infectious diseases for over a century. The earliest pieces by Ross (1916, 1917) and Ross and Hudson (1917) laid out the mathematical foundation. The economic adaptation of SIR naturally includes behavioral responses; see Dasaratha (2023) for a very recent advance. However, the focus of SIR has been about the spread of diseases and its time path as described by differential equations; the economic perspective includes how human behaviors would change the dynamics. Hence, the emphasis is on the short run. The model here offers alternative, long-run and seasonal perspectives.

Brito, Sheshinski, and Intriligator (1991) examine competitive equilibria and their welfare properties. Their model has a continuum of consumers, each of whom may suffer an illness. A consumer may avoid the

[1] However, the first vaccine was reportedly administered in 1796 for smallpox. See https://www.who.int/health-topics/vaccines-and-immunization#tab=tab_1

[2] See https://www.cdc.gov/vaccines/

[3] https://www.cdc.gov/vaccines/schedules/

illness by taking a vaccine, but taking the vaccine may be personally costly. They point out that universal vaccination is suboptimal; in fact, the utilitarian welfare optimum does not prescribe that. Taxes and subsidies can eliminate the discrepancy between the competitive equilibria and social optimum.

In my model, consumers experience variable benefits and costs from taking the vaccine. Hence, I consider multi-dimensional effects of vaccination. The Vaccination Equilibrium may be below or above the social optimum. I provide comparative static on how changes in vaccine characteristics affect the Vaccination Equilibrium. I also let vaccines be imperfect. Vaccination then produces externalities for both vaccinated and unvaccinated consumers. All these characteristics here are absent in Brito, Sheshinski, and Intriligator (1991).

Several papers study strategic issues when a small or a finite number of consumers simultaneously make vaccination decisions. Xu (1999) studies effects of vaccine innovations. As in Brito, Sheshinski and Intriligator (1991), each consumer incurs a variable and private cost to get vaccinated. A vaccine may eliminate illness disutility, but only succeeds with some probability. Vaccine improvement is defined by an increase in the effectiveness probability. Xu (1999), however, assumes that there are only a finite number of consumers. The complex consumer strategic responses do not yield a monotone effect of vaccine improvement on vaccination rate. By contrast, I separate vaccine improvement into benefit enhancement, side-effect reduction, and infection-likelihood reduction. These different aspects allow me to draw unambiguous conclusions.

Heal and Kunreuther (2005) is an unpublished paper which continues with the study of a strategic game of a finite number of players. Vaccines are assumed to be fully effective. One consumer's vaccination may have a strong effect on the infection environment among the remaining consumers (simply because there are only a few consumers around). Hence, vaccination decisions carry both strategic effects and externalities. Following the same small-finite consumer methodology, Heal and Kunreuther claimed to have shown that under certain conditions, every possible vaccination configuration among a finite population of consumers may be an equilibrium outcome.

Recently, Sorensen (2023) expanded on Heal and Kunreuther (2005) with an analysis on welfare effect when vaccines are not fully effective. The paper shows that an improvement in vaccine effectiveness can reduce vaccination rate, and welfare may actually decrease. In the model, consumers incur a constant vaccine cost and obtain a constant benefit from the vaccine. The study focuses on the infection probability and the associated infection probability, either in an equilibrium or in a social optimum.

My model here views vaccines as a good with many attributes. A vaccine benefits a consumer by reducing illness severity, and is costly for a consumer, in monetary and nonmonetary terms. A vaccine also produces externalities in that population infection probability depends on the vaccination rate. However, the benefit and cost vary across the population consumers. I do not view externality as the only issue in the determination of an equilibrium; personal cost and benefit are also relevant dimensions. Notably, I view benefits and costs as random and infection probability as a continuous function of vaccination rates. A consumer's decision to get vaccinated is the only binary aspect in the model.

I have adopted a full-rationality assumption; consumers are not misinformed about vaccine benefits and cost, and they do free ride on the inoculation effect of others. Given this assumption, I feel that a utilitarian social welfare index as a benchmark is not unconvincing. This index can, of course, be changed to others. Manski (2010, 2017) studied partial information and how such welfare criteria as min-max, min-max-regret will guide policies; these criteria are motivated by consumers lacking information or not acting as benefit-cost calculus prescribes.

I am agnostic about whether consumers actually require cognitive assistance or guidance about vaccines. My purpose here is not to offer optimal policies, but provide a language for the sustainable vaccination rate, in a rich model. I am not aware of papers that address private benefits and costs, and infection probability externalities together.

The next section is the model description. Then I lay out the Vaccination Equilibrium in Section 3. The first best is derived in Section 4. I compare the Vaccination Equilibrium and first best in Section 5. I then offer some concluding remarks.

2 Model

There is a set of consumers, with total mass normalized at 1. Each may suffer from an infection with some probability. When ill, a consumer suffers a utility loss ℓ. This loss varies among consumers, so I let it be a random variable. There is a vaccine that can reduce the illness loss. If the vaccine is taken, an infected consumer's loss ℓ is reduced to $\theta\ell$, where $0 \leq \theta \leq 1$, and where θ can also be random. I call the reduction of illness loss upon vaccination, $\ell - \theta\ell$, the vaccine benefit $\beta \equiv \ell - \theta\ell$. Because ℓ and θ are random, so is β. Hence I let the benefit β be a random variable with a continuous distribution F on a positive support $[\underline{\beta}, \overline{\beta}]$.

A consumer suffers from a disutility when getting vaccinated. This disutility can be interpreted as side effects, anticipated or otherwise unforeseen, and any psychological discomfort from using medicine; this may also include time cost and inconvenience. Different consumers may experience the disutility differently, so I let the disutility δ be a random variable with a continuous distribution G on a positive support $[\underline{\delta}, \overline{\delta}]$. A consumer is identified by the vaccination benefit and the disutility, namely (β, δ). The F and G distributions are assumed to be independent; for simplicity, they also are assumed to be absolutely continuous, and so differentiable almost everywhere.

I assume that the chance of a consumer becoming infected depends negatively on the total mass, M, of consumers who have been vaccinated. The infection probability function is denoted by $P : [0, 1] \to [0, 1]$. The infection probability $P(M)$ is continuous, and decreasing. For certain diseases, contagion may be insignificant and P may be a constant function; more discussion on this possibility will follow. If vaccination does reduce infection likelihood differently for vaccinated and unvaccinated individuals, I will just have two infection probability functions, one for the vaccinated and another for the unvaccinated. Then I will just adopt appropriate adjustments in their infection probabilities or illness reduction to account for vaccination beneficial effects. To economize on notation, I will simply assume that infection probability is the same for vaccinated and unvaccinated consumers. Another interpretation is that the vaccine benefit β can be redefined to include the reduced infection likelihood upon vaccination.

The vaccination cost per consumer is normalized at 1; the cost of vaccine development is a sunk cost and ignored. I assume that consumers may be subsidized or covered by health insurance. I let τ be a consumer's out-of-pocket expense for getting vaccinated. In the case of subsidy, τ may be negative.

Suppose that the mass of vaccinated consumers is M. The expected utility of a consumer with wealth W forgoing vaccination is defined as $U(W) - \ell \times P(M)$, where U is a continuous and increasing utility function. I will take up a vaccine mandate in the last section; in this setup, a mandate will correspond to a penalty if a consumer refuses vaccination. Wealth or income levels do vary among individuals. But I will brush aside wealth variations, so assume that W is a constant. In Subsection 5.2, I will take up the issue of wealth variations. A consumer gets infected with probability $P(M)$, and then suffers the loss ℓ if infected.

If a consumer gets vaccinated, the expected utility is defined as $U(W - \tau) - \theta\ell \times P(M) - \delta$. Here, the loss has been reduced to $\theta\ell$ but the consumer experiences the disutility δ. Consumer (β, δ) chooses to get vaccinated if that is the better choice than no vaccination:

$$U(W - \tau) - \theta\ell \times P(M) - \delta > U(W) - \ell \times P(M),$$

which simplifies to

$$\begin{aligned} U(W - \tau) + [\ell - \theta\ell] \times P(M) - \delta &> U(W), \\ U(W - \tau) - U(W) + \beta P(M) &> \delta. \end{aligned} \tag{1}$$

This inequality says that expected benefit from getting vaccinated, net of out-of-pocket expense, is higher than the disuility.

Next, define a function $\widehat{\delta} : [\underline{\beta}, \overline{\beta}] \times [0, 1] \to [\underline{\delta}, \overline{\delta}]$ by

$$\widehat{\delta}(\beta; M) = U(W - \tau) - U(W) + \beta P(M),$$

if, at β, there is a $\delta \in [\underline{\delta}, \overline{\delta}]$ to satisfy the equation; otherwise, $\widehat{\delta}(\beta; M)$ will simply take a value equal to either $\underline{\delta}$ or $\overline{\delta}$, according to whether vaccination benefit is higher or lower than vaccination disutility.[4] A consumer with benefit β facing a total vaccinated mass of M would get vaccinated if the disutility is below the threshold $\widehat{\delta}(\beta; M)$.

Given the mass of vaccinated consumers M, consumer (β, δ) chooses to get vaccinated if disutility δ is

[4] Alternatively, extend the support of δ but assign zero density to those values outside of $[\underline{\delta}, \overline{\delta}]$.

less than the threshold $\widehat{\delta}(\beta; M)$. Therefore, the set of consumers who choose to get vaccinated is given by

$$\int_{\underline{\beta}}^{\overline{\beta}} \int_{\underline{\delta}}^{\widehat{\delta}(\beta;M)} \mathrm{d}G(\delta)\mathrm{d}F(\beta) = \int_{\underline{\beta}}^{\overline{\beta}} G(\widehat{\delta}(\beta; M))\mathrm{d}F(\beta).$$

3 Vaccination Equilibrium

The primitives of the model are the distributions F and G, respectively those of vaccination benefits and disutilities, and the infection function P, as well as the consumer's utility function U and wealth W. Now I introduce an equilibrium concept:

Definition 1 (Vaccination Equilibrium) *The mass of vaccinated consumers M is said to be a Vaccination Equilibrium if*

$$\int_{\underline{\beta}}^{\overline{\beta}} \int_{\underline{\delta}}^{\widehat{\delta}(\beta;M)} dG(\delta)\,dF(\beta) = \int_{\underline{\beta}}^{\overline{\beta}} G(\widehat{\delta}(\beta; M))\,dF(\beta) = M, \quad \text{where } \widehat{\delta}(\beta; M) = U(W-\tau) - U(W) + \beta P(M) \quad (2)$$

or

$$\int_{\underline{\beta}}^{\overline{\beta}} G(U(W-\tau) - U(W) + \beta P(M))dF(\beta) = M. \quad (3)$$

That is, if each consumer expects that the total mass of vaccinated consumers is M, a total of M consumers will find it optimal to get vaccinated.

A Vaccination Equilibrium captures the idea of sustainability. The anticipation of M vaccinated consumers is self-fulfilling. This definition also can be stated as a fixed point. Any arbitrary value of M is mapped to the value of the left-hand side of the above equilibrium condition (3), and the value of this map at M must remain at M at the fixed point. Define $\phi : [0, 1] \to [0, 1]$ by

$$\phi(M) = \int_{\underline{\beta}}^{\overline{\beta}} G(U(W-\tau) - U(W) + \beta P(M))\mathrm{d}F(\beta).$$

A Vaccination Equilibrium is a fixed point of ϕ, say $\widehat{M}$, such that $\phi(\widehat{M}) = \widehat{M}$.

Proposition 1 *There is a unique Vaccination Equilibrium. In other words, ϕ has a unique fixed point $\widehat{M}$.*

Proof of Proposition 1: The function

$$\phi(M) = \int_{\underline{\beta}}^{\overline{\beta}} G(U(W-\tau) - U(W) + \beta P(M))\mathrm{d}F(\beta)$$

is continuous because, U, P, F, and G are continuous. Also, it is decreasing in M because P is decreasing and G is increasing. The range of ϕ is a subset in $[0, 1]$.

The identity map of $[0, 1]$ onto itself is obviously continuous, increasing, and has a range $[0, 1]$. The decreasing function $\phi(M)$ therefore intersects the increasing identity map exactly once. There is a unique $\widehat{M}$, where $\phi(\widehat{M}) = \widehat{M}$, which is the unique Vaccination Equilibrium. ■

Vaccination Equilibrium being a fixed point imbeds a notion of "rational expectation." Consumers are not systematically biased about what they think about the mass of vaccinated consumers. Perhaps, some "learning" is required to establish the equilibrium. Nevertheless, Vaccination Equilibrium is a useful benchmark because it describes what can be sustained in the long run. In the case of seasonal contagious illnesses, Vaccination Equilibrium is a compelling index to describe a long-run tendency.

The process of arriving at a Vaccination Equilibrium can be envisioned as follows. Start with $M = 0$ so no one in the population is vaccinated. The likelihood of infection is the highest. Some consumers with high benefit β and low disutility δ will find it optimal to get vaccinated. This shows that $M = 0$ is not a Vaccination Equilibrium. As the value of M increases, the benefit drops due to a reduction of infection probability. Fewer individuals will want to get vaccinated. The process continues until the Vaccination Equilibrium is reached. The opposite argument can be applied to the case of $M = 1$. In this case, the infection probability would be the lowest. Those consumers with low benefit and high disutility would refuse vaccination. Hence $M = 1$ is not a Vaccination Equilibrium.

The above arguments point to a dynamic process, much like one in a standard demand and supply framework. If consumers think that the current mass of vaccinated consumers is lower than the Vaccination Equilibrium, more will choose to be vaccinated. The opposite is true if consumers believe that the current mass is higher. The exact dynamics over time is beyond the scope of the paper.[5]

[5] It seems entirely possible for a time path to exhibit a cobweb cycle, just as in basic demand and supply models.

3.1 Vaccine effectiveness

How does vaccine effectiveness affect the Vaccination Equilibrium? I propose three notions of effectiveness, each focusing on one dimension of a vaccine. I then adopt a first-order dominance change in the relevant distribution or function to capture effectiveness. Consider two vaccines A and B. Let their benefit distributions, respectively, be F_A and F_B, their disutility distributions be G_A and G_B, and their infection probability functions be P_A and P_B.

Three notions of Vaccine Effectiveness are defined as follows.

1. Vaccine A is said to be more effective in benefit than vaccine B if $F_A(\beta) \leq F_B(\beta)$ for each $\beta \in [\underline{\beta}, \overline{\beta}]$, with strict inequality for some β.

2. Vaccine A is said to be more effective in disutility than vaccine B if $G_A(\delta) \geq G_B(\delta)$ for each $\delta \in [\underline{\delta}, \overline{\delta}]$, with strict inequality for some δ.

3. Vaccine A is said to be more effective in infection prevention than vaccine B if $P_A(M) \leq P_B(M)$ for each $M \in [0, 1]$, with strict inequality for some M.

A vaccine is more effective in benefit than another if its benefit distribution puts more weights on high benefit realizations. Analogously, it is more effective in disutility if it puts more weights on low disutility realizations. Finally, a vaccine is more effective in infection prevention if it yields a lower infection probability for each vaccination level among consumers. These effectiveness concepts use first-order dominance, so pairs of vaccines may not be always comparable. First order dominance is arguably the most basic comparison between distributions.

Corollary 1 *A vaccine that is more effective in benefit, disutility, or both will increase Vaccination Equilibrium. A vaccine that is more effective in infection prevention will not increase Vaccination Equilibrium, and may decrease it.*

Proof of Corollary 1: Vaccination Equilibria, M_A and M_B, from vaccine A and vaccine B are, respectively, defined by

$$\int_{\underline{\beta}}^{\overline{\beta}} G_A(U(W-\tau) - U(W) + \beta P_A(M_A))\mathrm{d}F_A(\beta) \quad = \quad M_A$$

$$\int_{\underline{\beta}}^{\overline{\beta}} G_B(U(W-\tau) - U(W) + \beta P_B(M_B))\mathrm{d}F_B(\beta) \quad = \quad M_B.$$

Suppose that vaccine A is more effective in benefit than vaccine B. The claim is that $M_A > M_B$. Suppose that it is not true. That is, suppose that $M_A \leq M_B$, so

$$0 \geq M_A - M_B$$
$$= \int_{\underline{\beta}}^{\overline{\beta}} G(U(W-\tau) - U(W) + \beta P(M_A))\mathrm{d}F_A(\beta) - \int_{\underline{\beta}}^{\overline{\beta}} G(U(W-\tau) - U(W) + \beta P(M_B))\mathrm{d}F_B(\beta). \tag{4}$$

If $M_A \leq M_B$, we have $P(M_A) \geq P(M_B)$ because the infection probability is decreasing. Hence, at each β, $G(U(W-\tau) - U(W) + \beta P(M_A)) \geq G(U(W-\tau) - U(W) + \beta P(M_B))$. Given that $F_A(\beta) \leq F_B(\beta)$ for each $\beta \in [\underline{\beta}, \overline{\beta}]$, with strict inequality for some β, the expression in (4) is strictly positive. This is a contradiction. The claim $M_A > M_B$ is valid.

Next, suppose that vaccine A is more effective in disutility than vaccine B. The claim is that $M_A > M_B$. Suppose that it is not true. That is, suppose that $M_A \leq M_B$, so

$$0 \geq M_A - M_B$$
$$= \int_{\underline{\beta}}^{\overline{\beta}} G_A(U(W-\tau) - U(W) + \beta P(M_A))\mathrm{d}F(\beta) - \int_{\underline{\beta}}^{\overline{\beta}} G_B(U(W-\tau) - U(W) + \beta P(M_B))\mathrm{d}F(\beta). \tag{5}$$

If $M_A \leq M_B$, we have $P(M_A) \geq P(M_B)$ because the infection probability is decreasing. Hence, at each β, $G_A(U(W-\tau) - U(W) + \beta P(M_A)) \geq G_A(U(W-\tau) - U(W) + \beta P(M_B))$. Given that $G_A(\delta) \geq G_B(\delta)$ for each $\delta \in [\underline{\delta}, \overline{\delta}]$, with strict inequality for some δ, the expression in (5) is strictly positive. This is a contradiction. The claim $M_A > M_B$ is valid.

Finally, suppose that vaccine A is more effective in infection prevention than vaccine B. The claim is that $M_A \leq M_B$. Suppose that it is not true. That is, suppose that $M_A > M_B$, so

$$0 < M_A - M_B$$
$$= \int_{\underline{\beta}}^{\overline{\beta}} G(U(W-\tau) - U(W) + \beta P_A(M_A))\mathrm{d}F(\beta) - \int_{\underline{\beta}}^{\overline{\beta}} G(U(W-\tau) - U(W) + \beta P_B(M_B))\mathrm{d}F(\beta). \tag{6}$$

If $M_A > M_B$, we have $P_A(M_A) < P_A(M_B)$ because the infection probability is decreasing. Hence, at each β, $G(U(W-\tau) - U(W) + \beta P_A(M_A)) < G(U(W-\tau) - U(W) + \beta P_A(M_B))$. Given that $P_A(M) \leq P_B(M)$ for each $M \in [0,1]$, with strict inequality for some M, the expression in (6) is negative. This is a contradiction. The claim that $M_A \leq M_B$ is valid. ■

Results in Corollary 1 are interpreted as follows. A vaccine that yields more benefits or less disutilities will attract all consumers, given any infection likelihood. The mass of vaccinated consumers must increase. Now when a vaccine is more effective in reducing infection, consumers with high disutilities or low benefits will not bother with vaccination. Thus, the mass of consumers choosing vaccination actually decreases. This last observation has caused some controversy in the literature because more effective vaccines are supposed to be more attractive, so should not be rejected by consumers more often. The model here is more general because vaccines have multiple attributes. Those that reduce infection should be expected to reduce consumption, because benefits are less likely to be relevant and vaccination disutility can be avoided. Only those attributes that contribute to higher benefits and lower disutilities should be expected to raise consumption. Corollary 1 clarifies earlier contributions.

The Vaccination Equilibrium $\widehat{M}$ depends on the vaccine price τ. The following result is immediate: the vaccination price can implement any Vaccination Equilibrium level.

Corollary 2 *Suppose that each consumer pays τ to get vaccinated, where τ can be negative or positive. For any fixed M there exist a τ such that M becomes a Vaccination Equilibrium.*

Proof of Corollary 2: Consider

$$\int_{\underline{\beta}}^{\overline{\beta}} G(U(W-\tau) - U(W) + \beta P(M))\mathrm{d}F(\beta) = M,$$

which defines a function of M in terms of τ implicitly. By straightforward total differentiation, we have

$$\begin{aligned} & -\int_{\underline{\beta}}^{\overline{\beta}} g(U(W-\tau) - U(W) + \beta P(M))\mathrm{d}F(\beta) \times U'(W-\tau)\mathrm{d}\tau \\ = & \left\{ \int_{\underline{\beta}}^{\overline{\beta}} - g(U(W-\tau) - U(W) + \beta P(M)) \times \beta P'(M)\mathrm{d}F(\beta) + 1 \right\} \mathrm{d}M, \end{aligned}$$

where g is the derivative of G. Because P is decreasing, the right-hand side is strictly positive. It follows that M and τ are strictly negatively related. For any fixed M, there is a value of τ such that M is the fixed point. ■

A subsidy, $\tau < 0$, encourages vaccination; conversely, a tax, $\tau > 0$, discourages it. If the utility function is strictly increasing, there is always some subsidy or tax that can make the Vaccination Equilibrium equal to some pre-determined level. Corollary 2 is not meant to be a positive or practical result. It demonstrates mathematically that the consumer's vaccination cost or penalty is strictly related to the vaccination rate. The simple-minded implication is that a Pigouvian tax or subsidy can be readily used to induce various vaccination rates. However, the penalty or bonus may not be socially acceptable. A very high vaccination rate may require an exorbitant payment to convince consumers to get vaccinated. I have not involved a social cost-benefit calculus to assess how much a bonus is needed for a vaccination target. Factors outside of conventional economic considerations may present significant (and binding) limits on carrots and sticks.

4 First best

Now I analyze the first best, or the efficient vaccination rate. First, I define a vaccination assignment; this is a function, $\alpha : [\underline{\beta}, \overline{\beta}] \times [\underline{\delta}, \overline{\delta}] \to [0, 1]$, which says that consumer (β, δ) should be vaccinated with probability $0 \leq \alpha(\beta, \delta) \leq 1$. Given assignment α, the total mass of vaccinated consumers is

$$\int_{\underline{\beta}}^{\overline{\beta}} \int_{\underline{\delta}}^{\overline{\delta}} \alpha(\beta, \delta) \mathrm{d}G(\delta) \mathrm{d}F(\beta) \equiv M(\alpha),$$

where $M(\alpha)$, the mass of consumers under assignment α, is understood to be a notation different from the mass of vaccinated consumers in the previous section.

The first-best perspective is a utilitarian benefit-cost comparison. The consumer's utility function, wealth, and payment are ignored. Under assignment α, if consumer (β, δ) is vaccinated, the payoff is the expected reduced severity $\theta\ell \times P(M(\alpha))$ less the vaccination disutility δ. If consumer (β, δ) remains unvaccinated, there is no illness severity reduction but the vaccination disutility δ is avoided. The assignment $\alpha(\beta, \delta)$ puts probabilities on these payoffs:

$$\begin{aligned}
& \alpha(\beta,\delta)[-\theta\ell \times P(M(\alpha)) - \delta] + [1-\alpha(\beta,\delta)][-\ell \times P(M(\alpha))] \\
= \; & \alpha(\beta,\delta)[(\ell - \theta\ell)P(M(\alpha)) - \delta] - \ell P(M(\alpha)) \\
= \; & \alpha(\beta,\delta)[\beta P(M(\alpha)) - \delta] - \ell P(M(\alpha)).
\end{aligned}$$

The social welfare from assignment α is defined as the sum:

$$\int_{\underline{\beta}}^{\overline{\beta}} \int_{\underline{\delta}}^{\overline{\delta}} \{\alpha(\beta,\delta)[\beta P(M(\alpha)) - \delta]\} \, \mathrm{d}G(\delta)\mathrm{d}F(\beta) - \mu P(M(\alpha)) - M(\alpha),$$

where μ is the mean of the loss ℓ, and vaccine cost is normalized at 1. The integral aggregates the private benefits from vaccination; it is the sum of individual benefits and disutility, $[\beta P(M(\alpha)) - \delta]$. Now, this part of the social welfare function is similar to an individual consumer's consideration. But the term $\mu P(M(\alpha))$ is the expected loss for the entire population when the infection probability is $P(M(\alpha))$; this comes from summing over all consumers, whether they are vaccinated or not, which is the externality aspect of vaccination. The social component $\mu P(M(\alpha))$ is not internalized by any consumer, and is the main difference between the first best and the Vaccination Equilibrium.

Definition 2 (First Best) *The first-best vaccination assignment is:*

$$\alpha^*(\beta,\delta) \equiv \underset{\alpha(\beta,\delta)}{\operatorname{argmax}} \int_{\underline{\beta}}^{\overline{\beta}} \int_{\underline{\delta}}^{\overline{\delta}} \{\alpha(\beta,\delta)[\beta P(M(\alpha)) - \delta]\} \, dG(\delta) dF(\beta) - \mu P(M(\alpha)) - M(\alpha),$$

where

$$\int_{\underline{\beta}}^{\overline{\beta}} \int_{\underline{\delta}}^{\overline{\delta}} \alpha(\beta,\delta) dG(\delta) dF(\beta) \equiv M(\alpha).$$

I characterize the first best with a two-step procedure:

1. For a given vaccination mass, find the optimal assignment to maximize the objective function with $M(\alpha)$ set at, say, $\overline{M}$.

2. Choose the optimal level of $\overline{M}$ given the optimal assignment conditional on $\overline{M}$.

For a given $\overline{M}$, $0 \leq \overline{M} \leq 1$, define

$$V(\overline{M}) - \overline{M} \equiv \max_{\alpha(\beta,\delta)} \left\{ \int_{\underline{\beta}}^{\overline{\beta}} \int_{\underline{\delta}}^{\overline{\delta}} \{\alpha(\beta,\delta)[\beta P(\overline{M}) - \delta]\} \, \mathrm{d}G(\delta)\mathrm{d}F(\beta) - \mu P(\overline{M}) \right\} - \overline{M}, \tag{7}$$

subject to

$$\int_{\underline{\beta}}^{\overline{\beta}} \int_{\underline{\delta}}^{\overline{\delta}} \alpha(\beta, \delta) \mathrm{d}G(\delta) \mathrm{d}F(\beta) = \overline{M}. \tag{8}$$

This is the indirect social welfare function under the (conditional) optimal vaccination assignment for a given vaccination mass $\overline{M}$.

Lemma 1 *At every $\overline{M}$, the (conditional) optimal assignment satisfies: $\alpha(\beta, \delta) = 1$ if $[\beta P(\overline{M}) - \delta] > \lambda(\overline{M})$, for some function λ that depends on $\overline{M}$.*

Proof of Lemma 1: Use pointwise optimization with respect to α at (β, δ). The Lagrangian is

$$\left\{\alpha(\beta, \delta)[\beta P(\overline{M}) - \delta]\right\} - \mu P(\overline{M}) - \overline{M} - \lambda \left[\alpha(\beta, \delta) - \overline{M}\right],$$

where λ is the multiplier of the constraint (8). The first-order derivative with respect to α is

$$[\beta P(\overline{M}) - \delta] - \lambda$$

for a nonnegative multiplier λ. The derivative is independent of α, so it is optimal to set $\alpha(\beta, \delta)$ to 1 when the above expression is positive:

$$\alpha(\beta, \delta) = 1 \qquad \Longleftrightarrow \qquad [\beta P(\overline{M}) - \delta] - \lambda \geq 0,$$

which is the expression in the statement of the Lemma.■

From (1) and Lemma 1, the socially optimal vaccination assignment is similar to one that would have been chosen by consumers themselves. This stems from comparing the benefit net of the disutility and the shadow price due to the limited vaccination mass. Hence, a consumer with payoff of an expected benefit less disutility above a threshold should get vaccinated: formally consumer (β, δ) should be vaccinated if $[\beta P(\overline{M}) - \delta] \geq \lambda$, which is the shadow price of the given vaccination mass $\overline{M}$.

Using the optimal assignment in Lemma 1, I rewrite the indirect social benefit in (7) by

$$V(\overline{M}) - \overline{M} = \int_{\underline{\beta}}^{\overline{\beta}} \int_{\underline{\delta}}^{\widehat{\delta}} [\beta P(\overline{M}) - \delta] \mathrm{d}G(\delta) \mathrm{d}F(\beta) - \mu P(\overline{M}) - \overline{M}, \quad \text{where } \widehat{\delta} = \beta P(\overline{M}) - \lambda(\overline{M}). \tag{9}$$

I assume that the indirect social welfare function $V(\overline{M}) - \overline{M}$ is concave, and that at very small values of $\overline{M}$, it is increasing, and eventually, it turns negative. The optimal vaccination mass is one where the marginal (indirect) benefit $V'(\overline{M})$ is equal to the marginal cost, which is 1.

Changing $\overline{M}$ has both direct and indirect effects. First, the direct effect is reflected in a partial derivative of $\overline{M}$ in the integrand of $V(\overline{M})$ and in the additive term $\mu P(\overline{M})$ in (9); this is the externality due to vaccination reducing infection probability for all consumers. Second, the indirect effect works through the optimal assignment, which says that $\alpha(\beta,\delta)=1$ if and only if $\beta P(\overline{M})-\delta \geq \lambda(\overline{M})$. By the envelope theorem, that indirect effect from a change of $\overline{M}$ in (9) is zero because the assignment is already optimal.

Let M^* be the socially optimal vaccination mass: $M^* = \text{argmax}_{\overline{M}}[V(\overline{M}) - \overline{M}]$. At $\overline{M} = M^*$ the multiplier λ vanishes because the constraint (8) does not bind at the optimal vaccination mass. Hence

$$V(M^*) = \int_{\underline{\beta}}^{\overline{\beta}} \int_{\underline{\delta}}^{\widehat{\delta}} [\beta P(M^*) - \delta] \mathrm{d}G(\delta) \mathrm{d}F(\beta) - \mu P(M^*) \quad \text{with } \widehat{\delta} = \beta P(M^*).$$

Hence, at the socially optimal mass M^*, the derivative of V at M^* is

$$\begin{aligned} V'(M^*) &= \int_{\underline{\beta}}^{\overline{\beta}} \int_{\underline{\delta}}^{\widehat{\delta}} \beta P'(M^*) \mathrm{d}G(\delta) \mathrm{d}F(\beta) - \mu P'(M^*) \\ &= P'(M^*) \left\{ \int_{\underline{\beta}}^{\overline{\beta}} \beta G(\beta P(M^*)) \mathrm{d}F(\beta) - \mu \right\}, \end{aligned} \tag{10}$$

where the value of $\widehat{\delta}$ has been substituted by $\beta P(M^*)$, and where any indirect effect on the objective through $\widehat{\delta} = \beta P(M^*)$ has been ignored. The effect of a change in $\overline{M}$ is further separated into two changes. First, an increase in $\overline{M}$ reduces infection rate by $P'(M^*)$, which affects all consumers, so results in a welfare improvement of $-P'(M^*)\mu$. Second, an increase in $\overline{M}$ actually reduces the mass of consumers that should be vaccinated. The consumer (β,δ) who is just indifferent between getting vaccinated and not vaccinated is defined by $\delta = \beta P(M^*)$, or $\frac{\delta}{\beta} = P(M^*)$. If vaccination mass increases, the infection probability decreases so the ratio $\frac{\delta}{\beta}$ decreases. Fewer consumers get vaccinated, so there is cost saving, which is reflected in the integral of the $V'(M^*)$ expression in (10). I summarize the above in the following (its proof having been presented above).

Proposition 2 *The first-best vaccination mass of consumers M^* is given by*

$$P'(M^*) \left\{ \int_{\underline{\beta}}^{\overline{\beta}} \beta G(\beta P(M^*)) dF(\beta) - \mu \right\} = 1. \tag{11}$$

5 Comparing Vaccination Equilibrium and First Best

The efficient or optimal vaccination assignment $\alpha(\beta, \delta)$ in the first best employs a different kind of calculus than what consumers use to arrive at a Vaccination Equilibrium. For individual decisions, if loss is high and disutility is low, a consumer would choose to pay a fee to get vaccinated. Under social welfare maximization, if loss is high and disutility low, a consumer should get vaccinated. But the presence of externality makes the social and private perspectives differ.

5.1 With externality

In a Vaccination Equilibrium, each consumer takes vaccination mass as given, so the two effects of changing vaccination mass, described just before Proposition 2, are ignored. This is the source of inefficiency. There is no reason to expect that the socially efficient vaccination mass is a Vaccination Equilibrium at some exogenously given price τ.

Suppose that consumers do not have to pay for the vaccine. Consider the equation (3) for a Vaccination Equilibrium:

$$\int_{\underline{\beta}}^{\overline{\beta}} \int_{\underline{\delta}}^{\overline{\delta}} \mathrm{d}G(\delta)\mathrm{d}F(\beta) = M, \qquad \text{where } \widehat{\delta} = \beta P(M). \tag{12}$$

Suppose that $M = 0$ so no consumer is vaccinated. Consider consumer (β, δ); this consumer will choose to get vaccinated if $\delta < \beta P(0)$. Consumers who choose to be vaccinated form a positive mass, so the left hand side of (12) is strictly positive. Hence, it is not a Vaccination Equilibrium at a zero vaccination rate. Now, raising the value of M from zero will discourage some consumers from getting vaccinated: $P(M) < P(0)$, so the mass of consumers satisfying $\delta < \beta P(M)$ is smaller than those satisfying $\delta < \beta P(0)$. As M continues to increase from 0, the left-hand side of (12) continues to decrease, until (12) holds.

The point is this. The adjustment to the vaccination mass to arrive at a Vaccination Equilibrium is entirely through the reduction in the disutility threshold: $\widehat{\delta} = \beta P(M)$ decreases as $P(M)$ decreases. The effect is due to the changes in the marginal consumer (β, δ). This consideration is entirely driven by individual consumers' decisions. Imposing a fee (such as the vaccine cost) will change this calculus, but the vaccination effect on infection probability is only taken as given by consumers.

Next, the determination of the first-best vaccination rate is from the maximization of the social surplus:

$$V(M) = \int_{\underline{\beta}}^{\overline{\beta}} \int_{\underline{\delta}}^{\widehat{\delta}} [\beta P(M) - \delta] \mathrm{d}G(\delta)\mathrm{d}F(\beta) - \mu P(M), \qquad \text{where } \widehat{\delta} = \beta P(M).$$

There are two differences from the determination of the Vaccination Equilibrium. The first originates from the term in the integrand: $\beta P(M) - \delta$, which is the net gain for consumer (β, δ) from vaccination; the threshold $\widehat{\delta} = \beta P(M)$ ensures that the net gain is positive. The second originates from the term $-\mu P(M)$, which measures the expected illness disutility for the entire population.

Again, suppose that $M = 0$, then $P(0)$ is the maximum infection probability. As M rises from 0, the threshold decreases from $\beta P(0)$, so the total net gain, the integral of $\beta P(M) - \delta$, actually decreases. The marginal effect is measured by $\int_{\underline{\beta}}^{\overline{\beta}} \beta P'(M^*) G(\widehat{\delta}) \mathrm{d}F(\beta)$, the integral in (10), and this is negative. But an increase in vaccination rates reduces expected illness disutility, $-\mu P(M)$. The marginal effect is $-\mu P'(M)$, which is positive. The total marginal benefit due to vaccination is in (10). As M increases, the marginal benefit in (10) is compared to the marginal cost of vaccination, which is assumed to be 1.

Vaccination Equilibrium and first-best vaccination mass are arrived at via different considerations. Vaccination Equilibrium is a sort of accounting: the total mass of consumers who desire vaccination, given a vaccination rate, must be equal to that said rate. First-best vaccination is the balance of aggregate benefit and cost, which include the vaccine's effect on vaccinated and unvaccinated consumers. There is no reason to expect the Vaccination Equilibrium and the first best to be identical. However, Corollary 2 does say that there exists τ, a vaccination fee penalty or bonus such that the Vaccination Equilibrium is the same as first best. This, however, does not imply that consumers achieve the same utility because they have to pay the penalty or receive the bonus. The following example illustratives various results.

Example 1 *Uniform distributions and linear infection probabilities*

Suppose that the benefit distribution is uniform on $[\underline{\beta}, \overline{\beta}]$; the disutility distribution is uniform on $[\underline{\delta}, \overline{\delta}]$. Now the net benefit β is defined as $\beta = (1 - \theta)\ell$, so if I assume that θ is a constant, then ℓ is also uniform, and its mean is $\dfrac{\underline{\beta} + \overline{\beta}}{2(1 - \theta)} = \mu$. Suppose that the infection probability is $P(M) = 1 - kM$, where $0 < k \leq 1$. Assume now that consumers are risk neutral so that U is linear, and that each pays a vaccination fee τ, so now consumer (β, δ) prefers to be vaccinated when $\beta(1-kM)-\tau \geq \delta$. Define $\widehat{\delta} = \beta(1-kM)-\tau$. I assume that at the Vaccination Equilibrium $\widehat{M}$, for any β in its support, $\underline{\delta} \leq \widehat{\delta} \equiv \beta(1-kM)-\tau \leq \overline{\delta}$, so thresholds are inside the support $[\underline{\delta}, \overline{\delta}]$.

The Vaccination Equilibrium $\widehat{M}$ is calculated by substituting in (3) the uniform distributions for F and G, and linear function for P :

$$\begin{aligned} \int_{\underline{\beta}}^{\overline{\beta}} \int_{\underline{\delta}}^{\widehat{\delta}} \mathrm{d}G(\delta)\mathrm{d}F(\beta) &= \widehat{M}, \qquad \text{where } \widehat{\delta} = \beta(1-k\widehat{M})-\tau \\ \frac{\frac{1}{2}(\overline{\beta}+\underline{\beta})-\tau-\underline{\delta}}{\frac{1}{2}(\overline{\beta}+\underline{\beta})k+\overline{\delta}-\underline{\delta}} &= \widehat{M}. \end{aligned} \tag{13}$$

which is obtained by straightforward calculation using the uniform densities $(\overline{\beta}-\underline{\beta})^{-1}$ and $(\overline{\delta}-\underline{\delta})^{-1}$ on $F(\beta)$ and $G(\delta)$, respectively.[6] Corollary 1 can be verified. First, for benefit effectiveness increase, I consider raising $\overline{\beta}$, which corresponds to the benefit distribution putting more weights on higher values. It is readily verified from (13) that the derivative of $\widehat{M}$ with respect to $\overline{\beta}$ is positive. An increase in benefit effectiveness raises the Vaccination Equilibrium. Second, for disutility effectiveness increase, I consider reducing $\underline{\delta}$, which corresponds to the disutility distribution putting more weights on lower values. It is readily verified from (13) that the derivative of $\widehat{M}$ with respect to $\underline{\delta}$ is negative. An increase in disutility effectiveness raises the Vaccination Equilibrium. Third, the infection probability is $P(M) = 1-kM$, so as k increases, the vaccine becomes more effective in prevention. From (13), the Vaccination Equilibrium $\widehat{M}$ decreases as the vaccine becomes more effective in prevention.

Next, I use the expression (11) in Proposition 2 to calculate the first-best vaccination mass M^*:

$$P'(M^*)\left\{\int_{\underline{\beta}}^{\overline{\beta}} \int_{\underline{\delta}}^{\widehat{\delta}} \beta \mathrm{d}G(\delta)\mathrm{d}F(\beta) - \mu\right\} = 1, \qquad \text{where } \widehat{\delta} = \beta(1-M^*).$$

Using the uniform densities and substituting $P'(M)$ by $-k$, I simplify (11), and obtain the first-best vacci-

[6] The intermediate steps are $\int_0^1 \int_0^{\widehat{\delta}} \mathrm{d}G(\delta)\mathrm{d}F(\beta) = \frac{1}{(\overline{\beta}-\underline{\beta})(\overline{\delta}-\underline{\delta})}\int_0^1 [\beta(1-k\widehat{M})-\tau-\underline{\delta}]\mathrm{d}\beta$ $= \frac{1}{(\overline{\delta}-\underline{\delta})}\left[(1-k\widehat{M})\frac{(\overline{\beta}+\underline{\beta})}{2}-\tau-\underline{\delta}\right] = \widehat{M}$

nation mass:

$$\begin{aligned}-k\left\{\int_{\underline{\beta}}^{\overline{\beta}}\beta\left[\frac{\beta(1-kM^*)-\underline{\delta}}{\overline{\delta}-\underline{\delta}}\right]\frac{\mathrm{d}\beta}{\overline{\beta}-\underline{\beta}}-\mu\right\} &= 1\\ \left[1-k\mu-\frac{\frac{1}{2}(\overline{\beta}+\underline{\beta})\underline{\delta}k}{\overline{\delta}-\underline{\delta}}\right]+\frac{k\left(\overline{\beta}^2+\underline{\beta}\overline{\beta}+\underline{\beta}^2\right)}{3(\overline{\delta}-\underline{\delta})} &= \frac{\left(\overline{\beta}^2+\underline{\beta}\overline{\beta}+\underline{\beta}^2\right)k^2}{3(\overline{\delta}-\underline{\delta})}M^* \qquad (14)\end{aligned}$$

by straightforward calculation.[7] The Vaccination Equilibrium in (13) is linear in the vaccination fee τ. Hence, there is a fee τ such that $\widehat{M}$ in (13) is the same as M^*in (14).

Clearly, the Vaccination Equilibrium may be smaller or larger than the first best, according to the comparison between (13) and (14). The Vaccination Equilibrium depends on the out-of-pocket expense τ. The first-best vaccination rate depends on the expected loss μ. They are different parts of the model. There is no unambiguous comparison between these two vaccination rates. Precisely because the values of τ and μ are different, one can imagine changing one or the other so that the two vaccination rates become closer or even identical.

5.2 Without externality

In this subsection, I assume that infection probability is fixed, unaffected by the vaccination rate. Let the infection probability be constant at $\overline{P}$, between 0 and 1, independent of the mass of vaccinated consumers. For a given assignment $\alpha(\beta,\delta)$, the social objective function becomes

$$\begin{aligned}&\int_{\underline{\beta}}^{\overline{\beta}}\int_{\underline{\delta}}^{\overline{\delta}}\left\{\alpha(\beta,\delta)[-\theta\ell\times\overline{P}-\delta]+[1-\alpha(\beta,\delta)][-\ell\times\overline{P}]\right\}\mathrm{d}G(\delta)\mathrm{d}F(\beta)-\int_{\underline{\beta}}^{\overline{\beta}}\int_{\underline{\delta}}^{\overline{\delta}}\alpha(\beta,\delta)\mathrm{d}G(\delta)\mathrm{d}F(\beta)\\ =&\int_{\underline{\beta}}^{\overline{\beta}}\int_{\underline{\delta}}^{\overline{\delta}}\left\{\alpha(\beta,\delta)[\beta\overline{P}-\delta]-1\right\}\mathrm{d}G(\delta)\mathrm{d}F(\beta)-\mu\overline{P}.\end{aligned}$$

The optimal assignment is a straightforward application of Lemma 1.

Corollary 3 *If there is no externality so that the infection probability is a constant, say $\overline{P}$, the first-best assignment satisfies: $\alpha(\beta,\delta)=1$ if and only $\beta\overline{P}-\delta>1$.*

[7] The intermediate steps are

$$-k\left\{\frac{1}{\overline{\delta}-\underline{\delta}}\int_{\underline{\beta}}^{\overline{\beta}}\left[\frac{\beta^2(1-kM^*)-\beta\underline{\delta}}{\overline{\beta}-\underline{\beta}}\right]\mathrm{d}\beta-\mu\right\}=-k\left\{\frac{1-kM^*}{3(\overline{\delta}-\underline{\delta})}\left(\overline{\beta}^2+\underline{\beta}\overline{\beta}+\underline{\beta}^2\right)-\frac{\frac{1}{2}(\overline{\beta}+\underline{\beta})\underline{\delta}}{\overline{\delta}-\underline{\delta}}-\mu\right\}=1.$$

The first-best assignment remains a cost-benefit comparison, and does not consider consumer's utility function U or wealth W. However, given a fee τ, consumer (β, δ) chooses to get vaccinated if $U(W-\tau) + \beta\overline{P} - \delta > U(W)$, which simplifies to

$$\beta\overline{P} - \delta > U(W) - U(W - \tau).$$

This is a little different from the condition in Corollary 3. Wealth effect comes into play due to the curvature of the utility function U. Even if a consumer bears the full cost, say $\tau = 1$, the difference between $U(W)$ and $U(W-\tau)$ may not implement the first-best vaccination decision. When a consumer does not have to bear any cost, $\tau = 0$, the consumer's vaccination decision becomes more favorable than that in Corollary 3. By and large, except for wealth effect, individual and public objectives are very much aligned.

6 Conclusion

I study the sustainable outcome of individual consumers making decisions on vaccination. A vaccine gives benefits in terms of reduced illness severity, but it may result in side effects or disutilities. Benefits and disutilities vary across consumers. A vaccine also gives rise to an externality; an illness infection likelihood decreases when more individuals get vaccinated. A Vaccination Equilibrium is the vaccination rate that results from consumers' uncoordinated vaccination decisions. A single consumer cannot affect the total vaccination mass, so will not internalize the reduced-infection effect. A Vaccination Equilibrium is never first best precisely because of the lack of internalization.

Vaccine improvements may be in terms of more benefits, less side effects, or improved infection characteristics. I show that more benefits and less side effects raise the Vaccination Equilibrium, but improved infection characteristics reduce it. These are straightforward results, and tend to debunk earlier claims that vaccine improvements yield undesirable consumer responses.

I demonstrate how a first best can be calculated. The procedure takes into account the externality; it also incorporates consumers' benefit and cost comparisons. My framework and results cast doubt on the power of public-health policies. Public health authorities seldom consider consumers' benefit-cost calculus when constructing vaccine recommendations. The broad guidelines, if they aim for the first-best vaccination rate, are unlikely to match the Vaccination Equilibrium. An inducement for a higher vaccination rate must work through increased benefits or reduced disutilities.

An appropriately chosen vaccination fee can align the misaligned incentive. However, I have only offered an abstract, theoretical perspective; in practice, a fee or subsidy is not always practical. As of April 2024, the majority of vaccines are priced below US$100 per dose for CDC and private sector purchases. The most expensive at almost $500, the Respiratory Syncytial Virus vaccine, seems to be an outlier.[8] If fees cannot be so drastically different from purchases prices, then the range of financial incentives is limited.

Mandates are often used to enforce vaccination. Health care professionals often have to be vaccinated for employment. Children have to be vaccinated for attending schools. Such mandates impose either a monetary fine or a disutility on refusal to accept vaccination. However, mandates are not universal. Perhaps cultural and social responsibility, outside of the standard benefit-cost calculus, may be an alternative way to improve efficiency. A full analysis on how social responsibility plays into vaccination campaign may be interesting research. Convincing that individuals owe it to others to reduce infection may serve to change decisions.

[8] https://www.cdc.gov/vaccines/programs/vfc/awardees/vaccine-management/price-list/index.html

References

[1] Brito, Dagobert L., Eytan Sheshinski, and Michael D. Intriligator (1991), “Externalities and compulsory vaccinations,” Journal of Public Economics, 45, 69-90.

[2] Dasaratha, Krishna (2023), “Virus dynamics with behavioral responses,” Journal of Economic Theory, 214, 105739.

[3] Heal, Jeffrey and Howard Kunreuther (2005), “The vaccination game,” Risk Management and Decision Processes Center Working Paper.

[4] Manski, Charles F. (2010), “Vaccination with partial knowledge of external effectiveness,” Proceedings of the National Academy of Science,107, 3953-3960.

[5] Manski, Charles F. (2017), “Mandating vaccination with unknown indirect effects,” Journal of Public Economic Theory, 19, 603-619.

[6] Ross, Ronald (1916), “An application of the theory of probabilities to the study of a priori pathometry—Part I” Proceeding of the Royal Society A, https://doi.org/10.1098/rspa.1916.0007.

[7] Ross, Ronald (1917), “An application of the theory of probabilities to the study of a priori pathometry—Part III” Proceeding of the Royal Society A, https://doi.org/10.1098/rspa.1917.0015.

[8] Ross, Ronald, and Hilda P. Hudson (1917), “An application of the theory of probabilities to the study of a priori pathometry—Part II” Proceeding of the Royal Society A, https://doi.org/10.1098/rspa.1917.0014.

[9] Sorensen, Andrea (2023), “The role of vaccine effectiveness on individual vaccination decisions and welfare,” Journal of Public Economic Theory, 2023;1-17.

[10] Xu, Xiaopeng (1999) “Technological improvements in vaccine efficacy and individual incentive to vaccinate,” Economics Letters, 65, 359-364.