VACCINATION WITH PARTIAL KNOWLEDGE OF EXTERNAL EFFECTIVENESS

Charles F. Manski
Department of Economics and Institute for Policy Research
Northwestern University

Draft: October 2009

Abstract

Economists studying public policy have generally assumed that the relevant social planner knows how policy affects population behavior. Planners typically do not possess all of this knowledge, so there is reason to consider policy formation with partial knowledge of policy impacts. Here I consider choice of a vaccination policy when a planner has partial knowledge of the effect of vaccination on illness rates. To begin I pose a planning problem whose objective is to minimize the utilitarian social cost of illness and vaccination. The consequences of candidate vaccination rates depend on the extent to which vaccination prevents illness. I study the planning problem when the planner has partial knowledge of the external-response function, which expresses how the illness rate of unvaccinated persons varies with the vaccination rate. I suppose that the planner observes the illness rate of a study population whose vaccination rate has previously been chosen. He knows that the illness rate of unvaccinated persons weakly decreases as the vaccination rate increases, but he does not know the magnitude of the preventive effect of vaccination. In this setting, I first show how the planner can eliminate dominated vaccination rates and then how he can use the minimax or minimax-regret criterion to choose an undominated vaccination rate.

This research was supported in part by National Science Foundation grant SES-0911181. I am grateful for comments from Mike Fu.

1. Introduction

Economists and other researchers performing normative study of public policy have typically assumed that the policy maker, or *planner*, knows how policy affects population outcomes. A standard exercise specifies a set of feasible policies and a social welfare function. The planner is presumed to know the welfare achieved by each policy. The objective of the exercise is to determine the optimal policy.

There are many examples. Economists have studied optimal income taxation under the assumption that the planner knows how the tax schedule affects the population distribution of labor supply (Mirrlees, 1971). They have studied optimal criminal justice systems under the assumption that the planner knows how policing and sanctions affect crime rates (Polinsky and Shavell, 2000). Researchers studying optimal vaccination against infectious disease have typically assumed the planner knows how vaccination affects illness rates. See Ball and Lyne (2002), Becker and Starczak (1996), Brito, Sheshinski, and Intriligator (1991), Boulier, Datta, and Goldfarb (2007), Hill and Longini (2003), Patel, Longini, and Halloran (2005), and Scuffham and West (2002).

Whatever the policy choice may be, inferential problems commonly make it difficult to learn how policy affects outcomes. Perhaps the most fundamental difficulty is the identification problem arising from the unobservability of counterfactual outcomes. At most one can observe the outcomes that occur under realized policies—the outcomes of unrealized policies are logically unobservable. Yet determination of an optimal policy requires comparison of all feasible policies. For this and many other reasons, planners usually have only partial knowledge of the welfare achieved by alternative policies. This limits the relevance of the standard exercise to actual policy analysis.

This paper considers choice of a vaccination policy when a planner has partial knowledge of the effect of vaccination on illness rates. There are at least two sources of partial knowledge. First, the planner may only partially know the *internal* effectiveness of vaccination in generating an immune response that prevents a vaccinated person from become ill or infectious. Second, the planner may only partially know the *external* effectiveness of vaccination in preventing transmission of disease to members of the population who are unvaccinated or unsuccessfully vaccinated. I focus on the second issue, which commonly is more problematic.

To see the problem, consider an idealized setting where the members of the population are observationally identical and where the planner chooses the population vaccination rate. A standard randomized clinical trial, which vaccinates a small experimental group of individuals, enables evaluation of the internal effectiveness of vaccination. However, the trial does not reveal the external effect of applying different vaccination rates to the population. If the experimental group is small, the population vaccination rate is essentially zero. If a trial vaccinating a non-negligible fraction of the population is undertaken, the resulting outcome data only reveal the external effectiveness of the chosen vaccination rate. The outcomes with other vaccination rates remain counterfactual, yet choice of a vaccination policy requires comparison of alternative rates.

Attempting to cope with the absence of empirical evidence, researchers have used epidemiological models to forecast the outcomes that would occur with counterfactual vaccination policies. The articles on optimal vaccination cited earlier use a variety of such models. However, authors typically provide little information that would enable one to assess the accuracy of their assumptions about individual behavior, social interactions, and disease transmission. Hence, it is prudent to view their forecasts more as thought experiments predicting outcomes under specific assumptions than as accurate predictions of policy impacts.

If one wants to formally consider vaccination policy with partial knowledge, how might one go about it? When economists have studied planning with partial knowledge, it has been standard to assert a subjective probability distribution over the feasible states of nature and propose choice of an action that maximizes subjective expected welfare. For example, Nordhaus (2008) used this approach to express partial knowledge of parameter values in his assessment of global warming policy, writing (page 27): "This book takes the standard economic approach to uncertainty known as the expected utility model, which relies on an assessment with subjective or judgmental probabilities." Researchers studying vaccination have recently begun to use the expected utility model. Tanner, Sattenspeil, and Ntaimo (2008) exposit the model, under the name *stochastic programming*. They admonish vaccination researchers about the need to recognize uncertainty in policy evaluation, writing (pages 150–151): "We believe that since accurate parameter estimation can be extremely difficult for epidemic models, ignoring parameter uncertainty is not a good assumption to make when creating a vaccination policy." [1]

Use of the expected utility model to make policy choices with partial knowledge is reasonable when a planner has a credible basis for asserting a subjective probability distribution on unknown decision-relevant quantities. However, a subjective probability distribution is itself a form of knowledge, and a planner may have no credible basis for asserting one. I have previously studied problems of this type, where a planner faces *ambiguity*.[2] I have mainly focused on ambiguity arising from partial identification of policy impacts; see Manski (2000, 2002, 2005, 2006, 2007a, 2007b, 2009a, 2009b). I have also used the Wald (1950) formulation of statistical decision theory to examine policy choice using sample data; see Manski (2004, 2005) and Manski and Tetenov (2007).

Here I study choice of vaccination policy under ambiguity. To show concretely how basic principles may be applied, I mainly consider a relatively simple scenario in which a planner must choose a vaccination rate for a population of observationally identical persons. The planner has no knowledge of the external effectiveness of vaccination except that it is monotone in the vaccination rate. That is, the rate of illness

---

[1] Although vaccination researchers have generally not studied policy choice as a problem of planning with partial knowledge, they have often performed sensitivity analyses in which they determine optimal policy under alternative assumptions. See, for example, Weycker *et al.* (2005). Sensitivity analysis can be instructive, but it does not provide a criterion for choice with partial knowledge.

[2] Use of the term *ambiguity* to describe the absence of a basis for assertion of a subjective probability distribution appears to have originated in Ellsberg (1961). Keynes (1921) and Knight (1921) used the term *uncertainty*, and some modern authors refer to ambiguity as *Knightian uncertainty*. Other authors have used *ignorance* as a synonym for ambiguity (Arrow and Hurwicz, 1972), while still others refer to *robust Bayesian analysis* (Berger, 1985) or *imprecise probabilities* (Walley, 1991).

among unvaccinated persons weakly decreases as the vaccination rate rises. The planner observes the outcomes realized under a status quo policy that vaccinates an observed fraction of the population. I show that this empirical evidence, combined with the weak assumption of monotone external effectiveness, implies that certain vaccination rates are strictly dominated. That is, there exist alternative vaccination rates that yield lower social cost whatever the true external effectiveness of vaccination may be. I then show how the minimax or the minimax-regret criteria may be used to choose a reasonable undominated vaccination rate.

The analysis in this paper builds on parts of my previous work on planning under ambiguity. I have earlier studied the criminal-justice problem of choosing a search policy when the planner has partial knowledge of the deterrent effect of search on crime (Manski, 2006). The formal structure of the vaccination problem has much in common with that of the search problem, the substantive difference between the two notwithstanding. I have also previously considered vaccination as an example of a class of planning problems with social interactions, in a context where the planner assumes external effectiveness is monotone but has no empirical evidence on a status quo policy (Manski, 2009a, 2009b).

To begin, I pose in Section 2 a planning problem whose objective is to minimize the utilitarian social cost of illness and vaccination. Vaccination is costly, as is illness. Vaccination is beneficial to the extent that it prevents illness. The external effect of vaccination is expressed through an *external-response function* that describes how the illness rate of unvaccinated persons varies with the vaccination rate. I use the analytically simple case of a linear external-response function to illustrate how the optimal vaccination rate may depend on costs and the external-response function.

Section 3 examines the planning problem when the planner has only partial knowledge of the external-response function and, hence, is unable to determine the optimal policy. I suppose that the planner observes the illness rate of a study population whose vaccination rate has previously been chosen. The study population may, for example, be a past cohort of the population now about to be treated. The planner knows (or finds it credible to assume) that the study population and the population of interest have the same external-response function. He also knows that the external-response function weakly decreases as the vaccination rate increases. However, the planner does not know the magnitude of the external effect of vaccination.

In this setting, I first show how the planner can eliminate dominated vaccination rules, ones which are inferior whatever the actual external-response function may be. Broadly speaking, low (high) vaccination rates are dominated when the cost of vaccination is low (high); Lemmas 1 and 2 make this precise. I then show how the planner can use the minimax or minimax-regret criterion to choose an undominated vaccination rule. Both criteria are reasonable and tractable, but they yield different policies; Lemmas 3 and 4 derive their explicit forms. The proofs of all lemmas are given in an Appendix.

The scenario studied in Section 3 is realistic enough to demonstrate key ideas about vaccination under ambiguity, but is idealized enough to yield simple analytical findings. Section 4 considers several extensions that are more realistic but more complex. I show how the analysis of Section 3 extends to settings where vaccination has imperfect but known internal effectiveness. I discuss generalization of the planning problem to settings where population members have observable covariates. I consider provision of incentives for private vaccination when the planner cannot mandate vaccination. And I discuss dynamic planning problems where a planner vaccinates a sequence of cohorts, using observation of past outcomes to inform present decisions. Section 5 concludes.

2. Optimal Vaccination with a Known External-Response Function

As prelude to consideration of vaccination under ambiguity, this section specifies the optimization problem that the planner wants to solve and derives the solution in an illustrative case.

2.1. Basic Concepts

Suppose there exists a large population of observationally identical persons and that a planner must choose the population vaccination rate. Assuming that persons are observationally identical implies that the planner must choose persons to vaccinate at random. He cannot systematically vaccinate persons who are particularly susceptible to illness or infectious. Assuming that the population is large implies that the planner can ignore finite-sample statistical variation as an issue when comparing alternative vaccination rates.

Assume that vaccination always prevents a vaccinated person from becoming ill. Let p(t) be the external-response function, giving the fraction of unvaccinated persons who become ill when the vaccination rate is t. Then the fraction of the population who become ill is $p(t)(1 - t)$.

Suppose the planner wants to minimize a social cost function with two additive components. These are the harm caused by illness and the cost of vaccination. Let $a > 0$ denote the monetized mean harm caused by illness and let $c > 0$ denote the mean cost per vaccination. The social cost of vaccination rate t is

$$(1) \quad S(t) = ap(t)(1 - t) + ct.$$

The first term on the right-hand side gives the aggregate cost of illness, and the second gives the aggregate cost of vaccination. The planner wants to solve the problem $\min_{t \in [0, 1]} S(t)$.

The simple social cost function specified here expresses the core tension of vaccination policy: a higher vaccination rate is more effective in preventing illness but is more costly. Similar welfare functions have been assumed in some past research on optimal vaccination, such as Brito *et.al.* (1992). However, it has been more common to pose a susceptible-infectious-removed (SIR) or other dynamic model of disease transmission and to assume that the social objective is to keep the transmission rate below the threshold at which an epidemic occurs; see, for example, Ball and Lyne (2002) or Hill and Longini (2003). The objective of preventing onset of an epidemic differs from minimization of social cost.

2.2. Optimal Vaccination with a Linear External-Response Function

The planner's problem is solvable if the external-response function is known. Suppose it is linear, with $p(t) = \rho(1 - t)$ and $0 < \rho \leq 1$. Thus, the illness rate of unvaccinated persons is ρ if no one is vaccinated, and it decreases linearly to zero as the vaccination rate rises. Then the optimal vaccination rate is

$$(2)\quad t^* = \underset{t \in [0, 1]}{\operatorname{argmin}}\ a\rho(1 - t)^2 + ct.$$

The cost parameters being non-negative, the quadratic first term of the social cost function is minimized at $t = 1$and the linear second term at $t = 0$. The optimal vaccination rate must resolve this tension.

The social cost function has a unique minimum at the value of t that solves the first order condition

$$(3)\quad 0 = 2a\rho t - 2a\rho + c.$$

Solving (3), the minimum is at $1 - c/(2a\rho)$. Hence, the optimal vaccination rate is

$$(4)\quad \begin{aligned} t^* &= 0 && \text{if } 2a\rho < c. \\ &= 1 - c/(2a\rho) && \text{if } 2a\rho \geq c. \end{aligned}$$

Observe that for no value of the parameters (a, c, ρ) is it optimal to vaccinate the entire population. It may, however, be optimal to vaccinate no one. This occurs if the cost of vaccination, expressed by c, is higher than the benefit in preventing illness, expressed by a and ρ.

3. Vaccination with Partial Knowledge of The External-Response Function

3.1. Empirical Evidence and Maintained Assumptions

Solution of the planning problem of Section 2 is possible given full knowledge of the external-response function. However, this knowledge generally is unavailable in practice. Empirical evidence and credible assumptions may restrict the form of the external-response function, but they rarely if ever pin it down fully.

To demonstrate how a planner with partial knowledge of the external-response function may choose a vaccination policy, I consider decision making in a particular informational setting. I suppose that the planner observes the illness rate of a study population whose vaccination rate has previously been chosen. The study population may, for example, be a past cohort of the population now about to be treated. The planner thinks it credible to assume that the study population and the treatment population have the same external-response function. He also thinks it credible to assume that the illness rate of unvaccinated persons weakly decreases as the vaccination rate increases. However, he does not know the magnitude of the external effect of vaccination.

Let r denote the vaccination rate in the study population and let q denote the realized illness rate. Formally, I maintain

Assumption 1 (Study Population): The planner observes r and q. The planner knows that q = p(r).

Assumption 2(Vaccination Weakly Prevents Illness): The planner knows that p(t) is weakly decreasing in t.

Taken together, Assumptions 1 and 2 imply that

(5) $t \le r \;\Rightarrow\; p(t) \ge q,$

$t \ge r \;\Rightarrow\; p(t) \le q.$

This knowledge of the external response function is much weaker than the traditional assumption that the planner knows the function. Moreover, Assumptions 1 and 2 often are credible. It is often possible to observe the vaccination and illness rates of a study population, and reasonable to suppose that the study population and the population of interest have (at least approximately) the same external-response function. It also usually is reasonable to think that vaccination weakly prevents illness. Assumption 2 is a specific instance of the general idea of *monotone treatment response* developed in Manski (1997).

Observe that Assumptions 1 and 2 partially reveal the social cost function. The planner knows that $S(r) = aq(1 - r) + cr$ and that $S(1) = c$. He can bound $S(\cdot)$ at other vaccination rates. It follows from (5) that $aq(1 - t) + ct \le S(t) \le a(1 - t) + ct$ when $0 \le t < r$ and that $ct \le S(t) \le aq(1 - t) + ct$ when $r < t < 1$.

3.2. Dominated Vaccination Rates

How should a planner use Assumptions 1 and 2 to choose a policy? Although there is no one "correct" answer to this question, a planner clearly should not choose a vaccination rate that is inferior whatever the actual external-response function may be.

Let Γ be the set of feasible external-response functions under Assumptions 1 and 2. For $\gamma \in \Gamma$, let

(6) $S(t, \gamma) \;\equiv\; a\gamma(t)(1 - t) + ct$

be the social cost of vaccination rate t when the external-response function is γ. Rule t is *strictly dominated* if and only if there exists another vaccination rate t' such that $S(t, \gamma) > S(t', \gamma)$ for all $\gamma \in \Gamma$.

Let $t \in [0, 1]$ and $s \in [0, 1]$ designate feasible choices for the vaccination rate. Let

$$(7) \quad d(t, s; \gamma) \equiv a\gamma(t)(1 - t) + ct - a\gamma(s)(1 - s) - cs$$

be the difference in the social cost of rates t and s when the external-response function is γ. Rate s is strictly dominated if and only if there exists a t such that $d(t, s; \gamma) < 0$ for all $\gamma \in \Gamma$. Let $D(t, s) \equiv \sup_{\gamma \in \Gamma} d(t, s; \gamma)$. An easily verifiable sufficient condition for strict dominance is that $D(t, s) < 0$.

Lemma 1 evaluates $D(t, s)$. Then Lemma 2 uses the result to determine a set of dominated vaccination rates.

Lemma 1: Let Assumptions 1 and 2 hold. Then

(i) $t < s \le r \Rightarrow D(t, s) = a(1 - t) - aq(1 - s) + c(t - s)$,

(ii) $t < r < s \Rightarrow D(t, s) = a(1 - t) + c(t - s)$,

(iii) $s \le t < r \Rightarrow D(t, s) = (c - aq)(t - s)$,

(iv) $r \le t < s \Rightarrow D(t, s) = aq(1 - t) + c(t - s)$,

(v) $r \le s \le t \Rightarrow D(t, s) = c(t - s)$,

(vi) $s < r \le t \Rightarrow D(t, s) = (c - aq)(t - s)$. □

Lemma 2: Let Assumptions 1 and 2 hold. Vaccination rate s is strictly dominated if these conditions hold:

(a) Let $c < aq$. Then s is strictly dominated if $s < r$.

(b) Let $c > aq$. Then s is strictly dominated if $s > r + aq(1 - r)/c$.

(c) Let $c > a$. Then s is strictly dominated if $(a - aq)/(c - aq) < s \le r$ or if $s > \max(r, a/c)$. □

It might have been thought that Assumptions 1 and 2 are too weak to yield interesting dominance findings. However, Lemma 2 shows these assumptions have considerable power. The broad finding is that small (large) values of s are dominated when the vaccination cost c is sufficiently small (large). Parts (a) through (c) give the specifics. Part (a) shows that vaccination rates lower than the rate r observed in the study population are dominated if $c < aq$. Thus, when the cost parameters satisfy this inequality, the optimal vaccination rate cannot be smaller than the observed rate r. Part (b) shows that vaccination rates sufficiently larger than r are dominated when $c > aq$. Part (c) shows that other vaccination rates are dominated if $c > a$.

3.3. Vaccination to Minimize Expected Social Cost

Elimination of dominated vaccination rates takes the planner part way toward solution of the vaccination problem. The literature on decision theory does not provide a consensus prescription for a complete solution. Instead, it offers alternative criteria that ensure choice of an undominated alternative.

Particularly familiar to economists is the expected utility model. In the present setting, this recommends that the planner place a subjective distribution on $p(\cdot)$, say Ψ, and minimize subjective expected social cost. The social cost function (1) is linear in $p(\cdot)$, so subjective expected social cost is

$$(8) \quad E_\Psi[S(t)] = a\pi(t)(1 - t) + ct,$$

where $\pi(\cdot) \equiv E_\Psi[p(\cdot)]$ is the subjective mean of $p(\cdot)$. Thus, the planner acts as a pseudo-optimizer, using the expected external-response function π as if it were the actual external-response function p. For example, if $\pi(\cdot)$ is the linear function $\pi(\cdot) = \rho(1 - t)$, the planner solves the optimization problem of Section 2.

3.4. Minimax Vaccination

Minimization of subjective expected cost is sensible if a planner can substantiate his choice of π, but it has no special appeal otherwise.[3] A planner who cannot substantiate choice of π can reasonably apply the minimax or minimax-regret criterion to choose a vaccination rate. These criteria choose rates that, in different senses, perform uniformly well across all feasible external-response functions. The minimax criterion chooses a vaccination rate that minimizes the maximum social cost that may feasibly occur. The minimax-regret criterion chooses a rate that minimizes the maximum loss to welfare that results from not knowing the external-response function. I derive the minimax vaccination rate under Assumptions 1 and 2 here and the minimax-regret rate in Section 3.5. [4]

For each candidate vaccination rate t, compute the maximum social cost that can occur across all feasible external-response functions; that is, $\max_{\gamma \in \Gamma} S(t, \gamma)$. The minimax criterion selects the vaccination rate that minimizes this maximum social cost. Thus, the minimax vaccination rate solves the problem

$$(9) \quad \min_{t \in [0, 1]} \max_{\gamma \in \Gamma} a\gamma(t)(1 - t) + ct.$$

Lemma 3 derives the vaccination rate that solves this problem.

---

[3] Decision theorists have sometimes asserted pre-eminence for maximization of expected utility (minimization of expected cost here), asserting not only that a decision maker *might* use this decision criterion but that he *should* do so. Reference is often made to representation theorems deriving the expected utility criterion from consistency axioms on hypothetical choice behavior, famously von Neumann and Morgenstern (1944) and Savage (1954). However, the theorems of axiomatic decision theory are interpretative rather than prescriptive. See Manski (2009c).

[4] Minimization of expected cost, minimax, and minimax-regret are particularly well-known criteria for decision making with partial knowledge, but they are not the only ones that may warrant consideration. For example, a decision maker who feels able to assert a partial subjective distribution on the states of nature could minimize maximum expected cost or minimize maximum expected regret. These ideas have a long history in the literature on statistical decision theory, which refers to them as the Γ-maximin and Γ-minimax regret criteria (see Berger, 1985).

Lemma 3: Under Assumptions 1 and 2, the minimax vaccination rate is

$$
\text{(10)} \quad t^m = \begin{cases} 0 & \text{if } c > a \text{ and } a \le aq(1-r) + cr, \\ r & \text{if } c > a \text{ and } a \ge aq(1-r) + cr \\ & \text{or if } aq < c \le a, \\ \text{all } t \in [0, 1] & \text{if } c = aq \text{ and } q = 1, \\ \text{all } t \in [r, 1] & \text{if } c = aq \text{ and } q < 1, \\ 1 & \text{if } c < aq. \end{cases} \qquad \square
$$

Lemma 3 shows that the minimax vaccination rate generically takes one of the three values (0, r, 1), the only exception being when $c = aq$, which has multiple maximin rates. All else equal, the vaccination rate weakly increases with the cost parameter a and decreases with c. It weakly increases with the realized illness rate q if $c < a$ and decreases with q otherwise.

3.5. Minimax-Regret Vaccination

For each feasible external-response function $\gamma \in \Gamma$, let $S^*(\gamma) \equiv \min_{t \in [0, 1]} S(t, \gamma)$ denote the smallest social cost achievable when the external-response function is γ. The regret of vaccination rate t in state of nature γ is $S(t, \gamma) - S^*(\gamma)$. Thus, regret measures the difference between the social cost delivered by rate t and that delivered by the best possible rate. The minimax-regret criterion selects the vaccination rate that minimizes maximum regret across all states of nature. Thus, the minimax-regret vaccination rate solves

$$
\text{(11)} \quad \min_{t \in [0, 1]} \ \sup_{\gamma \in \Gamma} S(t, \gamma) - S^*(\gamma)
$$

$$= \min_{t \in [0,1]} \sup_{\gamma \in \Gamma} \left\{a\gamma(t)(1-t) + ct - \min_{s \in [0,1]} [a\gamma(s)(1-s) + cs]\right\}$$

$$= \min_{t \in [0,1]} \sup_{\gamma \in \Gamma,\, s \in [0,1]} a\gamma(t)(1-t) + ct - a\gamma(s)(1-s) - cs$$

$$= \underset{t \in [0,1]}{\operatorname{argmin}} \max_{s \in [0,1]} D(t, s),$$

where D(t, s) was defined in Section 3.2. Lemma 4 derives the vaccination rate that solves this problem.

Lemma 4: Let Assumptions 1 and 2 hold.

(a) Let $c \le aq$. Then the minimax-regret vaccination rate is

$$(12) \quad t^{mr} = (aq + cr)/(aq + c).$$

(b) Let $c > aq$. Then the minimax-regret vaccination rate is

$$(13) \quad t^{mr} = \underset{t \in [0,1]}{\operatorname{argmin}}\ 1[t < r]\cdot\{\max [(a - aq)(1 - t), a(1 - t) + c(t - r), (c - aq)t]\}$$
$$+ 1[t \ge r]\cdot\{\max [aq(1 - t), c(t - r), (c - aq)t]\}.$$

If $r = 0$, $t^{mr} = aq/(aq + c)$. If $r = 1$, $t^{mr} = (a - aq)/[(a - aq) + (c - aq)]$. □

Lemma 4 shows that, as the cost c of vaccination increases from 0 to aq, the minimax-regret vaccination rate decreases continuously from 1 to $(1 + r)/2$. In contrast, Lemma 3 showed that the minimax vaccination rate equals 1 whenever $c \le aq$.

When $c > aq$, the solution to the minimax-regret problem generally does not have an explicit form of simplicity comparable to (12). However, the abstract finding (13) simplifies in two important polar cases.

Suppose that no one was vaccinated in the study population; thus, $r = 0$. Then t^{mr} decreases from ½ to 0 as c increases from aq to ∞. Now suppose that the entire study population was vaccinated; thus, $r = 1$. Then t^{mr} decreases from 1 to 0 as c increases from aq to ∞. [5]

3.6. Numerical Examples

Numerical examples are useful to illustrate the findings of Sections 2 and 3. I give three here, each modifying the preceding example in some respect.

First consider a scenario where the mean cost of vaccination is $c = 50$ dollars and the monetized mean cost of illness is $a = 1000$ dollars. The planner observes a study population with no vaccination ($r = 0$) and with illness rate $q = 1/5$. In this setting, a planner who believes the external-response function is linear would set $\rho = 1/5$ and use equation (4) to choose the vaccination rate $t^* = 7/8$. A planner who only knows the function to be weakly decreasing would not be able to use Lemma 2 to conclude that any vaccination rates are dominated, because $c < aq$ and $r = 0$. By Lemma 3, the minimax vaccination rate is $t^m = 1$. By Lemma 4, equation (12), the minimax-regret rate is $t^{mr} = 4/5$.

Next revise the scenario by supposing that the planner observes a study population where $r = 1/2$ and $q = 1/10$. Continue to assume that $c = 50$ and $a = 1000$. A planner who believes the external-response function is linear would still set $\rho = 1/5$ and choose $t^* = 7/8$. A planner who only knows the function to be weakly decreasing can use Lemma 2 to conclude that any vaccination rate smaller than 1/2 is strictly dominated. By Lemma 3, the minimax vaccination rate remains $t^m = 1$. By Lemma 4, equation (12), the minimax-regret rate is $t^{mr} = 5/6$.

Now revise the scenario again by supposing that vaccination is more costly, say $c = 250$. Continue to assume that $a = 1000$, $r = 1/2$, and $q = 1/10$. In this case, a planner who believes the external-response

---

[5] Manski (2009a, 2009b) give an alternative statement and proof of the results for $r = 0$ and $r = 1$.

function is linear would choose $t^* = 3/8$. A planner who only knows the function to be weakly decreasing can use Lemma 2 to conclude that any vaccination rate larger than 7/10 is strictly dominated. By Lemma 3, the minimax vaccination rate is $t^m = 1/2$. By Lemma 4, equation (13), the minimax-regret rate is $t^{mr} = 1/2$ as well.

4. Extensions of the Analysis

The analysis of Section 3 developed key ideas about vaccination under ambiguity in a relatively simple scenario. The findings may be used to form vaccination policy when a public health agency has the authority to mandate treatment with an internally effective vaccine, observes the outcome of a status quo policy, and is reluctant to assume more than that the external-response function is weakly decreasing.

There are innumerable variations on this scenario that may warrant attention. I discuss here four extensions of the analysis. I first consider vaccination with imperfect internal effectiveness (Section 4.1), next stratification of the population when its members have observable covariates (Section 4.2), then provision of incentives when it is infeasible to mandate vaccination (Section 4.3), and finally dynamic choice of vaccination rates for a sequence of cohorts (Section 4.4).

4.1. Vaccination with Imperfect Internal Effectiveness

I have assumed that vaccination always prevents a vaccinated person from becoming ill. This assumption is realistic for some vaccines but not for others. Suppose that the vaccine under study is internally effective with rate $\lambda \in (0, 1]$, and confer no immunity when administered to the remaining fraction $1 - \lambda$ of vaccinated persons. Then equation (1) does not give the social cost of vaccination. Instead, the social cost is

(14) $S(t) = ap(\lambda t)(1 - \lambda t) + ct.$

The cost per vaccination remains c. However, the fraction of the population who are effectively vaccinated is now λt rather than t.

When the planner has partial knowledge of both $p(\cdot)$ and λ, analysis of vaccination under ambiguity is beyond the scope of this paper. When λ is known, most of the analysis of Section 3 applies to the new planning problem with minor modification. I explain below.

Letting $\tau \equiv \lambda t$ be the *effective vaccination rate*, we may rewrite cost function (14) as

(15) $S'(\tau) = ap(\tau)(1 - \tau) + (c/\lambda)\tau.$

Solving the problem $\min_{\tau \in [0, \lambda]} S'(\tau)$ yields the optimal effective vaccination rate, say τ^*. This done, the optimal raw vaccination rate is $t^* = \tau^*/\lambda$. Similarly, we may use the reasoning of Section 3 to determine dominated values of the effective vaccination rate as well as the minimax and minimax-regret effective rates. The findings may then be divided by λ to obtain the corresponding raw vaccination rates.

If τ could range over the entire unit interval, the analysis of Section 3 would apply to the problem $\min_{\tau \in [0, 1]} S'(\tau)$, with the effective vaccination cost c/λ replacing c and with the study population's effective vaccination rate λr replacing its raw rate r. We need only determine whether the findings for this problem continue to hold when the domain of τ is restricted to $[0, \lambda]$.

Consider strict dominance. The proof to Lemma 2 shows that whenever a vaccination rate is dominated, there exists a dominating rate whose value is less than or equal to r. Applied to the new planning problem, this means that when some effective rate σ is dominated, there exists a dominating effective rate

$\tau \leq \lambda r$. It follows that when σ is dominated, the corresponding raw rate $s = \sigma/\lambda$ is dominated by $t = \tau/\lambda$.

Consider the minimax criterion. Adaptation of the proof to Lemma 3 shows that the minimax effective vaccination rate is

$$
\begin{aligned}
(16)\quad \tau^{m} &= 0 && \text{if} \quad c/\lambda \geq a \text{ and } a \leq aq(1 - \lambda r) + cr, \\
&= \lambda r && \text{if} \quad c/\lambda \geq a \text{ and } a \geq aq(1 - \lambda r) + cr \\
& && \text{or if} \quad aq \leq c/\lambda < a, \\
&= \lambda && \text{if} \quad c/\lambda \leq aq.
\end{aligned}
$$

Hence, the minimax raw rate is τ^{m}/λ.

Consider the minimax-regret criterion. The proof to Lemma 4 is more complex than those of the other Lemmas, and I will not attempt a full adaptation here. However, a simple result emerges when application of Lemma 4 to cost function $S'(\cdot)$ with τ permitted to range over the interval [0, 1] yields a solution $\tau^{mr} \leq \lambda$. Then this solution is a feasible value of the effective vaccination rate. In these cases, the minimax-regret value of the raw vaccination rate is τ^{mr}/λ.

4.2. Vaccination of a Population with Observable Covariates

I have assumed that the members of the population are observationally identical. This does not mean that persons actually are identical, only that the planner cannot distinguish them.

Suppose now that the planner observes some health-relevant covariates for each person, say x taking values in a covariate space X. Thus, suppose that persons with the same value of x are observationally identical, but the planner can distinguish persons with different values of x. Then the planner may want to choose vaccination rates that vary with x.

To formalize the planning problem, let the social cost function be

$$(17) \quad S(t_x, x \in X) = \sum_{x \in X} a_x p_x(t_w, w \in X)(1 - t_x) + c_x t_x.$$

The mean harm per illness, mean vaccination cost, and the external-response function may all vary with x. The external effect of vaccination on unvaccinated persons with covariates x may depend on the entire vector $(t_w, w \in X)$ of vaccination rates. The planner wants to choose $(t_x, x \in X)$ to minimize social cost.

The analysis of Section 3 applies directly to this planning problem if the planner assumes that the external effect of vaccination occurs only within groups and not between groups; that is, if the social cost function has the form

$$(18) \quad S(t_x, x \in X) = \sum_{x \in X} a_x p_x(t_x)(1 - t_x) + c_x t_x,$$

where $p_x(\cdot)$ is a function only of t_x, not of $(t_w, w \in X)$. This assumption may be credible if covariate x indexes sub-populations who do not come into physical contact with one another.

Cost function (18) is separable in x. Hence, the planner may treat each x-specific sub-population of persons separately, as if were the entire population. Doing so yields the desired result for the population as a whole. For example, the vector of vaccination rates that minimizes population-wide maximum regret may be obtained by solving the minimax-regret problem for each sub-population separately.

When cost function (17) does not have the separable structure (18), new analysis is required to determine what vectors of vaccination rates are dominated, and what vectors solve the population-wide minimax and minimax-regret problems. The natural extensions of Assumptions 1 and 2 to settings with observable covariates are

Assumption 1′ (Study Population): For $x \in X$, the planner observes (r_x, q_x) and knows that $q_x = p_x(r_w, w \in X)$.

Assumption 2′ (Vaccination Weakly Prevents Illness): For $x \in X$, the planner knows that $p_x(t_w, w \in X)$ is weakly decreasing in $(t_w, w \in X)$.

Taken together, Assumptions 1′ and 2′ imply that, for each $x \in X$,

$$(19) \quad t_w \le r_w, w \in X \;\Rightarrow\; p_x(t_w, w \in X) \ge q_x,$$
$$t_w \ge r_w, w \in X \;\Rightarrow\; p_x(t_w, w \in X) \le q_x.$$

Observe that Assumptions 1′ and 2′ reveal nothing about $p_x(t_w, w \in X)$ when $t_w < r_w$ for some $w \in X$ and $t_w > r_w$ for other w. Hence, a planner who observes covariates and wants to let the vaccination rate vary across sub-populations must cope with more ambiguity than one who can only choose a uniform population-wide vaccination rate.

4.3. Provision of Incentives for Private Vaccination

I have assumed that the planner has the power to mandate vaccination of the population. This assumption is realistic in some settings, such as vaccination of health care workers, military personnel, or students in public schools. However, it is more common for vaccination to be a private decision, which a public health agency may seek to influence through provision of incentives.

Suppose that a planner selects an incentive policy from a set D of feasible policies. Given a policy, members of the population individually choose whether or not to be vaccinated. Then the resulting social cost may depend not only on the fraction of the population who choose to be vaccinated but also on the composition of the vaccinated group. The reason is that members of the population may vary in their susceptibility to illness and in the extent to which they can infect other members of the population. Hence, the effectiveness of an incentive policy may depend on which as well as how many persons choose to be vaccinated. Public health agencies typically have limited knowledge of how incentive policies affect private vaccination choices and of the resulting implications for disease transmission. Hence, a planner choosing an incentive policy typically faces more ambiguity than does one who mandates vaccination.

General analysis of choice among incentive policies is beyond the scope of this paper. However, the analysis of Section 3 does apply to a special (albeit not very realistic) case. Suppose that the planner knows the fraction t(d) of the population who choose to be vaccinated under each policy $d \in D$. Suppose as well that the distribution of susceptibility and infectiousness among persons who choose to be vaccinated under each policy is the same as within the population as a whole. Then choice of an incentive policy from D is equivalent to choice of a vaccination rate from the set $[t(d), d \in D]$. The analysis of Section 3 applies directly if the policy set D is sufficiently rich that $[t(d), d \in D] = [0, 1]$. If $[t(d), d \in D]$ is a proper subset of [0, 1], then the analysis applies with the caveat that the feasible vaccination rates are $[t(d), d \in D]$ rather than all elements of the unit interval.

4.4. Vaccination of a Sequence of Cohorts

I have assumed that the planner observes outcomes in one study population and chooses a vaccination rate for one treatment population. Often, a planner chooses vaccination rates for a sequence of cohorts. Examples are vaccination of successive cohorts of military recruits or students entering kindergarten. In such cases, outcome data may be observed for multiple past cohorts and used to inform choice of vaccination rates for later cohorts.

Vaccination of a sequence of cohorts differs in two respects from the planning problem studied in this paper, one respect increasing the complexity of the analysis and the other being more fundamental. Observation of outcomes in multiple past cohorts increases the complexity of the analysis relative to the case where only one study population is observed. When vaccination rates vary across past cohorts, the planner knows more about the external-response function than was the case in Section 3. Determination of dominated vaccination rates, the minimax rate, and the minimax-regret rate does not differ conceptually from the analysis of Section 3, but these tasks become increasingly laborious as the number of past cohorts increases. It would be useful to develop analytical or numerical methods that ease the derivations.

The more fundamental difference between sequential planning problems and the static problem studied here is that a forward-looking planner may want to choose a vaccination rate for the present cohort that enhances the information available for treatment of future cohorts. The minimax and minimax-regret vaccination rates derived in Section 3 presumed that the planner is concerned only with the welfare of the current cohort. These need not be the minimax and minimax-regret rates when the planner wants to optimize the combined welfare of the current and future cohorts. To enhance combined welfare, it may be better to choose a current vaccination rate that is not desirable for the current cohort but that reduces ambiguity about the external-effectiveness function, enabling improved treatment of future cohorts. For example, a planner who wants to minimize the maximum combined social cost for the current and future cohorts may want to choose a current vaccination rate other than (0, r, 1), which Lemma 3 showed to be the only possible solutions to the static minimax problem. Development of reasonable strategies for dynamic vaccination under ambiguity is a wide-open question that warrants serious attention.

5. Conclusion

Vaccination policy has regularly been studied under the assumption that the planner knows the internal and external response of illness to vaccination. Consideration of vaccination with partial knowledge has been considered only recently, and then only under the assumption that the planner can credibly place a subjective probability distribution on the response function.

This paper has shown how choice of a vaccination rate may be addressed as a problem of planning under ambiguity. To demonstrate general ideas in a simple setting, I developed a planning problem where a planner has partial knowledge of the external-response function, which expresses how the illness rate of unvaccinated persons varies with the vaccination rate. I supposed that the planner observes the illness rate of a study population whose vaccination rate has previously been chosen. I supposed that the planner knows the illness rate of unvaccinated persons weakly decreases as the vaccination rate increases, but he does not know the magnitude of the preventive effect of vaccination. Under these assumptions, I showed how the planner can eliminate dominated vaccination rates. I then showed how he can use the minimax or minimax-regret criterion to choose an undominated rate.

Many other scenarios of vaccination with partial knowledge warrant attention. This paper has discussed some of them, giving a rich agenda for future research.

Appendix: Proofs of the Lemmas

Lemma 1

(i) $t < s \leq r \Rightarrow \gamma(t) \geq \gamma(s) \geq q$. Maximization over Γ occurs by setting $\gamma(t) = 1$ and $\gamma(s) = q$. This gives $D(t, s) = a(1 - t) + ct - aq(1 - s) - cs$.

(ii) $t < r < s \Rightarrow \gamma(t) \geq q \geq \gamma(s)$. Maximization over Γ occurs by setting $\gamma(t) = 1$ and $\gamma(s) = 0$. This gives $D(t, s) = a(1 - t) + ct - cs$.

(iii) $s \leq t < r \Rightarrow \gamma(s) \geq \gamma(t) \geq q$. Maximization over Γ occurs by setting $\gamma(t) = \gamma(s) = \delta$ for some $\delta \geq q$. This gives $D(t, s) = (c - a\delta)(t - s)$. Given that $a > 0$ and $t \geq s$, the maximum is attained by setting $\delta = q$. Hence, $D(t, s) = (c - aq)(t - s)$.

(iv) $r \leq t < s \Rightarrow q \geq \gamma(t) \geq \gamma(s)$. Maximization over Γ occurs by setting $\gamma(t) = q$ and $\gamma(s) = 0$. This gives $D(t, s) = aq(1 - t) + ct - cs$.

(v) $r < s \leq t \Rightarrow q \geq \gamma(s) \geq \gamma(t)$. Maximization over Γ occurs by setting $\gamma(t) = \gamma(s) = \delta$ for some $\delta \leq q$. This gives $D(t, s) = (c - a\delta)(t - s)$. Given that $a > 0$ and $t \geq s$, the maximum is attained by setting $\delta = 0$. Hence, $D(t, s) = c(t - s)$.

(vi) $s \leq r \leq t \Rightarrow \gamma(s) \geq q \geq \gamma(t)$. Maximization over Γ occurs by setting $\gamma(t) = \gamma(s) = q$. This gives $D(t, s) = (c - aq)(t - s)$.

Q. E. D.

Lemma 2

(a) Part (vi) of Lemma 1 showed that if $s < r \leq t$, then $D(t, s) = (c - aq)(t - s)$. Hence, $D(t, s) < 0$ for all such (t, s).

(b) Part (iv) of Lemma 1 showed that if $r \leq t < s$, then

$$D(t, s) = aq(1 - t) + c(t - s) = aq + (c - aq)t - cs.$$

Consider the right hand side as a function of t. The function is minimized at $t = r$, giving $D(r, s) = aq + (c - aq)r - cs$. If $s > r + aq(1 - r)/c$, then $D(r, s) < 0$.

(c) Part (i) of Lemma 1 showed that if $t < s \leq r$, then

$$D(t, s) = a(1 - t) - aq(1 - s) + c(t - s) = a(1 - q) + (c - a)t - (c - aq)s.$$

Consider the right-hand side as a function of t. The function is minimized at $t = 0$, giving $D(0, s) = a(1 - q) - (c - aq)s$. If $s > (a - aq)/(c - aq)$, then $D(0, s) < 0$.

Part (ii) of Lemma 1 showed that if $t < r < s$, then

$$D(t, s) = a(1 - t) + c(t - s) = a + (c - a)t - cs.$$

Consider the right hand side as a function of t. The function is minimized at $t = 0$, giving $D(0, s) = a - cs$. If $s > a/c$, then $D(0, s) < 0$.

Q. E. D.

Lemma 3

For each value of t, the inner maximization problem in (9) is solved by setting the illness rate to its largest feasible value; that is, $\gamma(t) = 1[t < r] + q{\cdot}1[t \geq r]$. Hence,

$$t^m \equiv \underset{t \in [0,\, 1]}{\text{argmin}} \quad a\{1[t < r] + q{\cdot}1[t \geq r]\}(1 - t) + ct.$$

To solve this problem, I first consider the domains $t < r$ and $t \geq r$ separately, and then combine them.

First let $t < r$. The minimization problem is $\min_{t < r}\ a(1 - t) + ct$. If $c > a$, the solution is $t = 0$ and the minimax value is a. If $c = a$, all $t < r$ are solutions and the minimax value is $c = a$. If $c < a$, the criterion function decreases as $t \to r$, with limit value $a(1 - r) + cr$.

Next let $t \geq r$. Now the minimization problem is $\min_{t \geq r}\ aq(1 - t) + ct$. If $c > aq$, the solution is $t = r$ and the minimax value is $aq(1 - r) + cr$. If $c = aq$, all $t \geq r$ are solutions and the minimax value is $c = aq$. If $c < aq$, the solution is $t = 1$ and the minimax value is c.

Now combine the two domains. If $c \geq a$, the solution on $t \in [0, 1]$ is $t = 0$ if $a \leq aq(1 - r) + cr$ and is $t = r$ if $a \geq aq(1 - r) + cr$. If $aq < c < a$, the solution is $t = r$. If $c = aq$, all $t \geq r$ are solutions. If $c < aq$, the solution is $t = 1$.

Q. E. D.

Lemma 4

Recall that Lemma 1 derived D(t, s). In each of parts (a) and (b), I first consider $t < r$, maximize D(t, s) over $s \in [0, 1]$ to obtain maximum regret for a specified t, and then find the value of t that minimizes maximum regret. I next do the same for $t \geq r$. Finally I combine the findings.

(a). Let $c \leq aq$. Let $t < r$. There are three cases to consider:

(i) $\max_{s:\, t < s \leq r} D(t, s)$ occurs at s = r. Hence, $\max_{s:\, t < s \leq r} D(t, s) = a(1 - t) - aq(1 - r) + c(t - r)$.

(ii) $\sup_{s:\, t < r < s} D(t, s)$ occurs at s = r. Hence, $\sup_{s:\, t < r < s} D(t, s) = a(1 - t) + c(t - r)$.

(iii) $\max_{s \leq t < r} D(t, s)$ occurs at s = t. Hence, $\max_{s \leq t < r} D(t, s) = 0$.

The supremum in (ii) exceeds the maxima in (i) and (iii). Hence, $\sup_{s \in [0, 1]} D(t, s) = a(1 - t) + c(t - r)$.

Minimization over t < r of the expression $a(1 - t) + c(t - r)$ yields the minimax-regret vaccination rate within this restricted range of rates. Given that $c \leq a$, the infimum occurs at t = r. Hence, the restricted minimax-regret value is $a(1 - r)$.

Now let $t \geq r$. There are again three cases to consider:

(iv) $\sup_{s:\, r \leq t < s} D(t, s)$ occurs at s = t. Hence, $\sup_{s:\, r \leq t \leq s} D(t, s) = aq(1 - t)$.

(v) $\sup_{s:\, r < s \leq t} D(t, s)$ occurs at s = r. Hence, $\sup_{s:\, r < s \leq t} D(t, s) = c(t - r)$.

(vi) $\max_{s:\, s \leq r \leq t} D(t, s)$ occurs at s = r. Hence, $\max_{s:\, s \leq r \leq t} D(t, s) = (c - aq)(t - r)$.

The supremum in (iii) is non-positive. Hence, $\sup_{s \in [0, 1]} D(t, s) = \max[aq(1 - t), c(t - r)]$.

Minimization over $t \geq r$ of $\max[aq(1 - t), c(t - r)]$ yields the minimax-regret vaccination rate within this restricted range of rates. Expression $aq(1 - t)$ falls from $aq(1 - r)$ to 0 as t rises from r to 1. Expression $c(t - r)$ rises from 0 to $c(1 - r)$ as t rises from r to 1. Hence, $\max[aq(1 - t), c(t - r)]$ is minimized when t solves the equation $aq(1 - t) = c(t - r)$. The solution is $t = (aq + cr)/(aq + c)$. The minimax-regret value is $caq(1 - r)/(aq + c)$.

Finally, compare the minimax-regret values over the two ranges t < r and $t \geq r$. The latter is smaller than the former. Hence, $(aq + cr)/(aq + c)$ is the overall minimax-regret vaccination rate.

(b) Let c > aq. Let t < r. There are three cases to consider:

(i) $\max_{s:\, t < s \leq r} D(t, s)$ occurs at s = t. Hence, $\max_{s:\, t < s \leq r} D(t, s) = (a - aq)(1 - t)$.

(ii) $\sup_{s:\, t < r < s} D(t, s)$ occurs at s = r. Hence, $\sup_{s:\, t < r < s} D(t, s) = a(1 - t) + c(t - r)$.

(iii) $\max_{s \leq t < r} D(t, s)$ occurs at s =0. Hence, $\max_{s \leq t < r} D(t, s) = (c - aq)t$.

Hence, maximum regret for $t < r$ is $\max [(a - aq)(1 - t), a(1 - t) + c(t - r), (c - aq)t]$.

Now let $t \geq r$. There are again three cases to consider:

(iv) $\sup_{s: r \leq t < s} D(t, s)$ occurs at $s = t$. Hence, $\sup_{s: r \leq t \leq s} D(t, s) = aq(1 - t)$.

(v) $\sup_{s: r < s \leq t} D(t, s)$ occurs at $s = r$. Hence, $\sup_{s: r < s \leq t} D(t, s) = c(t - r)$.

(vi) $\max_{s: s \leq r \leq t} D(t, s)$ occurs at $s = 0$. Hence, $\max_{s: s \leq r \leq t} D(t, s) = (c - aq)t$.

Hence, maximum regret for $t \geq r$ is $\max [aq(1 - t), c(t - r), (c - aq)t]$.

Combining the above findings, the minimax-regret vaccination rate is

$$t^{mr} = \underset{t \in [0, 1]}{\operatorname{argmin}}\ 1[t < r]\cdot\{\max [(a - aq)(1 - t), a(1 - t) + c(t - r), (c - aq)t]\} + 1[t \geq r]\cdot\{\max [aq(1 - t), c(t - r), (c - aq)t]\}.$$

If $r = 0$, the general expression for t^{mr} reduces to

$$t^{mr} = \underset{t \in [0, 1]}{\operatorname{argmin}}\ \{\max [aq(1 - t), ct]\}.$$

This problem has unique solution $aq/(aq + c)$. If $r = 1$, the general expression for t^{mr} reduces to

$$t^{mr} = \underset{t \in [0, 1]}{\operatorname{argmin}}\ 1[t < 1]\cdot\{\max [(a - aq)(1 - t), (c - aq)t]\} + 1[t = 1]\cdot(c - aq)\}.$$

This problem has unique solution $(a - aq)/[(a - aq) + (c - aq)]$.

Q. E. D.

References

Arrow, K. and L. Hurwicz (1972), “An Optimality Criterion for Decision-Making Under Ignorance,” in Carter, D. and J. Ford (editors), *Uncertainty and Expectations in Economics*, Oxford: Blackwell.

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.

Berger, J. (1985), *Statistical Decision Theory and Bayesian Analysis*, New York: Springer-Verlag.

Boulier, B., T. Datta, and R. Goldfarb (2007), “Vaccination Externalities,” *The B.E. Journal of Economic Analysis & Policy*, 7, issue 1, article 23. http://www.bepress.com/bejeap/vol7/iss1/art23.

Brito, D., E. Sheshinski, and M. Intriligator (1991), “Externalities and Compulsory Vaccinations,” *Journal of Public Economics*, 45, 69-90.

Ellsberg, D. (1961), “Risk, Ambiguity, and the Savage Axioms,” *Quarterly Journal of Economics*, 75, 643-669.

Hill, A. and I. Longini (2003), “The Critical Vaccination Fraction for Heterogeneous Epidemic Models,” *Mathematical Biosciences*, 181, 85–106.

Keynes, J.(1921), *A Treatise on Probability*, London: MacMillan.

Knight, F. (1921), *Risk, Uncertainty, and Profit*, Boston: Houghton-Mifflin.

Manski, C. (1997), “Monotone Treatment Response,” *Econometrica*, 65, 1311-1334.

Manski, C. (2000), “Identification Problems and Decisions Under Ambiguity: Empirical Analysis of Treatment Response and Normative Analysis of Treatment Choice,” *Journal of Econometrics*, 95, 415-442.

Manski, C. (2002), “Treatment Choice Under Ambiguity Induced by Inferential Problems,” *Journal of Statistical Planning and Inference*, 105, 67−82.

Manski, C. (2004), “Statistical Treatment Rules for Heterogeneous Populations,” *Econometrica*, 72, 1221-1246.

Manski, C. (2005), *Social Choice with Partial Knowledge of Treatment Response*, Princeton: Princeton University Press.

Manski, C. (2006), “Search Profiling with Partial Knowledge of Deterrence.” *Economic Journal*, 116, F385-F401.

Manski, C. (2007a), "Minimax-Regret Treatment Choice with Missing Outcome Data," *Journal of Econometrics*, 139, 105-115.

Manski, C. (2007b), *Identification for Prediction and Decision*, Cambridge, Mass.: Harvard University Press.

Manski, C. (2009a), "Social Planning with Partial Knowledge of Social Interactions," *Journal of Institutional and Theoretical Economics*, 165, 146-158.

Manski, C. (2009b), "Diversified Treatment under Ambiguity," *International Economic Review*, 50, 1013-1041.

Manski, C. (2009c), "Actualist Rationality," *Theory and Decision*, forthcoming.

Manski, C. and A. Tetenov (2007), "Admissible Treatment Rules for a Risk-Averse Planner with Experimental Data on an Innovation," *Journal of Statistical Planning and Inference*, 137, 1998-2010.

Mirrlees, J. (1971), "An Exploration in the Theory of Optimal Income Taxation," *Review of Economic Studies*, 38, 175-208.

Nordhaus, W. (2008), *A Question of Balance: Weighing the Options on Global Warming Policy*, New Haven, CT: Yale University Press.

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.

Savage, L. (1954), *The Foundations of Statistics*, New York: Wiley.

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

Tanner, M., L. Sattenspiel, and L. Ntaimo (2008), "Finding Optimal Vaccination Strategies Under Parameter Uncertainty Using Stochastic Programming," *Mathematical Biosciences*, 215, 144-151.

von Neumann, J. and O. Morgenstern (1944), *Theory of Games and Economic Behavior*, Princeton: Princeton University Press.

Wald, A. (1950), *Statistical Decision Functions*, New York: Wiley.

Walley. P. (1991), *Statistical Reasoning with Imprecise Probabilities*, London: Chapman & Hall.

Weycker, D., J. Edelsberg, E. Halloran, I. Longini, A. Nizam, V. Ciuryla, and G. Oster (2005), "Population-wide Benefits of Routine Vaccination of Children Against Influenza," *Vaccine*, 23, 1284-1293.