

Vaccination with partial knowledge of external effectiveness

Charles F. Manski¹

Department of Economics and Institute for Policy Research, Northwestern University, Evanston, IL 60208

This contribution is part of the special series of Inaugural Articles by members of the National Academy of Sciences elected in 2009.

Contributed by Charles F. Manski, December 30, 2009 (sent for review October 26, 2009)

Economists studying public policy have generally assumed that the relevant 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 been chosen previously. 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.

partial identification | planning under ambiguity | social interactions | vaccination policy

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 (1). Economists have also studied optimal criminal justice systems under the assumption that the planner knows how policing and sanctions affect crime rates (2). Researchers studying optimal vaccination against infectious disease have typically assumed the planner knows how vaccination affects illness rates (3–9).

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 becoming ill or infectious. Second, the planner may only partially know the *external* effectiveness of vac-

ination 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 nonnegligible 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 computational 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 unknown decision-relevant quantities (also known as states of nature) and propose choice of an action that maximizes subjective expected welfare. For example, Nordhaus used this approach to express partial knowledge of parameter values in his recent assessment of global warming policy (10), writing (p. 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 et al. exposit the model in ref. 11, under the name *stochastic programming*. They admonish vaccination researchers about the need to recognize uncertainty in policy evaluation, writing (pp. 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.”[†]

Author contributions: C.F.M. designed research, performed research, and wrote the paper. The author declares no conflict of interest.

¹E-mail: cfmanski@northwestern.edu.

[†]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, ref. 12. Sensitivity analysis can be instructive, but it does not provide a criterion for choice with partial knowledge.

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 not have a credible basis for asserting one. I have previously studied problems of this type, where a planner faces *ambiguity*.[‡] I have mainly focused on ambiguity arising from partial identification of policy impacts (19–26). I have also used the Wald formulation of statistical decision theory (27), to examine policy choice using sample data; see refs. 21, 28, 29.

Here I study choice of vaccination policy under ambiguity. To show concretely how basic principles may be applied, I consider in Section 2 a relatively simple scenario in which a planner must choose a vaccination rate for a population of observationally identical persons. The planner knows that vaccination is fully effective internally, always preventing a vaccinated person from becoming ill. However, he has no knowledge of the external effectiveness of vaccination except that it is monotone in the vaccination rate. That is, the rate of illness 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 only 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 and the minimax-regret criteria may be used to choose reasonable undominated vaccination rates. These criteria choose rates that, in different senses, perform uniformly well over all feasible states of nature. The former chooses an action that minimizes maximum cost across all states. The latter chooses an action that minimizes maximum regret across all states—the regret of a specified vaccination rate in a given state of nature is the cost of this rate minus the cost of the best possible rate. Both criteria protect the planner from poor outcomes, but in different formal ways.

I only refer to minimax and minimax-regret as “reasonable” decision criteria because there is no uniquely correct way to choose among undominated actions. After all, the crux of the problem in decision making under ambiguity is that the planner does not know which action is best. Wald, who studied minimax in abstraction in ref. 27, wrote (p. 18): “A minimax solution seems, in general, to be a reasonable solution of the decision problem.” Savage, who introduced the minimax-regret criterion in ref. 30, differed with Wald, writing (p. 63): “Application of the minimax rule . . . is indeed ultra-pessimistic; no serious justification for it has ever been suggested.” Savage observed with favor that the minimax-regret criterion is not similarly “ultra-pessimistic.” However, he later abandoned interest in minimax-regret and strongly endorsed the expected utility criterion instead (31).[§]

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 policy of search for evidence of

crime when the planner has partial knowledge of the deterrent effect of search on the rate of crime commission (22). 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 (25, 26).

To begin, I pose in Section 1 a planning problem whose objective is to minimize the utilitarian social cost of illness and vaccination. Vaccination is socially 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 2 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 rates, 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); Lemma 1 and Proposition 2 make this precise. After a brief discussion of minimization of expected cost, I show how the planner can use the minimax or minimax-regret criterion to choose an undominated vaccination rate. These criteria yield different policies; Propositions 3 and 4 derive their explicit forms. The proofs of all results are given in the *Appendix*.

The scenario studied in Section 2 is realistic enough to demonstrate key ideas about vaccination under ambiguity, but is idealized enough to yield simple analytical findings. Section 3 considers several extensions that are more realistic but more complex. I show how the analysis of Section 2 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 4 concludes.

1. 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.

1.1. Basic Concepts. Let there exist a large population of observationally identical persons. Suppose that a planner must choose the vaccination rate for this population. When persons are observationally identical, the planner must choose persons to vaccinate at random. He cannot systematically vaccinate persons who are particularly susceptible or infectious.

[‡]Use of the term *ambiguity* to describe the absence of a basis for assertion of a subjective probability distribution appears to have originated in ref. 13. The term *uncertainty* was used in refs. 14, 15, and some modern authors refer to ambiguity as *Knightian uncertainty*. Other authors have used *ignorance* as a synonym for ambiguity (16), whereas still others refer to *robust Bayesian analysis* (17) or *imprecise probabilities* (18).

[§]Decision theorists have also suggested blending the minimax and minimax-regret criteria with the expected utility model. A decision maker who feels able to assert a partial distribution on the states of nature can 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 -minimax and Γ -minimax regret criteria (17). The Γ -minimax approach has recently drawn attention from economists, such as in refs. 32 and 33.

Consideration of a large population simplifies analysis in important respects. First, the planner can ignore finite-sample statistical variation in the outcome of random treatment assignment. Second, the *ex ante* probability that a randomly drawn person becomes ill is identical to the *ex post* illness rate in the population. Formally, these statements hold if the population contains infinitely many members. Standard arguments using Laws of Large Numbers enable their approximate application to large finite populations.

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 mean social harm caused by illness and let $c > 0$ denote the mean social cost per vaccination, measured in commensurate units. The social cost of vaccination rate t is

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

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 optimization problem is invariant to the scale of $S(\cdot)$. Therefore, without loss of generality, I shall henceforth let $a = 1$ and interpret c as the ratio of the mean social cost of vaccination to the mean social cost of illness.

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 ref. 6. However, it has been more common to pose a susceptible-infectious-removed 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, refs. 3 and 7. The latter authors ask (p. 86): “What minimal fraction of each age group should be vaccinated to eliminate the possibility of an influenza epidemic in the whole population?”

It is important to understand that the objective of preventing onset of an epidemic generally differs from minimization of social cost as specified in Eq. 1. In epidemiology, an epidemic is defined to occur when the infected fraction of the population increases with time. In contrast, our social cost function abstracts from the time path of illness and considers the prevalence of illness in the population. The analysis of this paper applies to endemic diseases as well as to epidemic-prone ones. Application of the analysis to epidemic-prone diseases is most appropriate when social cost depends mainly on the prevalence of illness in the population rather than on the timing of new cases.

1.2. 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

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

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

$$0 = 2\rho t - 2\rho + c. \quad [3]$$

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

$$t^* = 0 \quad \text{if } 2\rho < c. \\ = 1 - c/(2\rho) \text{ if } 2\rho \geq c. \quad [4]$$

Observe that for no value of parameters (c, ρ) is it optimal to vaccinate the entire population. It is, however, optimal to vaccinate no one if $2\rho < c$.

2. Vaccination with Partial Knowledge of the External-Response Function

2.1. Empirical Evidence and Maintained Assumptions. Solution of the planning problem of Section 1 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 external effectiveness may choose a vaccination policy, I consider decision making in a particular informational setting. I suppose that the planner observes the vaccination and illness rates of a study population whose vaccination rate has been chosen previously to be some value less than one. The planner maintains two assumptions. First, he assumes that the study population and the treatment population have the same external-response function. Second, he assumes that the illness rate of unvaccinated persons weakly decreases as the vaccination rate increases. However, he makes no assumption about the magnitude of the external effect of vaccination.

Let $r < 1$ denote the vaccination rate in the study population and let $q(1 - r)$ denote the realized illness rate. Formally, I maintain *Assumption 1 (Study Population)*: The planner observes r and $q(1 - r)$. He 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

$$t \leq r \Rightarrow p(t) \geq q, \\ t \geq r \Rightarrow p(t) \leq q. \quad [5]$$

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 often is possible to observe the vaccination and illness rates of a study population, and credible to assume that the study and treatment populations have similar if not identical external-response functions.[†] It usually is credible to assume that vaccination weakly prevents illness. Assumption 2 is a specific instance of the general idea of *monotone treatment response* developed in ref. 34.

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

[†]This assumption may be particularly credible when the study population is a recent cohort of the population now about to be treated, within which treatments were randomly assigned. However, the assumption is not innocuous. If the study population differs from the treatment population in size or social structure, or if assignment to vaccination was nonrandom, its external-response function may differ as well.

2.2. 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. Let γ index the elements of Γ , the symbol p being reserved for the unknown true function. For $\gamma \in \Gamma$, let

$$S(t, \gamma) = \gamma(t)(1-t) + ct \quad [6]$$

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 alternative vaccination rates. Let

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

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 $D(t, s) \equiv \sup_{\gamma \in \Gamma} d(t, s; \gamma)$. A verifiable sufficient condition for strict dominance is that $D(t, s) < 0$.

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

Lemma 1: Let Assumptions 1 and 2 hold. Then

- (i) $t < s \leq r \Rightarrow D(t, s) = (1-t) - q(1-s) + c(t-s)$,
- (ii) $t < r < s \Rightarrow D(t, s) = (1-t) + c(t-s)$,
- (iii) $s \leq t < r \Rightarrow D(t, s) = (c-q)(t-s)$,
- (iv) $r \leq t < s \Rightarrow D(t, s) = q(1-t) + c(t-s)$,
- (v) $r \leq s \leq t \Rightarrow D(t, s) = c(t-s)$,
- (vi) $s < r \leq t \Rightarrow D(t, s) = (c-q)(t-s)$.

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

- (a). Let $c < q$. Then s is strictly dominated if $s < r$.
- (b). Let $c > q$. Then s is strictly dominated if $s > r + q(1-r)/c$.
- (c). Let $c > 1$. Then s is strictly dominated if $(1-q)/(c-q) < s \leq r$ or if $s > \max(r, 1/c)$.

It might have been thought that Assumptions 1 and 2 are too weak to yield interesting dominance findings. However, Proposition 2 shows that 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)–(c) give the specifics. Part (a) shows that vaccination rates lower than the rate r observed in the study population are dominated if $c < q$. Thus, when this inequality holds, 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 $c > q$. Part (c) shows that additional vaccination rates are dominated if $c > 1$.

2.3. 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 generically ensure choice of an undominated alternative. I say “generically” because commonly studied criteria (including expected utility, minimax, and minimax-regret) necessarily yield undominated actions when they have unique optima. However, when these criteria have multiple optima, some may be dominated.

Particularly familiar to economists is the expected utility model. In the present setting, this recommends that the planner

place a subjective distribution on the feasible $p(\cdot)$ and minimize subjective expected social cost. The cost function in Eq. 1 is linear in $p(\cdot)$, so subjective expected social cost is

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

where Ψ is the subjective distribution and $\pi(\cdot) \equiv E_{\Psi}[p(\cdot)]$ 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 1.

Applications of the expected utility model have typically assumed that the support of a subjective distribution is a finite-dimensional real space. However, Assumptions 1 and 2 only require that $p(\cdot)$ be an element of the abstract space of weakly decreasing functions that satisfy Eq. 5. A planner who assumes that $p(\cdot)$ belongs, with subjective probability one, to a finite-dimensional subset of this function space asserts considerable knowledge beyond Assumptions 1 and 2. Permitting the support of the subjective distribution to be a larger subset of the function space is mathematically delicate. I will not address the matter here.

2.4. 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.^{||} A planner who cannot substantiate choice of π can reasonably apply the minimax or minimax-regret criterion to choose a vaccination rate. I derive the minimax vaccination rate under Assumptions 1 and 2 here and the minimax-regret rate in Section 2.5.

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

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

Proposition 3 derives the vaccination rate that solves this problem.

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

$$\begin{aligned} t^m &= 0 && \text{if } c > 1 \text{ and } 1 \leq q(1-r) + cr, \\ &= r && \text{if } c > 1 \text{ and } 1 \geq q(1-r) + cr \\ &= \text{all } t \in [0, 1] && \text{if } c = q \text{ and } q = 1, \\ &= \text{all } t \in [r, 1] && \text{if } c = q \text{ and } q < 1, \\ &= 1 && \text{if } c < q. \quad \square \end{aligned} \quad [10]$$

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

2.5. 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

^{||}Decision theorists have sometimes asserted preeminence 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 in refs. 31 and 35. However, I have argued that the theorems of axiomatic decision theory should not be viewed as prescriptive (36).

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

$$\begin{aligned} & \min_{t \in [0,1]} \max_{\gamma \in \Gamma} S(t, \gamma) - S^*(\gamma) \\ &= \min_{t \in [0,1]} \max_{\gamma \in \Gamma} \{ \gamma(t)(1-t) + ct - \min_{s \in [0,1]} [\gamma(s)(1-s) + cs] \} \\ &= \min_{t \in (0,1)} \max_{\gamma \in \Gamma, s \in (0,1)} \gamma(t)(1-t) + ct - \gamma(s)(1-s) - cs \quad [11] \\ &= \min_{t \in [0,1]} \max_{s \in [0,1]} D(t, s), \end{aligned}$$

where $D(t, s)$ was defined in Section 2.2. Proposition 4 derives the vaccination rate that solves this problem.

Proposition 4: Let Assumptions 1 and 2 hold.

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

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

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

$$t^{mr} = \operatorname{argmin}_{t \in [0,1]} \{ 1[t < r] \bullet \{ \max[(1-q)(1-t), (1-t) + c(t-r), (c-q)t] \} + 1[t \geq r] \bullet \{ \max[q(1-t), c(t-r), (c-q)t] \} \}.$$

If $r = 0$, then $t^{mr} = q/(q + c)$. \square

[13]

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

When $c > q$, the solution to the minimax-regret problem generally does not have an explicit form of simplicity comparable to Eq. 12. However, the abstract finding in Eq. 13 simplifies in the polar case $r = 0$, where no one was vaccinated in the study population. Then t^{mr} decreases from 1/2 to 0 as c increases from q to ∞ .

2.6. Numerical Examples. Numerical examples are useful to illustrate the findings of Sections 1 and 2. I give three here, each modifying the preceding example in some respect.

First, consider a scenario where the mean cost of vaccination (relative to illness) is $c = 0.02$. 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 conclude that $\rho = 1/5$ and would use Eq. 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 Proposition 2 to conclude that any vaccination rates are dominated, because $c < q$ and $r = 0$. By Proposition 3, the minimax vaccination rate is $t^m = 1$. By Proposition 4, Eq. 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 $c = 0.02$. A planner who believes the external-response function is linear would still conclude that $\rho = 1/5$ and choose $t^* = 7/8$. A planner who only knows the function to be weakly decreasing can use Proposition 2 to determine that any vaccination rate smaller than 1/2 is strictly dominated. By Proposition 3, the minimax vaccination rate remains $t^m = 1$. By Proposition 4, Eq. 12, the minimax-regret rate is $t^{mr} = 5/6$.

Now revise the scenario again by supposing that vaccination is more costly relative to illness, say $c = 0.25$. Continue to assume that $r = 1/2$ and $q = 1/10$. In this case, a planner who believes the

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

3. Extensions of the Analysis

The analysis of Section 2 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. Recall that the social cost function in Eq. 1 considers only the prevalence of illness in the population, abstracting from the time path of illness. Hence, our analysis applies to endemic diseases and to epidemic-prone diseases in which social cost does not depend on the timing of cases. A broad direction for extension of the analysis is to consider epidemic-prone diseases in which social cost does depend on the timing of cases. This requires respecification of the social cost function to stipulate how timing matters.

Within the domain where the cost function in Eq. 1 is appropriate, the knowledge available to the planner may differ in many ways from Assumptions 1 and 2. In general, the set of dominated vaccination rates varies with the assumptions maintained. If Assumptions 1 and 2 are strengthened, all rates found to be dominated in Proposition 2 continue to be dominated and other rates may become dominated. If either Assumption 1 or 2 is weakened, the findings of Proposition 2 need not hold.

Modifying Assumption 1, the planner may believe that the study population differs from the treatment population in known or partially known ways. Or he may only observe sample data from the study population, not the illness rate of the entire population. In the latter case, the planner may apply the finite-sample versions of dominance, minimax, and minimax-regret developed in statistical decision theory (27).

Modifying Assumption 2, the planner may have knowledge that makes it credible to assume more about the external-response function than monotonicity. The research on optimal vaccination policy cited in the Introduction makes specific assumptions about disease transmission, implying specific conclusions about external-response functions. I have observed that authors typically provide little information that would enable one to assess the accuracy of their assumptions. Nevertheless, some assumptions may be sufficiently credible that a planner should use them in decision making.

I discuss here four extensions of the analysis. I first consider vaccination with imperfect internal effectiveness (Section 3.1), next stratification of the population when its members have observable covariates (Section 3.2), then provision of incentives when it is infeasible to mandate vaccination (Section 3.3), and finally dynamic choice of vaccination rates for a sequence of cohorts (Section 3.4).

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

$$S(t) = p(\lambda t)(1 - \lambda t) + ct. \quad [14]$$

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. I will show here that, when λ is known, most of the analysis of Section 2 applies to the new planning problem with minor modification. Knowledge of λ sometimes is realistic. In particular, λ might be learned empirically (or at least estimated) by measurement of the immune response of a random sample of vaccinated persons.

Let $\tau \equiv \lambda t$ be the *effective vaccination rate*. We may rewrite the cost function in Eq. 14 as

$$S'(\tau) = p(\tau)(1 - \tau) + (c/\lambda)\tau. \quad [15]$$

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 2 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 2 would apply directly 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 to 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 Proposition 2 shows that whenever a vaccination rate s is dominated, there exists a dominating rate, say $t(s)$, such that $t(s) \leq r$. Applied to the new planning problem, this means that when an effective rate σ is dominated, there exists a dominating effective rate, say $\tau(\sigma)$, such that $\tau(\sigma) \leq \lambda r$. It follows that when σ is dominated, the corresponding raw rate σ/λ is dominated by $\tau(\sigma)/\lambda$.

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

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

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

Consider the minimax-regret criterion. The proof to Proposition 4 is more complex than those of the other propositions, and I will not attempt a full adaptation here. However, a simple result emerges when application of Proposition 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 .

3.2. 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 . I will assume that X is a finite space and that the persons with each value of x form a large subpopulation.

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

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

I do not normalize the mean harm per illness to equal one here because mean harm may vary with x . The mean vaccination cost and the external-response function may also 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 Sections 1 and 2 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, assume the social cost function has the form

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

where $p_x(\cdot)$ is a function only of t_x , not of $(t_w, w \in X)$. This assumption is credible if x indexes geographically isolated groups who do not come into physical contact with one another.

The cost function in Eq. 18 is separable in x . Hence, the planner may treat each x -specific subpopulation separately, as if it 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 subpopulation separately.

When the cost function in Eq. 17 does not have the separable structure in Eq. 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.

3.3. 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.

Economists often refer to choice of an incentive policy as a *mechanism design* problem. Research on mechanism design has routinely assumed that the planner maximizes subjective expected welfare. Here we are concerned with mechanism design under ambiguity.

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 may 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 may face more ambiguity than does one who mandates vaccination.

3.4. 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 and young children. 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 2. Determination of dominated vaccination rates, the minimax rate, and the minimax-regret rate does not differ conceptually from the analysis of Section 2, 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 2 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.

4. Conclusion

Vaccination policy has regularly been studied under the assumption that the planner knows the internal and external responses of illness to vaccination. Consideration of vaccination with partial knowledge has received little attention, and then only using the expected utility model.

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. The dominance result in Proposition 2 is the most important finding of the paper. It is useful to any planner who feels it credible to maintain Assumptions 1 and 2 or any strengthening thereof. The derivations of minimax and minimax-regret vaccination rates in Propositions 3 and 4 will be useful to planners who wish to use one of these decision criteria and who feel it credible to maintain Assumptions 1 and 2 but no more. Strengthening these assumptions may change the minimax and minimax-regret rates.

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 Lemma and Propositions

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 $D(t, s) = (1 - t) + ct - q(1 - s) - cs$.
- (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 $D(t, s) = (1 - t) + ct - cs$.
- (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 $D(t, s) = (c - \delta)(t - s)$. Given that $t \geq s$, the maximum is attained by setting $\delta = q$. Hence, $D(t, s) = (c - q)(t - s)$.
- (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 $D(t, s) = q(1 - t) + ct - cs$.
- (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 $D(t, s) = (c - \delta)(t - s)$. Given that $t \geq s$, 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 $D(t, s) = (c - q)(t - s)$.

Q. E. D.

Proposition 2.

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

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

- (c). Part (i) of Lemma 1 showed that if $t < s \leq r$, then $D(t, s) = (1 - t) - q(1 - s) + c(t - s) = (1 - q) + (c - 1)t - (c - q)s$.

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

Part (ii) of Lemma 1 showed that if $t < r < s$, then $D(t, s) = (1 - t) + c(t - s) = 1 + (c - 1)t - cs$.

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

Q. E. D.

Proposition 3. For each value of t , the inner maximization problem in Eq. 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,

$$t^m = \operatorname{argmin}_{t \in [0, 1]} \{1[t < r] + q \cdot 1[t \geq r]\}(1 - t) + ct.$$

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

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

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

Q. E. D.

Proposition 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 $c \leq q$. Let $t < r$. There are three cases to consider:

- (i) $\max_{s: t < s \leq r} D(t, s)$ occurs at $s = r$. Hence, $\max_{s: t < s \leq r} D(t, s) = (1 - t) - q(1 - r) + c(t - r)$.
- (ii) $\sup_{s: t < r < s} D(t, s)$ occurs at $s = r$. Hence, $\sup_{s: t < r < s} D(t, s) = (1 - t) + c(t - r)$.
- (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, $\sup_{s \in [0, 1]} D(t, s) = (1 - t) + c(t - r)$.

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

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, $\sup_{s: r \leq t < s} D(t, s) = q(1 - t)$.
- (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, $\max_{s: s \leq r \leq t} D(t, s) = (c - q)(t - r)$.

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

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

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

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

- (i) $\max_{s: t < s \leq r} D(t, s)$ occurs at $s = t$. Hence, $\max_{s: t < s \leq r} D(t, s) = (1 - q)(1 - t)$.
- (ii) $\sup_{s: t < r < s} D(t, s)$ occurs at $s = r$. Hence, $\sup_{s: t < r < s} D(t, s) = (1 - t) + c(t - r)$.
- (iii) $\max_{s \leq t < r} D(t, s)$ occurs at $s = 0$. Hence, $\max_{s \leq t < r} D(t, s) = (c - q)t$.

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

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, $\sup_{s: r \leq t < s} D(t, s) = q(1 - t)$.
- (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, $\max_{s: s \leq r \leq t} D(t, s) = (c - q)t$.

Hence, maximum regret for $t \geq r$ is $\max[q(1 - t), c(t - r), (c - q)t]$. Combining the above findings, the minimax-regret vaccination rate is

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

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

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

This problem has unique solution $q/(q + c)$.
Q. E. D.

ACKNOWLEDGMENTS. I am grateful for comments from Larry Blume, Buz Brock, Mike Fu, Guido Imbens, and Matthew Tanner. This research was supported in part by National Science Foundation Grant SES-0911181.

1. Mirrlees J (1971) An exploration in the theory of optimal income taxation. *Rev Econ Stud* 38:175–208.
2. Polinsky M, Shavell S (2000) The economic theory of public enforcement of law. *J Econ Lit* 38:45–76.
3. Ball F, Lyne O (2002) Optimal vaccination policies for stochastic epidemics among a population of households. *Math Biosci* 177–178:333–354.
4. Becker NG, Starczak DN (1997) Optimal vaccination strategies for a community of households. *Math Biosci* 139:117–132.
5. Boulier B, Datta T, Goldfarb R (2007) Vaccination externalities. *BE J Econ Anal Policy* 7(Issue 1):Article 23. Available at <http://www.bepress.com/bejeap/vol7/iss1/art23>.
6. Brito D, Sheshinski E, Intriligator M (1991) Externalities and compulsory vaccinations. *J Public Econ* 45:69–90.
7. Hill AN, Longini IM, Jr (2003) The critical vaccination fraction for heterogeneous epidemic models. *Math Biosci* 181:85–106.
8. Patel R, Longini IM, Jr, Halloran ME (2005) Finding optimal vaccination strategies for pandemic influenza using genetic algorithms. *J Theor Biol* 234:201–212.
9. Scuffham PA, West PA (2002) Economic evaluation of strategies for the control and management of influenza in Europe. *Vaccine* 20:2562–2578.
10. Nordhaus W (2008) *A Question of Balance: Weighing the Options on Global Warming Policy* (Yale Univ Press, New Haven, CT).
11. Tanner MW, Sattenspiel L, Ntamo L (2008) Finding optimal vaccination strategies under parameter uncertainty using stochastic programming. *Math Biosci* 215:144–151.
12. Weycker D, et al. (2005) Population-wide benefits of routine vaccination of children against influenza. *Vaccine* 23:1284–1293.
13. Ellsberg D (1961) Risk, ambiguity, and the Savage axioms. *Q J Econ* 75:643–669.
14. Keynes J (1921) *A Treatise on Probability* (MacMillan, London).
15. Knight F (1921) *Risk, Uncertainty, and Profit* (Houghton-Mifflin, Boston, MA).
16. Arrow K, Hurwicz L (1972) An optimality criterion for decision-making under ignorance. *Uncertainty and Expectations in Economics*, eds Carter D, Ford J (Blackwell, Oxford).
17. Berger J (1985) *Statistical Decision Theory and Bayesian Analysis* (Springer, New York).
18. Walley P (1991) *Statistical Reasoning with Imprecise Probabilities* (Chapman & Hall, London).
19. Manski C (2000) Identification problems and decisions under ambiguity: Empirical analysis of treatment response and normative analysis of treatment choice. *J Econom* 95:415–442.
20. Manski C (2002) Treatment choice under ambiguity induced by inferential problems. *J Statist Plann Inference* 105:67–82.
21. Manski C (2005) *Social Choice with Partial Knowledge of Treatment Response* (Princeton Univ Press, Princeton, NJ).
22. Manski C (2006) Search profiling with partial knowledge of deterrence. *Econ J* 116: F385–F401.
23. Manski C (2007) Minimax-regret treatment choice with missing outcome data. *J Econom* 139:105–115.
24. Manski C (2007) *Identification for Prediction and Decision* (Harvard Univ Press, Cambridge, MA).
25. Manski C (2009) Social planning with partial knowledge of social interactions. *J Inst Theor Econ* 165:146–158.
26. Manski C (2009) Diversified treatment under ambiguity. *Int Econ Rev* 50:1013–1041.
27. Wald A (1950) *Statistical Decision Functions* (Wiley, New York).
28. Manski C (2004) Statistical treatment rules for heterogeneous populations. *Econometrica* 72:1221–1246.
29. Manski C, Tetenov A (2007) Admissible treatment rules for a risk-averse planner with experimental data on an innovation. *J Statist Plann Inference* 137:1998–2010.
30. Savage L (1951) The theory of statistical decision. *J Am Stat Assoc* 46:55–67.
31. Savage L (1954) *The Foundations of Statistics* (Wiley, New York).
32. Gilboa I, Schmeidler D (1989) Maxmin expected utility with non-unique prior. *J Math Econ* 18:141–153.
33. Hansen L, Sargent T (2008) *Robustness* (Princeton Univ Press, Princeton, NJ).
34. Manski C (1997) Monotone treatment response. *Econometrica* 65:1311–1334.
35. von Neumann J, Morgenstern O (1944) *Theory of Games and Economic Behavior* (Princeton Univ Press, Princeton, NJ).
36. Manski C (2010) Actualist rationality. *Theory Decis*, in press.