

NBER WORKING PAPER SERIES

VACCINATION PLANNING UNDER UNCERTAINTY, WITH APPLICATION TO
COVID-19

Charles F. Manski

Working Paper 28446
<http://www.nber.org/papers/w28446>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2021

I am grateful to Valentyn Litvin for excellent research assistance and to Michael Gmeiner for helpful comments. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Charles F. Manski. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Vaccination Planning under Uncertainty, with Application to Covid-19
Charles F. Manski
NBER Working Paper No. 28446
February 2021
JEL No. D81,H75,I18

ABSTRACT

Vaccination against infectious disease may be beneficial to reduce illness in vaccinated persons and disease transmission across the population. The welfare-economic practice of specifying a social welfare function and considering a planner who seeks to optimize welfare provides a constructive framework to evaluate vaccination policy. This paper characterizes choice of vaccination policy as a planning problem that aims to minimize the social cost of illness and vaccination. Manski (2010, 2017) studied vaccination as a problem of planning under uncertainty, assuming that a planner can choose any vaccination rate or that the planner has only two options: mandate or decentralize vaccination. The analysis focused on uncertainty regarding the effect of vaccination on disease transmission. Here I weaken the assumptions to recognize multiple uncertainties relevant to evaluation of policy for vaccination against COVID-19. These include uncertainty not only about the effect of vaccination on disease transmission, but also about the fraction of susceptible persons in the population, the effectiveness of vaccination in reducing illness and infectiousness, and the health risks associated with vaccination. The paper considers planning under ambiguity using the minimax and minimax-regret criteria, as well as planning using a subjective probability distribution on unknown quantities. It develops algorithms that may be applied flexibly to determine policy choices with specified degrees and types of uncertainty.

Charles F. Manski
Department of Economics
Northwestern University
2211 Campus Drive
Evanston, IL 60208-2600
and NBER
cfmanski@northwestern.edu

1. Introduction

Mitigation of infectious disease through vaccination is a core concern of health policy. Vaccination may be beneficial to reduce illness in vaccinated persons and disease transmission across the population.

In the American federal system, states make vaccination policy. States have legal power to mandate vaccination of specified groups against specific diseases. Immunization Action Coalition (2020) lists current state mandates. The development of COVID-19 vaccines has opened the possibility that states may mandate vaccination to reduce transmission of this novel disease.¹ The alternative to a mandate is decentralization, with vaccination decisions being made by families, schools, care facilities, and employers. In decentralized regimes, states may encourage vaccination without requiring it.

Mandatory vaccination has been a subject of controversy. Vamos, McDermott, and Daley (2008) give arguments for and against mandatory administration of the HPV vaccine to middle-school girls. Thomas (2009) and May and Silverman (2005) do likewise for mandatory vaccination of health care workers against influenza and children against multiple diseases. Reiss and Caplan (2020) discuss legal and ethical considerations regarding mandatory COVID-19 vaccination. These and other articles in the legal and health literatures acknowledge that society has conflicting objectives when deciding whether to mandate vaccination. However, they do not provide formal frameworks to help policy makers reconcile the conflicting objectives.

The welfare-economic practice of specifying a social welfare function and considering a planner who seeks to optimize welfare provides a constructive framework to evaluate vaccination policy. A broad objective is to minimize the social cost of illness and vaccination. Vaccination is beneficial to the extent that it prevents illness. However, vaccines are costly to produce and may have health side effects. Specifying a social cost function expresses quantitatively how society trades off the advantages and disadvantages of different policies.

¹ In September 2020, Wisconsin Governor Evers stated that requiring COVID-19 vaccination is the right thing to do, but he questioned if it is “feasible in this political environment” (Star Tribune, 2020). In November 2020, the New York State Bar Association recommended that the state consider mandating a COVID-19 vaccine once a scientific consensus emerges that it is safe, effective and necessary (New York State Bar Association, 2020).

This paper characterizes choice of vaccination policy as a planning problem that aims to minimize the social cost of illness and vaccination. Economists have usually studied vaccination as a deterministic planning problem, assuming that a planner knows the outcomes that alternative policies would yield. See, for example, Brito, Sheshinski, and Intriligator (1991), Francis (1997) and Chen and Toxvaerd (2014). In contrast, Manski (2010, 2017) studied vaccination as a problem of planning under uncertainty.

Manski (2010) assumed that a planner can choose any vaccination rate for a population. Manski (2017) assumed that the planner has only two options: mandate or decentralize vaccination. In both articles, the analysis focused on uncertainty regarding the effect of vaccination on disease transmission, assuming only that vaccination reduces transmission. With this weak assumption, planning is a problem of decision making under ambiguity. Planning was studied using the minimax and minimax-regret criteria. Explicit algebraic solutions were obtained.

Here I weaken the assumptions of Manski (2010, 2017) to recognize multiple uncertainties relevant to evaluation of policy for vaccination against COVID-19. These include uncertainty not only about the effect of vaccination on disease transmission, but also about the fraction of susceptible persons in the population, the effectiveness of vaccination in reducing illness and infectiousness, and the health risks associated with vaccination. As in the precedent work, the paper studies planning under ambiguity using the minimax and minimax-regret criteria. It also studies settings where the planner places a subjective probability distribution on unknown quantities and minimizes subjective expected social cost.

Algebraic solution of the planning problem is not feasible when multiple uncertainties are present, but computational solution is tractable. The paper develops algorithms that may be applied flexibly to express specified degrees and types of uncertainty. While the paper focuses on COVID-19 because of its immediate salience, the analytical framework developed here is applicable to other infectious diseases as well.

Section 2 specifies an optimization problem that a planner with complete information would be able to solve. Section 3 calls attention to sources of uncertainty that make it infeasible to solve this optimization problem. Section 4 considers how policy making might reasonably cope with these uncertainties.

1.1. Differentiating Epidemic Suppression from Social Cost Minimization

Before proceeding, I find it important to distinguish the welfare economic objective of minimization of social cost from the epidemiological objective of suppressing epidemics. Epidemiologists have used dynamic susceptible-infectious-removed (SIR) and related dynamic compartmental models of disease transmission to study epidemics and to evaluate policies that seek to prevent their onset or stop them once begun. An epidemic is a dynamic concept, defined to occur when the infected fraction of the population increases with time relative to some baseline (Centers for Disease Control and Prevention, 2006). Studies of vaccination policy performed by epidemiologists use SIR and related dynamic models to frame their research.²

Dynamic epidemiological modeling can help to evaluate vaccination policy, but such modeling has practical and conceptual drawbacks. A practical problem has been that the analytical and computational complexity of epidemiological models inhibits understanding of their mathematical properties and their realism. Researchers use models of disease transmission to forecast outcomes that would occur with various policies. However, authors generally provide little information that would enable one to assess the accuracy of the many assumptions their models place on individual behavior, social interactions, and disease transmission. Hence, it is prudent to view their forecasts as computational experiments rather than accurate predictions of policy impacts.

A conceptual drawback has been the prevailing focus of epidemiologists on the goal of suppressing epidemics rather than on the objective of minimizing the social cost of illness and vaccination. Epidemiologists often seek a policy that keeps the *effective reproduction rate*, a parameter governing the rate of disease transmission, below the rate at which a model forecasts that a new epidemic will begin or that an ongoing one will cease. An epidemic is forecast to cease when the fraction of the population who are immune,

² Some models applied to COVID-19 policy are described in Aguas *et al.* (2020), Ferguson *et al.* (2020), IHME (2020), and Paltiel *et al.* (2021). Some of the many instances of pre-covid research are Becker and Starczak (1997), Ball and Lyne (2002), Scuffham and West (2002), Hill and Longini (2003), Patel, Longini, and Halloran (2005), Boulier, Datta, and Goldfarb (2007), Althouse, Bergstrom, and Bergstrom (2010), and Keeling and Shattock (2012).

either through natural infection or vaccination, reaches the so-called *herd immunity threshold* (HIT). Formal definitions of the effective reproduction rate and the HIT vary with the epidemiological model studied. See Anderson and May (1991), Fine (1993), and Fine, Eames, and Heymann (2011).

The goal of suppressing epidemics may sometimes be congruent with the objective of minimizing social cost. However, epidemiological research does not show that achieving the former goal is necessary or sufficient for accomplishment of the latter objective.

1.2. Conceptualizing Choice of Vaccination Rate as a Static Optimization Problem

This paper builds on Manski (2010, 2017), weakening some assumptions to increase the relevance to COVID-19. As in these precedent papers, I abstract from the dynamics of infectious disease and study minimization of the social cost of illness and vaccination as a static optimization problem. I assume that vaccination occurs at a specified point in time, rather than as a process performed gradually over time. I assume that society is concerned with social cost over a specified horizon following vaccination. I study choice of an aggregate vaccination rate rather than disaggregated choice of whom to vaccinate. Differentiating itself from the earlier work, this paper simplifies study of the effect of vaccination on illness prevalence by assuming a simple representation of the HIT.

The static optimization problem studied here transparently expresses the core tension of the mandate decision: the higher vaccination rate achieved by a mandate reduces illness relative to decentralized vaccination, but it raises the social cost of vaccination. Analysis of the static problem is much simpler than study of policy making with a dynamic compartmental model.

The tradeoff is that abstracting from dynamics is an idealization. It disregards the possibility that the social cost of illness and vaccination may depend not only on the prevalence of illness and vaccination but also on the timing. For example, capacity constraints in hospitals might make the social cost of illness depend on the timing of severe cases, not just their totality. Heterogeneity in susceptibility and infectiousness may

make policy effectiveness depend on the temporal sequence of vaccination of persons with different attributes. Welfare analysis of a dynamic optimization problem could shed light on these matters.

2. Optimization with Complete Knowledge of the Social Cost Function

2.1. Direct and Indirect Effects of Vaccination on Illness

To formalize the planning problem, suppose that a vaccine has been developed, approved, and is available for use at a specified point in time. Using the compartmental concepts of SIR models, let $S > 0$ denote the fraction of the population who are susceptible to illness in the absence of vaccination. Let $I > 0$ denote the fraction who are currently infected. Let $R \geq 0$ denote the fraction who have recovered from previous infection. Assume that a recovered person has immunity and so cannot later become susceptible or infectious. Also assume that the death rate from illness is negligible. Then $S + I + R = 1$.

Vaccination may benefit a susceptible person by generating an immune response that prevents illness. It may also prevent this person from becoming infectious, thereby preventing transmission of the disease to susceptible persons who are unvaccinated or unsuccessfully vaccinated. Thus, vaccination may have a direct preventive effect on the person vaccinated and an indirect preventive effect on other persons. It is generally thought that the immune response sufficient to achieve a direct preventive effect is weaker than the response required to achieve an indirect preventive effect. Infectiousness is prevented if vaccination yields *sterilizing immunity*, which means an immune response that prevents the virus from replicating inside the body.

To formalize the direct effect, suppose that the available vaccine generates a protective immune response in a vaccinated person that lasts over the social horizon with probability $\lambda \geq 0$ and confers no immunity with probability $1 - \lambda$. Let v denote the rate of vaccination of susceptible persons. Then the *effective vaccination rate* is λv . The fraction of the population who are susceptible given vaccination rate v is $(1 - \lambda v)S$.

I maintain the usual assumption that vaccination has beneficial indirect effects. In principle, it might have negative indirect effects. In the context of COVID-19, persons receiving vaccines that are less than fully effective may perhaps behave in ways that increase disease transmission, foregoing mask-wearing and social distancing. Vaccination may encourage evolutionary selection of the virus towards variants resistant to the vaccine. These negative biological forces have previously been thought to be weak in practice (Mishra *et al.*, 2012) and will not be considered here.

To formalize the beneficial indirect effect of vaccination, suppose that it prevents a susceptible person from transmitting disease over the social horizon with probability $\mu \geq 0$. Then the fraction of the population who cannot transmit disease is $(\mu v)S + R$. Recall that prevention of infectiousness requires a stronger immune response than prevention of illness. Hence, we should expect that $\mu \leq \lambda$.

Let the infection rate p give the fraction of susceptible persons who become infected over the social horizon. In general, the infection rate may vary with $[(S, R), (\lambda, \mu), v]$. Section 2.2 poses a model with a specific functional form. For conciseness, I define $\theta \equiv [(S, R), (\lambda, \mu), v]$ to be the vector of parameters, and I write the infection rate as $p(\theta)$.

Illness *prevalence*, the fraction of the population who become ill over the social horizon, equals the fraction who are susceptible times the rate of infection among those who are susceptible. When the vaccination rate is v , illness prevalence is $p(\theta) \cdot [(1 - \lambda v)S]$. Increasing the vaccination rate reduces illness prevalence in two ways. It reduces the fraction $(1 - \lambda v)S$ of persons who are susceptible, and it may reduce the rate $p(\theta)$ of illness among those who are susceptible.

This formalization generalizes the one in Manski (2010, 2017). That paper considered vaccination prior to spread of disease and hence assumed that $S = 1$. It assumed that a vaccine generates a protective direct immune effect if and only if it generates an indirect effect that prevents a person from being infectious; hence, $\lambda = \mu$. It assumed that the infection rate is a weakly decreasing function of the effective vaccination rate λv . These assumptions implied that illness prevalence is $p(\lambda v)(1 - \lambda v)$, which decreases with λv .

The present paper does not assume that $S = 1$ because COVID-19 vaccines have been developed in the midst of the pandemic, when many persons have already been infected and recovered; thus, $S < 1$ and $R > 0$. Moreover, it may not be realistic to assume that $\lambda = \mu$ for COVID-19 vaccines.

2.2. Illness Prevalence with A Sharp Herd Immunity Threshold

Given values for $[(S, R), (\lambda, \mu), v]$, illness prevalence depends on the shape of the infection-rate function $p(\cdot)$. This function has a specific form under a simple model of disease transmission that makes these assumptions: (i) infection spreads until an HIT h^* is reached and then ceases immediately; (ii) the transmission process is rapid enough to occur entirely within the social horizon; (iii) the fraction of the population who are currently infectious when vaccination occurs is negligible, so $I \approx 0$.

These assumptions may not hold exactly, but they may be sufficiently reasonable to warrant application as approximations. The realism of assumption (ii) necessarily depends on the specified social horizon. Given this, dynamic epidemiological models can shed some light on the accuracy of assumptions (i) and (ii).

Assumption (iii) enables approximation of the compartmental equation by $S + R = 1$ and, hence, the fraction of non-transmitters in the population by $(\mu v)S + 1 - S$. Empirical evidence supports the assumption in the setting of COVID-19. It is thought that persons infected by the SARS-CoV-2 virus on average are infectious for about ten days. Combining this average infection period with an estimate of new cases in December 2020 implies that the fraction of the American population who were infectious when vaccination commenced in late December was about 0.033.³

³ *New York Times* (2020) reported that the number of newly confirmed cases per day in the week ending December 20, 2020 was 216,070. Confirmed cases undercount actual cases to an unknown degree. A plausible possibility is that the actual number of cases was five times the number of confirmed cases (Kalish *et al.*, 2021). If so, the actual number of new cases per day was 1,080,350. Multiplying this by a ten-day infection period implies an estimate of 10,803,500 currently infected persons. The U.S. population in December 2020 was estimated to be about 331,000,000. This yields the estimate $I = 0.033$.

Recall that h^* denotes the HIT. The assumptions imply that the infection rate is zero if $(\mu\nu)S + 1 - S \geq h^*$; hence, illness prevalence is zero as well. The infection rate may be positive if $(\mu\nu)S + 1 - S < h^*$. Then the sum of illness prevalence and the fraction of non-transmitters cannot exceed the HIT. That is,

$$(1) \quad p(\theta) \cdot [(1 - \lambda\nu)S] + [(\mu\nu)S + 1 - S] \leq h^*.$$

If the fraction $(1 - \lambda\nu)S$ of potentially ill persons exceeds the distance $h^* - [(\mu\nu)S + 1 - S]$ of society from herd immunity, infection spreads until the population attains herd immunity, making inequality (1) an equation. If $(1 - \lambda\nu)S < h^* - [(\mu\nu)S + 1 - S]$, all susceptible persons become infected. Combining these results, the infection rate among susceptible persons is

$$(2) \quad p(\theta) = \begin{cases} \{h^* - [(\mu\nu)S + 1 - S]\} / [(1 - \lambda\nu)S] & \text{if } (\mu\nu)S + 1 - S < h^* \text{ and } (1 - \lambda\nu)S \geq h^* - [(\mu\nu)S + 1 - S], \\ 1 & \text{if } (\mu\nu)S + 1 - S < h^* \text{ and } (1 - \lambda\nu)S \leq h^* - [(\mu\nu)S + 1 - S], \\ 0 & \text{if } (\mu\nu)S + 1 - S \geq h^*. \end{cases}$$

2.3. The Social Cost of Illness and Vaccination

I assume that the social cost of illness and vaccination is an additive function of illness prevalence and the vaccination rate. Let $b > 0$ denote the mean social cost resulting from a case of illness. Let $k > 0$ denote the mean social cost per vaccination, measured in commensurate units. The values of b and k are intended to express all socially relevant health, financial, and other costs generated by illness and vaccination. The additive structure of the social cost function implies that comparison of alternative vaccination rates only requires knowledge of the ratio of b to k , not the magnitudes of the two parameters.

The COVID-19 vaccination plans being formulated by state health agencies in the United States do not restrict vaccination to persons thought to be susceptible. Hence, I assume that a vaccination rate v is applied to the entire population rather than only to those who are susceptible. Then the social cost of rate v is

$$(3) \quad S(v) = b \cdot p(\theta) \cdot [(1 - \lambda v)S] + kv.$$

In the model with a sharp HIT, $p(\theta)$ given in (2) implies that, for $v \in [0, 1]$, (3) becomes

$$(4) \quad \begin{aligned} S(v) &= b \cdot \{h^* - [(\mu v)S + 1 - S]\} + kv && \text{if } (\mu v)S + 1 - S < h^* \text{ and } 1 + (\mu - \lambda)vS \geq h^*, \\ &= b \cdot (1 - \lambda v)S + kv && \text{if } (\mu v)S + 1 - S < h^* \text{ and } 1 + (\mu - \lambda)vS \leq h^*, \\ &= kv && \text{if } (\mu v)S + 1 - S \geq h^*. \end{aligned}$$

2.3.1. Optimal Unconstrained Choice of Vaccination Rate

A planner with complete knowledge of the social cost function and the power to choose any vaccination rate can choose a rate that minimizes $S(v)$ over all $v \in [0, 1]$. Minimization of the general form (3) of the social cost function is not conducive to analysis because the infection-rate function $p(\theta)$ is an unspecified function of v . Minimization yields an explicit solution when there is a sharp HIT.

Equation (4) has three regimes. In each regime, social cost is a linear function of v , subject to linear inequalities. The analysis below assumes that $\mu \leq \lambda$.

In regime 1, v satisfies the inequalities $(\mu v)S + 1 - S < h^*$ and $1 + (\mu - \lambda)vS \geq h^*$. The first inequality requires $v < [h^* - (1 - S)]/\mu S$, which is feasible if $h^* > 1 - S$. The second holds for all v if $\mu = \lambda$ and requires $v \leq (1 - h^*)/[(\lambda - \mu)S]$ if $\mu < \lambda$. For vaccination rates that satisfy these inequalities, social cost increases with v if $k \geq b\mu S$ and decreases if $k \leq b\mu S$. Hence, the optimal rate subject to the inequalities is $v = 0$ if $k \geq b\mu S$. It is $v = \min\{1, [h^* - (1 - S)]/\mu S, (1 - h^*)/[(\lambda - \mu)S]\}$ if $k \leq b\mu S$.

In regime 2, v satisfies the inequalities $(\mu v)S + 1 - S < h^*$ and $1 + (\mu - \lambda)vS \leq h^*$. The first inequality again requires $v < [h^* - (1 - S)]/\mu S$, which is feasible if $h^* > 1 - S$. The second requires $v \geq (1 - h^*)/[(\lambda - \mu)S]$, which is feasible if $1 - h^* \leq (\lambda - \mu)S$. Thus, feasible v satisfy $(1 - h^*)/[(\lambda - \mu)S] \leq v < [h^* - (1 - S)]/\mu S$ when this interval exists. Social cost increases with v if $k \geq b\lambda S$ and decreases if $k \leq b\lambda S$. Hence, the optimal vaccination rate in regime 2 is $v = \min\{1, (1 - h^*)/[(\lambda - \mu)S]\}$ if $k \geq b\lambda S$ and $v = \min\{1, [h^* - (1 - S)]/\mu S\}$ if $k \leq b\lambda S$.

In regime 3, v satisfies the inequality $(\mu v)S + 1 - S \geq h^*$. Hence, $v \geq [h^* - (1 - S)]/\mu S$, which is feasible if $h^* - (1 - S) \leq \mu S$. Social cost is kv , an increasing function of v , so the optimal vaccination rate is the smallest value satisfying the inequality. Thus, the optimal rate in this regime is $v = \max\{0, [h^* - (1 - S)]/\mu S\}$.

The overall optimal vaccination rate minimizes social cost across the regimes that are feasible.

2.3.2. Optimal Constrained Choice Between Mandate and Decentralization

A planner who is only empowered to choose between a mandate and decentralization makes a constrained choice. Let v_d denote the vaccination rate generated by decentralized decision making. Assume that $v = 1$ with a mandate.⁴ The constrained optimal decision is

$$\begin{aligned} (5) \text{ mandate if} \quad & b \cdot p(\theta_m) \cdot [(1 - \lambda)S] + k \leq b \cdot p(\theta_d) \cdot [(1 - \lambda v_d)S] + k v_d, \\ \text{do not mandate if} \quad & b \cdot p(\theta_m) \cdot [(1 - \lambda)S] + k \geq b \cdot p(\theta_d) \cdot [(1 - \lambda v_d)S] + k v_d, \end{aligned}$$

where $\theta_m = [(S, R), (\lambda, \mu), 1]$ and $\theta_d = [(S, R), (\lambda, \mu), v_d]$. If there is a sharp HIT, θ_m and θ_d are the values given in (2) with $v = 1$ and $v = v_d$ respectively.

⁴ In practice v may be less than one because states may exempt persons with religious objections to vaccination.

2.4. Optimization with Illustrative Parameter Values

Knowledge of $[p(\cdot), (S, R), (\lambda, \mu), (b, k)]$ enables a planner to choose an optimal unconstrained vaccination rate. Additionally, knowledge of the decentralized rate v_d enables choice between a mandate and decentralization. I illustrate here, using the model with a sharp HIT and values of its parameters $[h^*, S, (\lambda, \mu), (b, k)]$ that may perhaps be plausible when considering COVID-19 at the end of 2020. However, section 3 will argue that knowledge of these parameters is seriously incomplete and, hence, optimization is not feasible.

Empirical evidence and modeling provide partial grounding for selection of values for (h^*, S, λ, v_d) . Findings reported for the Pfizer and Moderna vaccine trials suggest $\lambda = 0.95$ as a realistic estimate of the probability that these vaccines generate immune responses that prevent illness; see Section 3.3 for further discussion. Epidemiologists have conjectured many values for the COVID-19 HIT. I use $h^* = 0.7$, which has been a common conjecture (McNeil, 2020). A Gallup poll surveying the American public in November 2020 about their willingness to be vaccinated suggests $v_d = 0.58$ as an estimate of the decentralized vaccination rate (Reinhart, 2020).

A perhaps plausible estimate of the fraction of susceptible persons in late December 2020 is $S = 0.73$. To obtain this estimate, I begin with the cumulative confirmed national case count on December 20, which was about 17.9 million. Epidemiologists have conjectured a wide range of possibilities for the degree to which confirmed cases undercount actual cases. A plausible possibility is that the actual number of cases is five times the number of confirmed cases (Kalish *et al.*, 2021), which implies that the actual cumulative number of cases is about 89.5 million. The size of the American population is about 331 million. Applying the approximation that $I \approx 0$, it follows that $R = 89.5/331$ and hence that $S = 1 - 89.5/331 = 0.73$.

Choice of exact values for (μ, b, k) must be speculative. The investigators in vaccine trials have not reported evidence that is directly informative about μ . I use the value $\mu = 0.8$. Careful analysis of the costs of illness and vaccination may in principle enable realistic estimation of (b, k) , but I am not aware of such

analysis to date. Comparison of vaccination rates only requires knowledge of the ratio of b and k . I use $b/k = 10$. Recognizing that these values are speculative, I will replace them by bounds in Section 4.

With these parameter values, the threshold value of v that separates regimes 1 and 2 from regime 3 is $[h^* - (1 - S)]/(\mu S) = 0.736$ and the threshold separating regime 1 from regime 2 is $(1 - h^*)/[(\lambda - \mu)S] = 2.74$. Thus, the second regime is not feasible. Regime 1 applies when $0 \leq v < 0.736$ and regime 3 applies when $0.736 \leq v \leq 1$. The inequality $k \leq b\mu S$ holds, so social cost decreases with v in regime 1, whereas it increases with v in regime 3. Hence, social cost is minimized at $v = 0.736$.

Figure 1 shows graphically how social cost and post-vaccination illness prevalence vary with the vaccination rate. The graph normalizes b/k by setting $b = 1$ and $k = 0.1$. The upper graph displays the social-cost function $S(v)$. The lower graph displays illness prevalence, which equals $h^* - (\mu v + 1 - S)$ in regime 1 and zero in regime 3.

Finally, consider choice between a mandate and decentralization. Social cost under a mandate is $k = 0.1$. It is 0.149 with the assumed decentralized vaccination rate $v_d = 0.58$. Hence, a mandate yields lower social cost than decentralization with the assumed parameter values.

3. COVID-19 Uncertainties

Optimization remains feasible if a planner lacks the complete knowledge assumed in Section 2 but knows enough to determine what vaccination rate minimizes social cost. Optimization is not possible with less information. In the setting of COVID-19, knowledge is seriously incomplete. I explain here.

3.1. The Infection-Rate Function

Manski (2010, 2017) considered policy choice with incomplete knowledge of the infection-rate function $p(\cdot)$. Epidemiological models predict that the future rate of infection of susceptible persons weakly decreases

with the fraction of the population who have already achieved immunity by vaccination or previous illness. This broad idea makes it credible to assume that $p(\cdot)$ decreases with $(\mu\nu)S + R$, but it does not per se suggest further structure on the function.

Specific models imply further structure. Assuming a hard HIT implies the functional form (2), which makes the infection rate a function of three quantities: the fraction $(\mu\nu)S + R$ of non-transmitters in the population, the fraction $(1 - \lambda\nu)S$ of persons susceptible to infection, and the HIT h^* .

The HIT has long been a prominent concept in epidemiology, but a persistent problem has been to estimate its magnitude credibly. A serious difficulty has been scarcity of empirical evidence revealing how disease transmission varies with the compartmental composition of the population. Studies of disease transmission have to rely on the limited observational data generated by actual epidemics. Randomized clinical trials (RCTs) may enable one to learn the direct effect of vaccination on illness in vaccinated persons, but RCTs typically do not provide evidence on the indirect effect of vaccination preventing transmission of disease. An RCT that randomly vaccinates a specified fraction of a population only reveals the illness outcomes that occur with the vaccination rate used in the trial. The outcomes that the population would experience with other vaccination rates remain counterfactual. Yet policy choice requires comparison of alternative vaccination rates.⁵

Scarcity of empirical evidence is problematic when studying any infectious disease, but it is particularly severe when studying COVID-19. This is a new disease whose transmission process differs in some respects from that of other infectious diseases. At the beginning of the pandemic essentially no data were available to provide an empirical foundation for modeling disease transmission. The influential early COVID-19 model developed in Ferguson *et al.* (2020) made forecasts using a modified version of a simulation model previously

⁵ Vaccine trials can reveal indirect effects if the population partitions into isolated groups (aka clusters) of persons. Then the members of each group may infect one another but not the members of other groups. In such cases, one can define treatment units to be groups rather than persons, randomly assign varying vaccination rates to different groups, and use the trial to learn about illness outcomes under alternative vaccination rates. Hudgens and Halloran (2008) develop methodology for analysis of RCTs performed in such settings. Loeb *et al.* (2010) report a trial performed on isolated Hutterite communities in Canada. However, populations rarely partition in modern societies. RCTs have no identifying power in the polar case of a fully connected society, where social interactions are global rather than local (Manski, 2013).

developed to support pandemic influenza planning. Observational data in various locations has gradually accumulated as the COVID-19 pandemic has progressed, but considerable uncertainty about disease transmission persists.

3.2. Compartmental Composition of the Population

Serious data problems have limited knowledge of the fraction S of local, regional, and national populations who are currently susceptible to COVID-19 and, correspondingly, the fraction R who have previously been ill and are now recovered.

Measurement of cumulative infection has been hampered by substantial missing data. Confirmed cases have been measured by rates of positive findings among persons who have been tested for infection. Infection data are missing for persons who have not been tested. The persons who have been tested differ considerably from those who have not been tested. Criteria used to determine eligibility for testing have typically required demonstration of symptoms associated with presence of infection or close contact with infected persons. This gives considerable reason to believe that some fraction of untested persons are asymptomatic or pre-symptomatic carriers of the COVID-19 disease. Hence, the actual cumulative rate of infection has been higher than the reported rate.

A second problem of data quality is that measurement of confirmed cases is imperfect because the prevalent nasal swab tests for infection are not fully accurate. There is basis to think that accuracy of nasal swab tests is highly asymmetric, with few false positive results but many false negative ones. Given this asymmetry, the actual rate of infection has again been higher than the reported rate.

Combining the problems of missing data and imperfect test accuracy yields the conclusion that reported cumulative rates of infections are lower than actual rates. Researchers have put forward widely varying point estimates derived in various ways. The estimates differ in the assumptions used. The assumptions vary and so do the findings. No assumption or estimate has been thought sufficiently credible as to achieve consensus.

Rather than make point estimates that lack foundation, Manski and Molinari (2021) combined available data with credible assumptions to bound the cumulative infection rate at specific locations and dates. The bounds were wide.

3.3. Direct and Indirect Effectiveness of Vaccination

The probabilities λ and μ respectively measure the direct and indirect preventive effects of vaccination on illness. We currently lack an empirical basis to measure μ . We have some capacity to measure λ .

The investigators in ongoing COVID-19 vaccine trials have to date reported no findings that are informative regarding μ . They could in principle measure *viral shedding*, the release of virus from a person's body into the environment, but they have not done so. Measurement of viral shedding would help in estimating μ , but it would not suffice. Infectiousness is determined by both biology and behavior. Viral shedding is a biological prerequisite for infectiousness but behavior, including lack of social distancing and mask wearing, determines whether shed virus is transmitted to other persons.

Ongoing vaccine trials report a measure of *efficacy* that equals λ given certain assumptions. Consider a trial in which persons are randomized into receipt of a vaccine ($z = 1$) or no vaccine ($z = 0$). Let $y(t) = 1$ if a person becomes ill within a time period of a specified length t after the date of vaccine receipt/non-receipt and let $y(t) = 0$ if the person does not become ill. Trial data enable estimation of $P[y(t) = 1|z = 1]$ and $P[y(t) = 1|z = 0]$, the probabilities of illness conditional on vaccination status. Efficacy is defined as one minus the risk ratio; that is, $1 - P[y(t) = 1|z = 1]/P[y(t) = 1|z = 0]$. Pfizer Inc. reported an efficacy estimate of 0.95 for its COVID-19 vaccine in a press release on November 18, 2020.⁶ Moderna has reported a similar estimate.

The measurements of $P[y(t) = 1|z = 1]$ and $P[y(t) = 1|z = 0]$ made public by Pfizer and Moderna have thus far concerned relatively short time periods, whereas the social horizon for vaccination policy presumably is much longer. If the trial and social horizons are identical, knowledge of $P[y(t) = 1|z = 1]$ and $P[y(t) = 1|z =$

⁶ <https://www.pfizer.com/news/press-release/press-release-detail/pfizer-and-biontech-conclude-phase-3-study-covid-19-vaccine>, accessed December 11, 2020.

0] reveals λ if one assumes that a vaccinated person becomes ill if and only if the person lacks immunity and is exposed to the virus. In the absence of natural immunity, the probability that a vaccinated person lacks immunity is $1 - \lambda$. In a perfect RCT, the observed illness rate $P[y(t) = 1|z = 0]$ of unvaccinated persons equals the probability of exposure to the virus. Hence, $P[y(t) = 1|z = 1] = (1 - \lambda)P[y(t) = 0|z = 0]$. Solving this equation gives $\lambda = 1 - P[y(t) = 1|z = 1]/P[y(t) = 0|z = 0]$. Thus, λ is the vaccine efficacy.

3.4. Social Costs of Illness and Vaccination

Uncertainty about the social costs of illness and vaccination is a generic problem when evaluating vaccination policies. There may be uncertainty about the relative frequencies of severe and mild cases, the former being more costly to treat and generating greater health harm than the latter. There may be uncertainty about the frequency and severity of long-term side effects of vaccination, which may occur after the observation period in trials end.

These uncertainties may be substantial in the context of COVID-19. It is widely thought that the fractions of asymptomatic cases and ones with mild symptoms are grossly underestimated in official statistics because persons with such cases tend not to be tested. However, there is no consensus on the extent of underestimation. Essentially nothing is presently known about the long-term side effects of vaccination because the observation periods in trials have thus far been short.

3.5. Decentralized Vaccination Rate

Finally, there is uncertainty about the fraction of the population who would choose to be vaccinated in a regime with decentralized decision making. Some suggestive data are available in surveys undertaken during the course of pandemic. Samples of persons have been asked whether they would choose to be vaccinated

should vaccines with specified characteristics be approved and made available. The findings have varied over time and place.

4. Policy Choice under Uncertainty

4.1. Some Basic Decision Theory

To cope with the uncertainties described above, I bring to bear some basic decision theory, applied previously to choice of vaccination policy in Manski (2010, 2017). To begin, suppose that a decision maker must choose among a set of feasible actions. The welfare achieved by an action depends on an unknown feature of the environment, called the *state of nature*. The decision maker lists all the states of nature that he believes could possibly occur. This list, the *state space*, expresses partial knowledge. The larger the state space, the less the decision maker knows about the outcome of each action.

In this paper, an unconstrained planner chooses a vaccination rate v in the set $[0, 1]$ of all possible rates. A constrained planner chooses between a mandate and decentralization. In either case, welfare is measured by social cost function (3). In the unconstrained case, the state space is the set of values for $[p(\cdot), (S, R), (\lambda, \mu), (b, k)]$ that the planner deems possible. In the constrained case, the decentralized vaccination rate v_d adds an additional dimension to the state space.

The fundamental difficulty of planning under uncertainty is clear even in a simple setting with two feasible actions and two states of nature. Suppose that one action yields lower social cost in one state of nature and the other action yields lower social cost in the other state. Then the planner does not know which action is better. Recognizing that optimization is impossible, decision theory offers various reasonable criteria for decision making, each with its own properties.

To economists, the most familiar decision criterion places a subjective probability distribution on the state space and minimizes subjective expected social cost. This approach, commonly called Bayesian

planning, is compelling if a planner feels able to select a credible subjective distribution. However, a subjective distribution is a form of knowledge, and a planner may not feel able to assert one.

When it is difficult to assert a credible subjective distribution, a reasonable way to act is to use a decision criterion that achieves uniformly satisfactory results, whatever the true state of nature may be. There are multiple ways to formalize the idea of uniformly satisfactory results. The two most commonly studied are the minimax and minimax-regret criteria.

The minimax criterion chooses an action that minimizes the maximum social cost that might possibly occur. The minimax-regret criterion considers each state of nature and computes the loss in welfare that would occur if one were to choose a specified action rather than the one that is best in this state. This quantity, called *regret*, measures the nearness to optimality of the specified action in the state of nature. The decision maker must choose without knowing the true state. To achieve a uniformly satisfactory result, he computes the maximum regret of each action; that is, the maximum distance from optimality that the action would yield across all possible states of nature. The criterion chooses an action that minimizes this maximum distance from optimality.⁷

4.2. Computational Determination of Bayesian, Minimax, and Minimax-Regret Vaccination Policies

Each of the decision criteria described above --- Bayesian, minimax, and minimax-regret --- merits consideration as an approach to choose a vaccination policy. Algebraic determination of the policies yielded by the criteria is not always feasible when multiple uncertainties are present, but computation is tractable. To demonstrate, I have implemented an algorithm that assumes there exists a sharp HIT. Then the state space lists the feasible values of $[h^*, S, (\lambda, \mu), (b, k), v_d]$, with v_d being relevant in the case of constrained planning. I summarize the main features of the algorithm here.

⁷ The minimax and minimax-regret criteria are sometimes confused with one another, but they yield the same choice only in certain special cases. Whereas the former criterion considers the worst outcome that an action may yield in absolute terms, the latter considers the worst outcome relative to what is achievable in a given state of nature.

The algorithm is applicable when the state space is a specified hyperrectangle of values for $[h^*, S, (\lambda, \nu), (1, k), \nu_d]$, where $\mu \equiv \nu\lambda$. Thus, ν expresses the indirect effect of vaccination as a fraction of the direct effect. Setting $b = 1$ is without loss of generality, because the policy chosen by each decision criterion depends only on the ratio b/k rather than on b and k separately.

The rectangular state space is formed by placing specified lower and upper bounds on each of the parameters $[h^*, S, (\lambda, \nu), k, \nu_d]$. All parameters are logically non-negative, so each lower bound must be non-negative. All parameters except k are logically less than or equal to one, so their upper bounds must be less than or equal to one. In principle k may exceed one, but such cases yield the trivial result that the optimal vaccination rate is zero whatever values the other parameters take.

Computation of the minimax policy is simple when the state space is rectangular. Inspection of equation (4) shows that, for all values of ν , the state that maximizes social cost places (h^*, S, k) at their upper bounds and (λ, ν) at their lower bounds.

To make computation of minimax-regret policies tractable, a user of the algorithm approximates the state space by a finite grid of parameter values. A selected density of the grid controls the closeness of the approximation. The performance of a given vaccination rate is evaluated by computing its maximum regret over the grid of values.

Bayesian planning requires specification of a subjective probability distribution on the state space. The algorithm enables the user to place a selected Beta distribution on the bound for each parameter and creates a joint distribution by assuming subjective statistical independence across parameters. The subjective expected social cost of a given vaccination rate is approximated by Monte Carlo integration. That is, the algorithm makes repeated simulated draws of parameter values from the joint distribution, computes social cost with each draw, and then averages the computed social cost across draws. The user-chosen number of simulated draws controls the closeness of the approximation.

4.3. Policy Choices with an Illustrative State Space

To illustrate, I determine minimax and minimax-regret vaccination rates for a specified state space. I do not present findings for Bayesian policy choice, which depends on the subjective distribution placed on the state space.

I let the state space be the hyperrectangle with $S \in [0.5, 0.8]$, $h^* \in [0.6, 0.8]$, $\lambda = 0.95$, $v \in [0.7, 0.9]$, and $k \in [0.05, 0.2]$. This space is a neighborhood of the state used in the illustration of Section 2.4. It contains intervals of parameter values that appeared credible in late December 2020. I use the single point $\lambda = 0.95$ because this quantity, alone among the parameters, has a relatively firm evidentiary grounding.

For this state space, the upper graph in Figure 2 shows how maximum social cost across different states of nature varies with the vaccination rate. The lower graph displays the post-vaccination illness rates associated with these social costs. The slight kinks in the graphs that occur at about the vaccination rate $v = 0.9$ arise because the determination of maximum social cost transitions from regime 1 to regime 2 here.

We find that maximum social cost and the associated illness prevalence both decrease as the vaccination rate increases. Hence, the policy that minimizes maximum social cost vaccinates the complete population. This implies that a vaccine mandate is socially preferable to decentralized vaccination, regardless of the fraction of the population who would choose to vaccinate with decentralization.

Figure 3 shows findings for maximum regret across the state space, as a function of the vaccination rate. Maximum regret initially decreases as the vaccination rate increases, attains an interior minimum at about $v = 0.74$, and then increases as the vaccination rate increases further. The minimum occurs at the threshold between vaccination regimes 1 and 3. The illness rate in the state yielding maximum regret initially decreases continuously and then drops discontinuously to zero at about $v = 0.74$, where the transition to regime 3 occurs. The values of (S, h^*, v, k) that maximize regret all change discontinuously at this value of v .

Considered from the perspective of maximum regret, the choice between a mandate and decentralization varies with the decentralized vaccination rate. A mandate yields maximum regret equal to 0.2. Maximum

regret with decentralization is greater than 0.2 for $v_d < 0.63$ and less than 0.2 when $0.63 \leq v_d < 1$. Hence, a mandate is preferable if one deems it feasible that $v_d < 0.63$. Decentralization is preferable if one is certain that $v_d \geq 0.63$.

References

- Aguas, R., R. Corder, J. King, G. Gonçalves, M. Ferreira, and M Gomes, (2020), Herd immunity thresholds for SARS-CoV-2 estimated from 2 unfolding epidemics," medRxiv, <https://www.medrxiv.org/content/10.1101/2020.07.23.20160762v3> .
- Althouse, B., T. Bergstrom, and C. Bergstrom (2010), "A Public Choice Framework for Controlling Transmissible and Evolving Diseases," *Proceedings of the National Academy of Sciences*, 107, 1696-1701.
- Anderson R. and R. May (1991), *Infectious Diseases of Humans: Dynamics and Control*, Oxford: Oxford University Press.
- Ball F. and O. Lyne (2002), "Optimal Vaccination Policies for Stochastic Epidemics among a Population of Households," *Mathematical Biosciences*, 177&178, 333–354.
- Becker, N. and D. Starczak (1997), "Optimal Vaccination Strategies for a Community of Households," *Mathematical Biosciences*, 139, 117-132.
- Boulier, B. T. Datta, and R. Goldfarb (2007), "Vaccination Externalities," *The B.E. Journal of Economic Analysis & Policy* 7, issue 1, article 23.
- Brito, D., E. Sheshinski, and M. Intriligator (1991), "Externalities and Compulsory Vaccinations," *Journal of Public Economics*, 45, 69-90.
- Centers for Disease Control and Prevention (2006), *Principles of Epidemiology in Public Health Practice, Third Edition*, www.cdc.gov/ophss/csels/dsepd/ss1978/ss1978.pdf, accessed March 14, 2016.
- Chen, F. and F. Toxvaerd (2014), "The Economics of Vaccination," *Journal of Theoretical Biology*, 363, 105-117.
- Ferguson, N. et al. (2020), "Report 9: Impact of Non-pharmaceutical Interventions (NPIs) to Reduce COVID-19 Mortality and Healthcare Demand," Imperial College London, <https://www.imperial.ac.uk/media/imperial-college/medicine/mrc-gida/2020-03-16-COVID19-Report-9.pdf> , accessed June 4, 2020.
- Fine, P. (1993), "Herd Immunity: History, Theory, Practice," *Epidemiological Reviews*, 15, 265-302.
- Fine, P, K. Eames, and D. Heymann (2011), "'Herd Immunity': A Rough Guide," *Clinical Infectious Diseases*, 52, 911–916.
- Francis, P. (1997), "Dynamic Epidemiology and the Market for Vaccinations," *Journal of Public Economics*, 63, 383-406.
- Hill, A. and I. Longini (2003), "The Critical Vaccination Fraction for Heterogeneous Epidemic Models," *Mathematical Biosciences*, 181, 85-106.
- Hudgens, M. and E. Halloran (2008), "Toward Causal Inference with Interference," *Journal of the American Statistical Association*, 103, 832-842.

IHME COVID-19 Health Service Utilization Forecasting Team (2020). "Forecasting COVID-19 impact on hospital bed-days, ICU-days, ventilator-days and deaths by US state in the next 4 months," Institute for Health Metrics and Evaluation, University of Washington,

<https://www.medrxiv.org/content/10.1101/2020.03.27.20043752v1.full.pdf>.

Immunization Action Coalition (2020), "State Laws and Mandates by Vaccine," <https://www.immunize.org/laws/>, accessed November 23, 2020.

Kalish, H. *et al.* (2021), "Mapping a Pandemic: SARS-CoV-2 Seropositivity in the United States," medRxiv; <https://doi.org/10.1101/2021.01.27.21250570>.

Keeling, M. and A. Shattock (2012), "Optimal but Unequitable Prophylactic Distribution of Vaccine," *Epidemics*, 4, 75-85.

Loeb, M., M. Russell, L. Moss, K. Fonseca, J. Fox, D. Earn, F. Aoki, G. Horsman, P. Van Caesele, K. Chokani, M. Vooght, L. Babiuk, R. Webby, and S. Walter (2010), "Effect of Influenza Vaccination of Children on Infection Rates in Hutterite Communities: A Randomized Trial," *Journal of the American Medical Association*, 303, 943-950.

Manski, C. (2010), "Vaccination with Partial Knowledge of External Effectiveness," *Proceedings of the National Academy of Sciences*, 107, 3953-3960.

Manski, C. (2013), "Identification of Treatment Response with Social Interactions," *The Econometrics Journal*, 16, S1-S23.

Manski, C. (2017), "Mandating Vaccination with Unknown Indirect Effects," *Journal of Public Economics Theory*, 19, 603-619.

Manski, C. and F. Molinari (2021) "Estimating the COVID-19 Infection Rate: Anatomy of an Inference Problem," *Journal of Econometrics*, 220, 181-192.

May, T. and R. Silverman (2005), "Free-riding, Fairness, and the Rights of Minority Groups in Exemption from Mandatory Childhood Vaccination," *Human Vaccines*, 1, 12-15.

McNeil, D. (2020), "How Much Herd Immunity Is Enough?" *New York Times*, December 24, <https://www.nytimes.com/2020/12/24/health/herd-immunity-covid-coronavirus.html>, accessed January 31, 2021.

Mishra, R., E. Oviedo-Orta, P. Prachi, R. Rappuoli and F. Bagnoli (2012), "Vaccines and Antibiotic Resistance," *Current Opinion in Microbiology*, 15, 596-602.

New York State Bar Association (2020), "New York State Bar Association Calls Upon State To Consider Mandating a Safe and Effective Vaccine if Voluntary Measures Fail To Protect Public Health," November 7, <https://nysba.org/new-york-state-bar-association-calls-upon-state-to-consider-mandating-a-safe-and-effective-vaccine-if-voluntary-measures-fail-to-protect-public-health/>, accessed November 25, 2020.

New York Times (2020), "Coronavirus in the U.S.: Latest Map and Case Count," December 21, <https://www.nytimes.com/interactive/2020/us/coronavirus-us-cases.html?action=click&module=Top+Stories&pgtype=Homepage>, accessed December 21, 2020.

Paltiel, D., J. Schwartz, A. Zheng, and R. Walensky (2021), "Clinical Outcomes of A COVID-19 Vaccine: Implementation over Efficacy," *Health Affairs*, 40, 42-52.

Patel, R., I. Longini, and E. Halloran (2005), "Finding Optimal Vaccination Strategies for Pandemic Influenza Using Genetic Algorithms," *Journal of Theoretical Biology*, 234, 201-212.

Reinhart, R. (2020), "More Americans Now Willing to Get COVID-19 Vaccine," <https://news.gallup.com/poll/325208/americans-willing-covid-vaccine.aspx> , accessed January 31, 2021.

Reiss, D. and A. Caplan (2020), "Considerations in Mandating a new Covid-19 Vaccine in the USA for Children and Adults," *Journal of Law and the Biosciences*, 7, 1–9, <https://doi.org/10.1093/jlb/ljaa025>.

Scuffham, P. and P. West (2002), "Economic Evaluation of Strategies for the Control and Management of Influenza in Europe," *Vaccine*, 20, 2562–2578.

Star Tribune (2020), "Evers won't say whether he'll seek mask mandate extension," September 9, <https://www.startribune.com/evers-won-t-say-whether-he-ll-seek-mask-mandate-extension/572363432/>, accessed November 23, 2020.

Thomas, A. (2009), "Mandatory Vaccination of Health Care Workers," *New England Journal of Medicine*, 2015-2017.

Vamos, C., R. McDermott, and E. Daley (2008), "The HPV Vaccine: Framing the Arguments FOR and AGAINST Mandatory Vaccination of All Middle School Girls," *Journal of School Health*, 78, 302-309.

Figure 1: Social Cost (upper graph) and Illness Prevalence (lower graph) as Functions of the Vaccination Rate

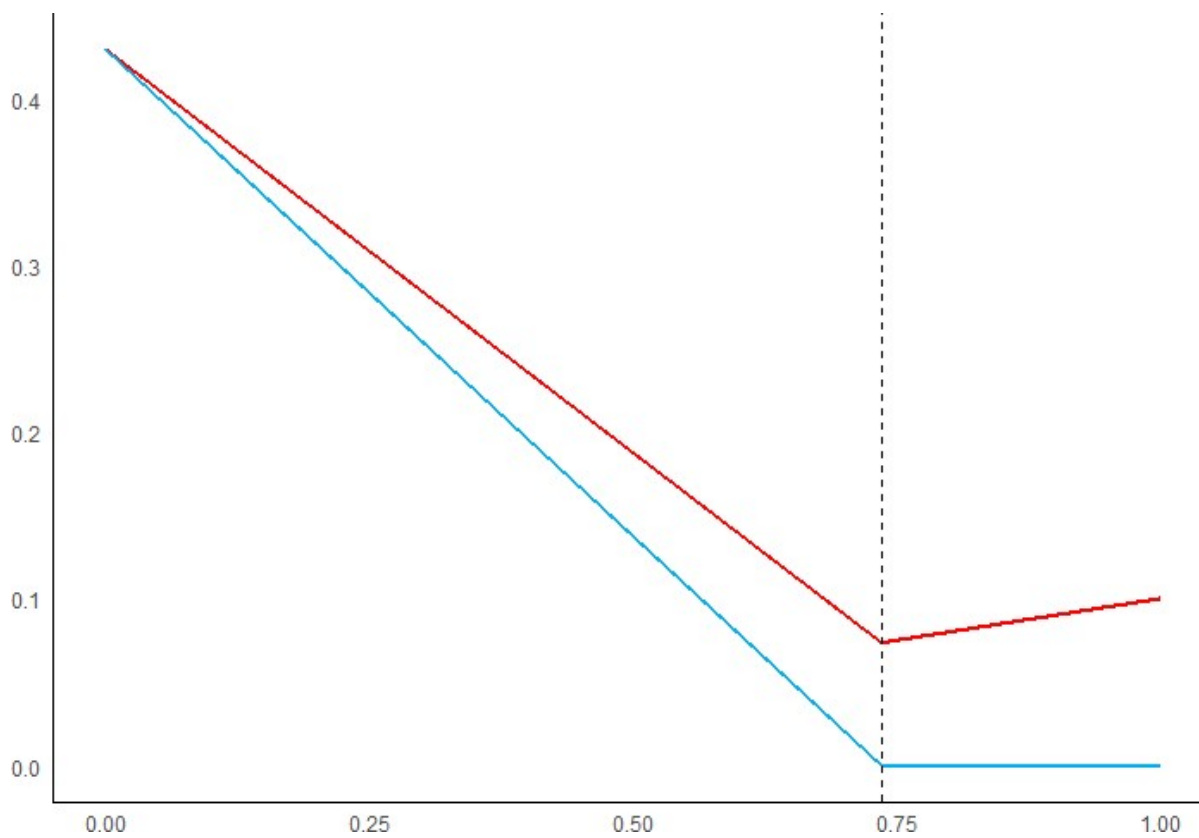


Figure 2: Maximum Social Cost (upper graph) and Associate Illness Prevalence (lower graph) as Functions of the Vaccination Rate




Figure 3: Maximum Regret (non-monotone graph) and Associate Illness Prevalence (decreasing graph) as Functions of the Vaccination Rate

