Vaccine Approvals and Mandates Under Uncertainty: Some Simple Analytics

Charles F. Manski
Board of Trustees Professor of Economics
Faculty Fellow, Institute for Policy Research
Northwestern University

Version: December 2014

DRAFT

Abstract

Social interactions make communicable disease a core concern of public health policy. A prevalent problem is scarcity of empirical evidence that is informative about how interventions affect population behavior and illness. Randomized trials, which have been important to evaluation of treatments for non-infectious diseases, are less informative about treatment of communicable diseases because they do not shed light on population-wide disease transmission. In particular, trials do not reveal the indirect preventive (herd immunity) effect of vaccination on persons who are not vaccinated or who are unsuccessfully vaccinated. This paper studies the decision problems faced by health planners who must choose whether to approve a new vaccine or mandate an approved one, but who do not know the indirect effect of vaccination. Manski studies vaccine approval as a choice between a zero vaccination rate (rejection of the new vaccine) and whatever vaccination rate the health-care system will yield if the vaccine is approved. He studies the decision to mandate an approved vaccine as a choice between vaccinating the entire population (the mandate) and the vaccination rate that would be generated by decentralized health-care decisions. Considering decision making with partial knowledge, he shows that it might be possible to determine optimal policies in some cases where the planner can only bound the indirect effect of vaccination. Considering settings where optimal policy is indeterminate, he poses several criteria for decision-making—expected utility, minimax, and minimaxregret— and derives the policies they yield. He suggests that performance of formal decision analysis can improve prevailing vaccine approval and mandate procedures.

1. Introduction

Social interactions in treatment response make communicable disease a core concern of public health policy. Spread of infection creates a negative externality. Preventive administration of vaccines, therapeutic administration of antimicrobial drugs, and separation of infected persons from the general population may reduce disease transmission. In a decentralized health care system, infected and at-risk persons may not adequately recognize the social implications of their actions. Hence, there may be a rationale for government to seek to influence treatment of communicable disease. See, for example, Brito, Sheshinski, and Intriligator (1991) and Althouse, Bergstrom, and Bergstrom (2010). In practice, policies range from quarantines of infected persons to mandatory vaccination of certain sub-populations to subsidization of vaccines and drugs.

A prevalent problem in policy choice is scarcity of empirical evidence that are informative about how alternative interventions affect population behavior and illness. Randomized clinical trials (RCTs), which have been so important to evaluation of treatments for non-infectious diseases, are less informative about treatment of communicable diseases. The classical argument for inference from RCTs assumes that the outcome experienced by each person or other treatment unit may vary only with his own treatment, not with those of other members of the population. This assumption—variously called non-interference (Cox, 1958), the stable unit treatment value assumption (Rubin, 1978), or individualistic treatment response (Manski, 2013)—does not hold when treating communicable diseases. RCTs have limited power to identify the effects of treatments for such diseases.

A leading case is the decision problem of a health planner who must choose a vaccination rate for a population. Vaccination of a particular person may benefit him directly by generating an immune response that reduces his susceptibility to the disease. It may also reduces the infectiousness of this person and thereby inhibit transmission of the disease to others who are unvaccinated or unsuccessfully vaccinated. Thus, vaccination may have both a direct preventive effect on the person vaccinated and an indirect preventive effect on other persons. The indirect effect is sometimes called herd or community immunity.

1

See Fine (1993) and Fine, Eames, and Heymann (2011) for discussions of the long history and use of this concept in epidemiology.[1]

An RCT that randomly vaccinates a specified fraction of the population may enable evaluation of the direct effect of vaccination on illness, but it does not reveal the indirect effect. The trial only reveals the illness outcomes that occur with the vaccination rate used in the trial. The outcomes that the population would experience with other vaccination rates remain counterfactual. Yet policy choice requires comparison of alternative vaccination rates.[2]

Attempting to cope with the dearth of empirical evidence, researchers have used epidemiological models of disease transmission to forecast the outcomes that would occur with counterfactual vaccination policies. Unfortunately, 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.

With this background, I have previously considered choice of a vaccination rate when a planner has partial knowledge of the indirect effect of vaccination on illness (Manski, 2010). I posed 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. I supposed that the planner

---

[1] While the usual presumption is that vaccines have beneficial indirect effects, they can in principle have negative effects. Vaccines based on attenuated live pathogens may infect persons other than the recipient. Vaccination may also encourage selection of pathogens towards variants resistant to the vaccine. However, these negatives forces are thought to be weak in practice. See Mishra *et al.* (2012).

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

observes the illness rate of a study population whose vaccination rate has been chosen previously. I assumed that the illness rate of unvaccinated or unsuccessfully vaccinated persons weakly decreases as the vaccination rate increases, but that the magnitude of this indirect preventive effect of vaccination is unknown. In this setting, I first showed how the planner can eliminate dominated vaccination rates and then how he can use various decision criteria to choose an undominated vaccination rate.

Whereas Manski (2010) studied decision making by a hypothetical powerful planner who can choose any vaccination rate, the present paper considers two prominent instruments of actual vaccination policy, vaccine approvals and mandates. To begin, Section 2 describes prevailing procedures for approval of new vaccines and for mandating approved ones in the United States. Vaccine approval is a federal responsibility performed by the Food and Drug Administration (FDA). The FDA approval process aims to measure the direct effects of vaccines by requiring performance of RCTs. It does not seek to measure indirect effects, ignores monetary costs, and rejects use of formal decision analysis as a tool in policy making. Vaccine mandates are primarily enacted by states. There is no uniform decision process across jurisdictions, but the indirect effects and monetary costs of vaccination are prevalent considerations.

To provide a formal framework for evaluation of vaccine approvals and mandates, I adapt the planning problem posed in Manski (2010). My analysis intentionally makes various simplifying assumptions in order to make basic points transparently. Relaxing these assumptions may complicate the specifics of the analysis but the general principles would remain intact.

Section 3 supposes that the planner has sufficient knowledge to choose an optimal policy. Among other things, the planner knows the *indirect-response function* that describes how the illness rate of unvaccinated or unsuccessfully vaccinated persons varies with the vaccination rate.[3] I earlier studied the unconstrained optimization problem where the planner can choose any vaccine rate. Here I view decisions

---

[3] The indirect-response function was called the *external-response* function in Manski (2010). I revise the word "external" to "indirect" here to make the terminology consistent with the norm in the literature on vaccination.

to approve or mandate vaccines as constrained optimization problems.

Vaccine approval presents a choice between a zero vaccination rate (rejection of the new vaccine) and whatever vaccination rate the health-care system will yield if the vaccine is approved. The decision to mandate an approved vaccine presents a choice between vaccinating the entire population (the mandate) and the vaccination rate that would be generated by decentralized health-care decisions. I show that a planner can solve the constrained optimization problem posed by vaccine approval if he knows the direct effect of vaccination on illness as well as the vaccination and illness rates that will be generated by the health-care system if the vaccine is approved. He can solve the constrained optimization problem that arises when contemplating a vaccine mandate if he knows the vaccination and illness rates that would be generated by decentralized decision making as well as the illness rate that would result if the entire population were vaccinated.

Section 4 considers decision making with partial knowledge. I discuss the knowledge that a planner may have in practice and conclude that incomplete knowledge of the indirect effect of vaccination is likely. I show that it may be possible to determine optimal policies in some cases where the planner can only bound the indirect effect. In other cases the optimal policy is indeterminate. I pose several criteria for decision making in such cases—expected utility, minimax, and minimax-regret—and derive the policies they yield.

The concluding Section 5 suggests that performance of formal decision analysis can improve prevailing vaccine approval and mandate procedures. I question the narrow scope of the FDA approval process, which uses a qualitative approach that ignores the monetary costs and indirect effects of vaccination.

2. Prevailing Vaccine Approval and Mandate Procedures

The American institutions for approving new vaccines and mandating approved ones differ in numerous respects. Whereas the federal Food and Drug Administration performs the former function, state public health agencies perform the latter. Whereas the FDA approval process focuses on the direct effect of vaccination on the person vaccinated, concern with indirect effects on unvaccinated or unsuccessfully vaccinated persons is a central rationale for state mandates. Whereas FDA approval decisions ignore the monetary costs associated with vaccination and illness, state mandate decisions explicitly consider these costs. I explain here.

2.1. FDA Approval

The FDA vaccine approval process is similar to the one used for drug approval. The FDA describes it this way (U.S. Food and Drug Administration, 2009):

> "FDA's Center for Biologics Evaluation and Research (CBER) is responsible for regulating vaccines in the United States. . . . Vaccine clinical development follows the same general pathway as for drugs and other biologics. A sponsor who wishes to begin clinical trials with a vaccine must submit an Investigational New Drug application (IND) to FDA. The IND describes the vaccine, its method of manufacture, and quality control tests for release. Also included are information about the vaccine's safety and ability to elicit a protective immune response (immunogenicity) in animal testing, as well as the proposed clinical protocol for studies in humans. Pre-marketing (pre-licensure) vaccine clinical trials are typically done in three phases, as is the case for any drug or biologic. Initial human studies, referred to as Phase 1, are safety and immunogenicity studies performed in a small number of closely monitored subjects. Phase 2 studies are dose-ranging studies and may enroll hundreds of subjects. Finally, Phase 3 trials typically enroll thousands of individuals and provide the critical documentation of effectiveness and important additional safety data required for licensing. . . . . If successful, the completion of all three phases of clinical development can be followed by the

submission of a Biologics License Application (BLA). To be considered, the license application must provide the multidisciplinary FDA reviewer team . . . . with the efficacy and safety information necessary to make a risk/benefit assessment and to recommend or oppose the approval of a vaccine." Thus, FDA approval of new vaccines rests primarily on evidence from RCTs administering the vaccine, as does FDA approval of drugs.

As discussed in the Introduction, RCTs are generally uninformative about the indirect effect of vaccination on illness. The FDA approval process focuses on the safety of the new vaccine and its direct preventive effect on the person vaccinated, with scant attention to indirect effects. A search of the extensive "Vaccines, Blood & Biologics" section of the FDA website reveals only five mentions of the term "herd immunity" and none of the terms "indirect effect" and "community immunity." Documents associated with FDA approval of specific vaccines contain at most brief discussions of herd immunity or related concepts. See, for example, the summary for approval of the varicella (chickenpox) vaccine (U.S. Food and Drug Administration, 1994).

The FDA description of the vaccine approval process, quoted above, refers to performance of a "risk/benefit assessment" en route to an approval decision. The word "risk" refers to the possibility that administration of a vaccine to a person may have deleterious side effects while "benefit" refers to the direct preventive effect of vaccination on illness. The FDA description of its drug approval process similarly refers to assessment of risks and benefits. The approval processes for both vaccines and drugs take no account of the monetary costs associated with treatment and illness. This explains use of the term risk/benefit rather than cost/benefit.

To an economist or decision analyst, the term "risk/benefit assessment" suggests quantification of risks and benefits and weighing the two to measure the net value of approving a vaccine or drug. However, FDA documents summarizing approval decisions do not quantify the tradeoff between risks and benefits. Indeed, the FDA administration has explicitly rejected use of formal decision analysis in drug and vaccine approval.

6

A 2012 amendment to the federal Food, Drug, and Cosmetic Act called on the FDA to institute a "structured benefit-risk assessment framework," stating (U. S. Congress, 2012, Section 905):

"Section 505(d) (21 U.S.C. 355(d)) is amended by adding at the end the following: 'The Secretary shall implement a structured risk-benefit assessment framework in the new drug approval process to facilitate the balanced consideration of benefits and risks, a consistent and systematic approach to the discussion and regulatory decision making, and the communication of the benefits and risks of new drugs."

In response to this Congressional mandate, the FDA subsequently released to the public a draft document that is so revealing of current FDA thinking that I will quote from it at length (U.S. Food and Drug Administration, 2013).

The FDA document begins this way (page 1):

"In the past, some FDA stakeholders have indicated that there is room for improvement in the clarity and transparency of FDA's benefit-risk assessment in human drug review. When FDA approves a new product, the agency publishes the various relevant documents, such as discipline reviews . . . and decision memoranda, on its website. While FDA takes great care to clearly explain the reasoning behind a regulatory decision in these documents, the clinical analysis may not always be readily understood by a broad audience who may wish to understand FDA's thinking. In addition, some have argued that drug regulatory decisions should be based on more formalized and quantitative approaches to benefit-risk assessment, including the assignment of weights to benefit and risk considerations. Others, however, are skeptical of fully quantitative approaches, and consider such attempts to be a highly subjective exercise that would add little clarity to regulatory decision making."

It goes on to side with those who are skeptical of quantitative approaches, stating (page 4):

"In the last few years, as other disciplines such as decision science and health economics have been applied to drug regulatory decision-making, there has been much discussion among regulators, industry, and other stakeholders regarding "qualitative" versus "quantitative" approaches to benefit-risk assessment. The term "quantitative benefit-risk assessment" can have various meanings depending on who is asked. Some hold the view that a quantitative benefit-risk assessment encompasses approaches that seek to quantify benefits and risks, as well as the weight that is placed on each of the components such that the entire benefit-risk assessment is quantitative. This approach

is typical of quantitative decision modeling. It usually requires assigning numerical weights to benefit and risk considerations in a process involving numerous judgments that are at best debatable and at worst arbitrary. The subjective judgments and assumptions that would inevitably be embodied in such quantitative decision modeling would be much less transparent, if not obscured, to those who wish to understand a regulator's thinking. Furthermore, application of quantitative decision modeling seems most appropriate for decisions that are largely binary. Many benefit-risk assessments are more nuanced and conditional based on parameters that could be used to effectively manage a safety concern in the post-market setting. There is significant concern that reliance on a relatively complex model would obscure rather than elucidate a regulator's thinking.

These concerns have led FDA to the conclusion that the best presentation of benefit-risk considerations involves focusing on the individual benefits and risks, their frequency, and weighing them appropriately. FDA believes that this can be accomplished by a qualitative descriptive approach for structuring the benefit-risk assessment that satisfies the principles outlined earlier in this section, while acknowledging that quantification of certain components of the benefit-risk assessment is an important part of the process to support decision-making. FDA considers it most important to be clear about what was considered in the decision, to be as quantitative as possible in characterizing that information, and to fully describe how that information was weighed in arriving at a conclusion. Quantitative assessments certainly underpin the qualitative judgments of FDA's regulatory decisions, but FDA has adopted a structured qualitative approach that is designed to support the identification and communication of the key considerations in FDA's benefit-risk assessment and how that information led to the regulatory decision."

Thus, the FDA has rejected performance of formal decision analysis en route to drug and vaccine approval.

Indeed, an aversion to decision analysis appears widespread in medicine. For example, a recent report of the Institute of Medicine (2011) proposing standards for development of clinical practice guidelines contains only a brief cautionary mention of decision analysis as a potential tool for evaluation of practice recommendations. The report states (p. 171): "A frontier of evidence-based medicine is decision analytic modeling in health care alternatives' assessment. . . . . Although the field is currently fraught with controversy, the committee acknowledges it as exciting and potentially promising, however, decided the state of the art is not ready for direct comment." See Manski (2013b) for further discussion.

2.2. State Mandates

In the American federal system, each state decides the circumstances in which to mandate administration of vaccines. Specific vaccines may be mandated for children in day care, pupils in schools, students in colleges, and residents of long-term care facilities.[4] While there is no uniform decision process across states, the indirect effect of vaccination is a prevalent consideration.

For example, the state of Washington lists this as one of nine criteria considered when contemplating a mandate for vaccination of children and young adults (Washington State Board of Health, 2006):

> "Vaccinating the infant, child, or adolescent against this disease reduces the risk of person-to-person transmission. Having some proportion of the population vaccinated with the antigen helps to stem person to person transmission of the disease (i.e, herd immunity). Even community members who are not vaccinated (such as newborns and those with chronic illnesses) are offered some protection because the disease has less opportunity to spread within the community. Vaccinating children in school and/or child care centers can increase the percentage of children in these groups who are immune and thus reduce the risk of outbreaks of the disease in these groups and in the community at large."

The Washington state criteria also consider the monetary cost of vaccination and the burden that compliance with the mandate places on the parents of children.

While documents describing the decision processes used by Washington and other states express concern with the indirect effects and monetary costs of vaccination, they do not explain the manner in which these considerations are measured and balanced. It is not clear whether formal decision analysis is used or whether, like the FDA, the states use qualitative approaches.

---

[4] The webpage www.immunize.org/laws/ describes the present mandates in each state.

3. Optimal Vaccination

Economists and decision analysts performing normative study of public policy have long analyzed the decision making of a real or hypothetical planner who seeks to choose a policy that maximizes social welfare or, equivalently, minimizes social cost. A standard exercise specifies a set of feasible policies and supposes that the planner knows the outcome yielded by each policy. The objective of the exercise is to determine the optimal policy.

Various instances of this exercise appear in the literature on optimal vaccination, which contains contributions by economists, epidemiologists, and others. Studies have varied in the policy options they consider, the welfare functions they pose, and the manner in which they model the transmission of disease. However, they have almost universally assumed that the planner has sufficient knowledge to determine an optimal policy. See, for example, Brito, Sheshinski, and Intriligator (1991), Becker and Starczak (1997), Ball and Lyne (2002), Scuffham and West (2002), Hill and Longini (2003), Patel, Longini, and Halloran (2005), Boulier, Datta, and Goldfarb (2007), and Althouse, Bergstrom, and Bergstrom (2010). To my knowledge, the only studies that have considered choice of vaccination policy under uncertainty are Tanner, Sattenspiel, and Ntaimo (2008) and Manski (2010).

In my earlier work and here, consideration of optimal policy is a prelude to study of vaccination under uncertainty. Section 3.1 poses the unconstrained optimization problem that I studied earlier. Section 3.2 formalizes decisions to approve and mandate vaccines as constrained optimization problems.

3.1. Unconstrained Optimization

Consider a planner who must choose a vaccination rate for a large population of observationally identical persons. The term "observationally identically" does not mean that persons actually are identical,

only that the planner cannot distinguish them. 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.[5]

Consideration of a large population simplifies analysis in important respects. The planner can ignore finite-sample statistical variation in the outcome of random treatment assignment. 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 contain infinitely many members. Standard arguments using Laws of Large Numbers enable their approximate application to large finite populations.

Suppose that the vaccine under consideration generates a protective immune response in a vaccinated person with probability $\lambda \in (0, 1]$, preventing this person from becoming ill and from infecting others. Contrariwise, the vaccine confers no immunity with probability $1 - \lambda$. Then λ measures the direct effect of vaccination.[6] If fraction v of the population are vaccinated, the *effective vaccination rate* is λv and the fraction $1 - \lambda v$ of the population are susceptible to the disease.

The indirect-response function $p(\cdot)$: $[0, 1] \to [0, 1]$ gives the probability that a susceptible person becomes ill when a specified fraction of the population are effectively vaccinated. If vaccination yields an indirect preventive effect, then $p(\cdot)$ is a decreasing function. When the effective vaccination rate is λv, the fraction of the population who become ill is $p(\lambda v)(1 - \lambda v)$.

---

[5] The assumption of observationally identical persons is maintained in some of the literature on optimal vaccination, such as Althouse, Bergstrom, and Bergstrom (2010) and Brito, Sheshinki, and Intriligator (1992). Findings there and in the present paper extend directly to settings where the population partitions into multiple isolated groups, each defined by a different value of observed covariates. Analysis is more complex when interacting persons with different observed covariates vary systematically in susceptibility or infectiousness. Studies of such settings often consider a constrained optimization problem in which the problem is to allocate a fixed supply of vaccine among persons with different covariates. See, for example, Patel, Longini, and Halloran (2005) and Keeling and Shattock (2012).

[6] Manski (2010) mainly assumed that $\lambda = 1$ and sketched extension of the analysis to cases with $\lambda < 1$. Here I study from the outset the general case where the vaccine may be imperfect. Imperfect vaccination is a substantive concern when considering whether to approve or mandate some vaccines.

Suppose that the planner wants to minimize a social cost function with two additive components, the harm caused by illness and the cost of vaccination. Let a > 0 denote the mean social harm resulting from a case of illness and let c > 0 denote the mean social cost per vaccination, measured in commensurate units. The social cost of vaccination rate v is

(1) $S(v) = ap(\lambda v)(1 - \lambda v) + cv.$

The first term measures the social cost of illness and the second gives the social cost of vaccination. The optimization problem is invariant to the scale of $S(\cdot)$. Hence, without loss of generality, I shall let a = 1 and interpret c as the ratio of the mean social cost of vaccination to the mean social cost of illness. This done, the planner wants to solve the problem

(2) $\min_{v \in [0, 1]} p(\lambda v)(1 - \lambda v) + cv.$

This simple static optimization problem expresses the core tension of vaccination policy: a higher vaccination rate is more effective in preventing illness but is more costly. Similar planning problems have been posed in some past research on optimal vaccination, such as Brito, Sheshinski, and Intriligator (1991) and Keeling and Shattock (2012). 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) and Hill and Longini (2003). 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?"

The objective of preventing onset of an epidemic differs from minimization of social cost. In epidemiology, an epidemic is defined to occur when the infected fraction of the population increases with

time. In contrast, social cost function (1) abstracts from the dynamics of illness and considers the prevalence of illness in the population. This perspective is reasonable when considering vaccination to reduce the incidence of an endemic disease. It may be less appropriate when considering vaccination as a measure intended to prevent a foreign disease from becoming established in a population.

From the perspective of dynamic models of disease transmission, $p(\lambda v)(1 - \lambda v)$ is the steady-state fraction of infectious persons. Classical SIR models imply that $p(\cdot)$ is a decreasing function that may attain the value zero at some effective vaccination rate less than one, say e^*. Hence, the disease is eradicated if the actual vaccination rate is greater than or equal to e^*/λ, provided that $e^*/\lambda \leq 1$. In this case, the steady-state fraction of infectious persons is zero. Vaccination rate e^*/λ is called the critical vaccination or herd immunity threshold. See for example, Anderson and May (1991) and Fine (1993).

3.2. Approvals and Mandates as Constrained Optimization Problems

I have thus far supposed that the planner can choose any vaccination rate in the unit interval. Decisions to approve or mandate a vaccine pose a sequence of constrained optimization problems in which the planner must choose between two vaccination rates.

The decision to mandate an approved vaccine presents a choice between vaccinating the entire population (the mandate) and the vaccination rate, say v_d, that would be generated by decentralized decision making in the absence of the mandate. Using social cost function (1), the optimal mandate decision is

(3)   mandate if $p(\lambda)(1 - \lambda) + c \leq p(\lambda v_d)(1 - \lambda v_d) + cv_d,$

do not mandate if $p(\lambda)(1 - \lambda) + c \geq p(\lambda v_d)(1 - \lambda v_d) + cv_d.$

Vaccine approval presents a choice between a zero vaccination rate (rejection of the new vaccine)

and the vaccination rate that the health-care system will yield if the vaccine is approved. The latter rate is either 1 or v_d, depending on whether approval would or would not be followed by imposition of a mandate. Let v_a denote the vaccination rate that would occur following approval. Again using social cost function (1), the optimal approval decision is

(4) approve if $p(\lambda v_a)(1 - \lambda v_a) + cv_a \leq p(0),$

   reject if $p(\lambda v_a)(1 - \lambda v_a) + cv_a \geq p(0).$

This formalization of the planning problem poses approval and mandate as a sequence of binary decisions. I do so to facilitate discussion of the existing process in which FDA vaccine approval precedes state consideration of a mandate. If the same social cost function is used in both decisions, one can equivalently formalize the planning problem as a single decision comparing three vaccination rates: 0, v_d, and 1.

Implicit in criteria (3) and (4) is the assumption that decentralized vaccination decisions are effectively random. Thus, vaccinated persons do not differ systematically from unvaccinated ones in their cost of vaccination, the probability that vaccination successfully generates a protective immune response, susceptibility to the disease, or infectiousness. I maintain this assumption, which may or may not be realistic, because it substantially simplifies the planning problem, making knowledge of v_d sufficient to determine the welfare outcome of decentralized decision making. Without the assumption, the welfare achieved by decentralization would depend not only on the vaccination rate but on the composition of the vaccinated sub-population.

Having expressed vaccine approval and mandate as constrained optimization problems, it is natural to ask why a society would limit policy making to these decisions when unconstrained choice of the vaccination rate may yield greater social welfare. A possible answer is that having a planner choose a

vaccination rate other than 0 or 1 would violate a version of the ethical principle of "equal treatment of equals." Vaccinating an interior fraction of the population is consistent with the equal-treatment principle in the *ex ante* sense that all observationally identical people have the same probability of vaccination. However, it violates equal treatment in the *ex post* sense that only some persons ultimately are vaccinated. The ex post sense of equal treatment expresses a deontological consideration that is absent from social cost function (1) and, more generally, from the consequentialist welfare functions usually assumed in economic analysis of planning. See Manski (2009) for further discussion.

4. Approvals and Mandates with Partial Knowledge

Viewing vaccine approvals and mandates as constrained optimization problems, Section 3.2 showed that a planner can make optimal decisions if he has sufficient knowledge of the social cost of vaccination, the direct and indirect effects, and the vaccination rate that would result from decentralized decision making. Criterion (3) showed that to make an optimal mandate decision, it suffices to know $[c, \lambda, p(\lambda), v_d, p(\lambda v_d)]$. Criterion (4) showed that to make an optimal approval decision, it suffices to know $[c, \lambda, p(0), v_a, p(\lambda v_a)]$. The criteria are inequalities, so it is not necessary to know the precise values of the relevant quantities. The planner only needs to know enough to determine which inequalities hold.

Section 4.1 considers what knowledge a planner might have in practice. Sections 4.2 and 4.3 apply basic elements of decision theory to suggest criteria for making approval and mandate decisions with partial knowledge.

4.1. What Knowledge Might a Planner Have?

Performance of RCTs ordinarily is a necessary prelude to approval of a vaccine. RCTs can measure the immune responses that a vaccine generates in vaccinated subjects, which provides evidence on the direct effect of vaccination. RCTs can also reveal the clinical side effects and monetary costs of vaccination, two primary determinants of social costs. Thus, planners making approval and mandate decisions typically have substantial knowledge of the parameters λ and c.

In settings where a disease is endemic, it seems realistic to expect that a planner deciding whether to approve a new vaccine observes $p(0)$, the illness rate in the absence of vaccination. It also seems realistic to expect that a planner deciding whether to mandate a previously approved vaccine observes the decentralized vaccination rate v_d and illness rate $p(\lambda v_d)$ that occur in the absence of the mandate. However, the former planner cannot observe v_a and $p(\lambda v_a)$; these quantities become observable only after a vaccine is approved. The latter planner cannot observe $p(\lambda)$; this becomes observable only after a mandate is enacted.

In the absence of direct observation, what knowledge might these planners realistically have of the illness rates $p(\lambda v_a)$ and $p(\lambda)$? A basic assumption of epidemiological analysis of infectious disease is that raising the vaccination rate increases the indirect protective effect of vaccination. Hence, the indirect-response function $p(\cdot)$ is decreasing. When combined with available empirical evidence, this monotonicity assumption yields inequalities on $p(\lambda v_a)$ and $p(\lambda)$. A planner in the approval setting knows $p(0)$ and hence can deduce that $0 \leq p(\lambda v_a) \leq p(0)$. A planner in the mandate setting knows $p(\lambda v_d)$ and hence can deduce that $0 \leq p(\lambda) \leq p(\lambda v_d)$. These inequalities apply the general idea of monotone treatment response (Manski, 1997).

Beyond the monotonicity assumption, epidemiological analysis yields no clearly credible assumption about the shape of $p(\cdot)$. Alternative models of disease transmission imply different shapes for the function. Researchers have typically chosen models to be analytically and computationally tractable, with only limited understanding of the realism of their assumptions. Hence, a planner may not know much about the shape

16

of $p(\cdot)$ beyond its monotonicity. I assume no further knowledge in the analysis below.

It remains to consider what a planner making an approval decision might know about the vaccination rate v_a that would occur if the vaccine is approved. Some relevant empirical evidence may be available from observation of vaccination rates for other diseases with previously approved vaccines, but there may be no firm basis for extrapolation from these data to new vaccines. Nevertheless, in the interest of simple analytics, I assume below that the planner knows v_a.

4.2. Mandates

The mandate decision is somewhat simpler than the approval decision, so I consider mandates first. The difficulty is incomplete knowledge of $p(\lambda)$. A mandate is consequential only if decentralized decision making would yield incomplete vaccination of the population. Hence, I assume that $v_d < 1$.

4.2.1. Optimization with Partial Knowledge

Rewriting criterion (3), a mandate is optimal if

$$(3') \qquad p(\lambda)(1 - \lambda) \ \leq \ p(\lambda v_d)(1 - \lambda v_d) - c(1 - v_d)$$

and is sub-optimal otherwise. Observation of $p(\lambda v_d)$ and the assumption of monotone treatment response imply that $0 \leq p(\lambda) \leq p(\lambda v_d)$. This bound on $p(\lambda)$ suffices to determine the optimality of a mandate if either of two conditions hold:

$$(5a) \qquad p(\lambda v_d)(1 - \lambda) \ \leq \ p(\lambda v_d)(1 - \lambda v_d) - c(1 - v_d) \ \Rightarrow \ \text{mandate is optimal,}$$

$$(5b) \qquad 0 \ > \ p(\lambda v_d)(1 - \lambda v_d) - c(1 - v_d) \qquad \qquad \Rightarrow \ \text{mandate is not optimal.}$$

The inequality on the left-hand side of (5a) implies that inequality (3') holds for all feasible values of $p(\lambda)$; hence, a mandate must be optimal. The inequality on the left-hand side of (5b) implies that (3') holds for no feasible value of $p(\lambda)$; hence, a mandate cannot be optimal. If neither inequality holds, a mandate is optimal for some feasible values of $p(\lambda)$ but not for others.

Conditions (5a) and (5b) show that a planner with incomplete knowledge of the illness rate that would occur with a mandate may nonetheless be able to determine if a mandate is optimal. One or the other condition necessarily holds if $\lambda = 1$. In this case, $p(\lambda)(1 - \lambda) = 0$ for all values of $p(\lambda)$. Hence, indeterminacy of the optimum can happen only if the vaccine has an imperfect direct effect.

To study cases with $\lambda < 1$, it is instructive to rewrite (5) in the equivalent form

(5a') $c \leq \lambda p(\lambda v_d)$ $\Rightarrow$ mandate is optimal,

(5b') $c > p(\lambda v_d)[(1 - \lambda v_d)/(1 - v_d)]$ $\Rightarrow$ mandate is not optimal.

Condition (5a') considers the worst-case scenario in which increasing the vaccination rate from v_d to 1 has no indirect preventive effect. The inequality shows that, all else equal, a mandate is still optimal if the cost c of vaccination is sufficiently small. Condition (5b') considers the best-case scenario in which increasing the vaccination from v_d to 1 has the strongest possible indirect effect. This inequality shows that, all else equal, a mandate cannot be optimal if c is sufficiently large. The optimality of a mandate is indeterminate if c has an intermediate value; that is, if

(6) $\lambda p(\lambda v_d) < c \leq p(\lambda v_d)[(1 - \lambda v_d)/(1 - v_d)]$.

To determine the optimal policy when (6) holds, the planner must know more about $p(\lambda)$ than the bound $0 \leq p(\lambda) \leq p(\lambda v_d)$. In some cases, a planner may judge available epidemiological models of disease

transmission to be sufficiently credible that he thinks it reasonable to use the models to predict $p(\lambda)$ and then possibly tighten the bounds. In the best case, a planner might use a model to make a point prediction and view the prediction as accurate. However, it typically is unrealistic to view any point prediction as accurate. It may be more credible to use epidemiological models to tighten the bound on $p(\lambda)$. One might make predictions using multiple a priori plausible models, determine the range of findings, and use this range as the bound on $p(\lambda)$.[7] This done, the analysis resulting in inequalities (5) and (6) can be repeated using the tightened bound.

4.2.2. Decision Making When the Optimal Policy is Indeterminate

It remains to consider decision making when the optimality of a mandate is indeterminate. Decision theory gives no consensus prescription, but it suggests various reasonable criteria.

When economists have studied planning with partial knowledge, it has been standard to assert a subjective probability distribution over unknown decision-relevant quantities (aka states of nature) and propose choice of an action that maximizes subjective expected utility. Researchers studying vaccination policy have recently begun to use this criterion to cope with uncertainty. Tanner *et al.*(2008) exposits it under the name *stochastic programming*.

In the present setting, one would minimize subjective expected cost. Let S denote the set of states of nature, indexing feasible values of the unknown quantity $p(\lambda)$. Let π denote the subjective distribution placed on S. For $s \in S$, let $p_s(\lambda)$ denote the value of $p(\lambda)$ in state s. Let $f(\cdot): R \to R$ be a specified increasing function expressing the risk preference of the planner. Then minimization of subjective expected cost means replacement of optimality criterion (3) with the computable criterion

---

[7] Studies of disease transmission sometimes report sensitivity analyses that conjecture multiple parameter values within a specified parametric model. See, for example, Weycker *et al.* (2005). As far as I am aware, epidemiological research has not sought to generate set-valued predictions conjecturing a broad set of plausible models, as is the norm in partial identification analysis (Manski, 2007).

(7)   mandate if $\int_S f[p_s(\lambda)(1 - \lambda) + c]d\pi \leq f[p(\lambda v_d)(1 - \lambda v_d) + cv_d]$,

     do not mandate if $\int_S f[p_s(\lambda)(1 - \lambda) + c]d\pi \geq f[p(\lambda v_d)(1 - \lambda v_d) + cv_d]$.

Criterion (7) simplifies if f(·) is the identity function, in which case the planner is risk neutral. It reduces to

(8)   mandate if $[\int_S p_s(\lambda)d\pi](1 - \lambda) + c \leq p(\lambda v_d)(1 - \lambda v_d) + cv_d$,

     do not mandate if $[\int_S p_s(\lambda)d\pi](1 - \lambda) + c \geq p(\lambda v_d)(1 - \lambda v_d) + cv_d$.

Thus, the planner acts as if the unknown p(λ) equals its known subjective mean $\int_S p_s(\lambda)d\pi$.

Use of the expected utility criterion to make policy choices with partial knowledge is reasonable when a planner has a credible basis for asserting a subjective probability distribution on unknown 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 various problems of this type, where a planner faces ambiguity; Manski (2011) reviews basic ideas and gives applications. In particular, Manski (2010) studied the unconstrained planning problem of Section 3.1 as choice of a vaccination rate under ambiguity, focusing on the minimax and the minimax-regret criteria. These criteria are complex to evaluate when the planner can choose any vaccination rate, but they simplify when he makes a mandate decision.

A planner who uses the minimax criterion and who assumes that p(·) is monotone acts as if he faces the worst-case scenario in which increasing the vaccination rate from v_d to 1 has no indirect preventive effect. Thus, he acts as if p(λ) = p(λv_d). This done, he replaces optimality criterion (3) with the computable criterion

(9)   mandate if $p(\lambda v_d)(1 - \lambda) + c \leq p(\lambda v_d)(1 - \lambda v_d) + cv_d$,

     do not mandate if $p(\lambda v_d)(1 - \lambda) + c \geq p(\lambda v_d)(1 - \lambda v_d) + cv_d$.

This criterion simplifies to

(9') mandate if $c \leq \lambda p(\lambda v_d)$,

do not mandate if $c \geq \lambda p(\lambda v_d)$.

Thus, the minimax criterion resolves indeterminacy in favor of no mandate.

To form the minimax-regret criterion, let S_m and S_d respectively be the subsets of S in which a mandate yields strictly lower and higher social cost than decentralized decision making. For states in S_m, regret is zero if the planner chooses the mandate, which is optimal in these states. Regret in state $s \in S_m$ takes the positive value $[p(\lambda v_d)(1 - \lambda v_d) + cv_d] - [p_s(\lambda)(1 - \lambda) + c]$ if he choose not to mandate the vaccine. Analogously, for states in S_d, regret is zero if the planner chooses not to mandate the vaccine and takes the positive value $[p_s(\lambda)(1 - \lambda) + c] - [p(\lambda v_d)(1 - \lambda v_d) + cv_d]$ if he choose to mandate it.

Maximum regret in each subset of states is

(10a) $R_m \equiv \underset{s \in S_m}{\max} \ [p(\lambda v_d)(1 - \lambda v_d) + cv_d] - [p_s(\lambda)(1 - \lambda) + c]$,

(10b) $R_d \equiv \underset{s \in S_d}{\max} \ [p_s(\lambda)(1 - \lambda) + c] - [p(\lambda v_d)(1 - \lambda v_d) + cv_d]$.

The maximum regret of mandate and no-mandate decisions across all states are R_d and R_m respectively. Hence, a minimax-regret decision is

(11) mandate if $R_d \leq R_m$,

do not mandate if $R_d \geq R_m$.

To obtain an explicit solution, observe that the maximum in (10a) occurs in the best-case scenario

21

for vaccination, where $p_s(\lambda) = 0$. The maximum in (10b) occurs in the worst-case scenario where $p_s(\lambda) = p(\lambda v_d)$. Hence,

(12a) $R_m = p(\lambda v_d)(1 - \lambda v_d) - c(1 - v_d),$

(12b) $R_d = (1 - v_d)[c - \lambda p(\lambda v_d)].$

Inserting these values into (11) yields the criterion

(13)   mandate if $2c(1 - v_d) \leq p(\lambda v_d)[1 - 2\lambda v_d + \lambda],$

       do not mandate if $2c(1 - v_d) \geq p(\lambda v_d)[1 - 2\lambda v_d + \lambda].$

Comparison of (13) with (9') shows that the minimax and minimax-regret criteria are broadly similar but differ in their specifics. The two criteria are identical when $\lambda = 1$. In that case, both prescribe adoption of the mandate if $c \leq p(v_d)$.

4.2.3. Numerical Example

Numerical examples are useful to illustrate the findings. First consider a scenario where $\lambda = 1$. The vaccine has been approved and the planner observes that decentralized decision making yields vaccination rate $v_d = 0.5$ and illness rate among the unvaccinated $p(v_d) = 0.2$. In this setting, inequality (5a') holds if the mean cost of vaccination (relative to illness) is sufficiently small, specifically if $c \leq 0.2$. Inequality (5b') holds if $c > 0.2$. Hence, the planner can conclude that a mandate is optimal if $c \leq 0.2$ and is not optimal otherwise.

Indeterminacy of optimal policy can occur if $\lambda < 1$. Suppose that $\lambda = 0.75$, $v_d = 0.5$ and $p(\lambda v_d) = 0.2$. Applying (5'), the planner can conclude that a mandate is optimal if $c \leq 0.15$ and not optimal if $c > 0.25$.

22

Optimal policy is indeterminate if $0.15 < c \leq 0.25$.

In cases of indeterminacy, policy choice depends on the decision criterion that the planner uses. The expected utility criterion (7) can yield either policy depending on the transformation $f(\cdot)$ and subjective distribution π. A planner using the minimax criterion (9') enacts the mandate if $c \leq 0.15$ but not if $c > 0.15$. One using the minimax-regret criterion (13) enacts the mandate if $c \leq 0.2$ but not if $c > 0.2$.

4.3. Approvals

4.3.1. Approval Followed by a Mandate

Analysis of the decision to approve a new vaccine is straightforward if the planner knows the vaccination rate v_a that would occur after approval. Suppose first approval would be followed by a mandate. Then approval is equivalent to enacting a mandate in a setting where the observed decentralized vaccination rate is $v_d = 0$ and the observed decentralized illness rate is $p(\lambda v_d) = p(0)$. The analysis of Section 4.2 applies directly to this case.

Application of (5) shows that the available knowledge suffices to determine the optimality of approval if either of two conditions hold:

(14a) $p(0)(1 - \lambda) \leq p(0) - c \implies$ approval is optimal,

(14b) $0 > p(0) - c \qquad \implies$ approval is not optimal.

The optimality of approval is indeterminate if

(15) $\lambda p(0) < c \leq p(0)$.

23

Application of (9') and (13) show that the minimax and minimax-regret approval decisions are

(16) approve if $c \leq \lambda p(0)$,

reject if $c \geq \lambda p(0)$

and

(17) approve if $2c \leq p(0)(1 + \lambda)$,

reject if $2c \geq p(0)(1 + \lambda)$.

4.3.2. Approval Followed by Decentralized Vaccination

Now suppose the planner knows that approval would be followed by decentralized decision making. Approval is not equivalent to a mandate in this case, but the reasoning underlying the analysis of Section 4.2 applies with straightforward modifications. Approval is consequential only if decentralized decision making would yield a positive vaccination rate, so I assume that $v_d > 0$.

Applying (4), approval is optimal if

(18) $p(\lambda v_d)(1 - \lambda v_d) \leq p(0) - cv_d$

and is sub-optimal otherwise. Observation of $p(0)$ and the monotonicity assumption imply that $0 \leq p(\lambda v_d) \leq p(0)$. This bound on $p(\lambda v_d)$ suffices to determine the optimality of approval if either of two conditions hold:

(19a) $c \leq \lambda p(0)$ $\Rightarrow$ approval is optimal,

(19b) $c > p(0)/v_d \;\;\Rightarrow$ approval is not optimal.

The optimality of approval is indeterminate if neither inequality holds; that is, if

(20) $\lambda p(0) < c \leq p(0)/v_d.$

A planner who uses the minimax criterion to cope with indeterminacy acts as if $p(\lambda v_d) = p(0)$. Thus, he replaces optimality criterion (4) with the computable criterion

(21) approve if $p(0)(1 - \lambda v_d) + cv_d \leq p(0),$

reject if $p(0)(1 - \lambda v_d) + cv_d \geq p(0).$

This simplifies to

(21') approve if $c \leq \lambda p(0),$

reject if $c \geq \lambda p(0).$

Thus, the minimax criterion resolves indeterminacy in favor of rejection of the vaccine.

To form the minimax-regret criterion, let S_a and S_r respectively be the subsets of S in which approval yields strictly lower and higher social cost than rejection. For states in S_a, regret is zero if the planner approves the vaccine, which is optimal in these states. Regret in state $s \in S_a$ takes the positive value $p(0) - p_s(\lambda v_d)(1 - \lambda v_d) - cv_d$ if he rejects the vaccine. Analogously, for states in S_r, regret is zero if the planner rejects the vaccine and takes the positive value $p_s(\lambda v_d)(1 - \lambda v_d) + cv_d - p(0)$ if he approves it.

Maximum regret in each subset of states is

(22a) $R_a \equiv \max\limits_{s \in S_a} p(0) - p_s(\lambda v_d)(1 - \lambda v_d) - cv_d$

(22b) $R_r \equiv \max\limits_{s \in S_r} p_s(\lambda v_d)(1 - \lambda v_d) + cv_d - p(0).$

The maximum regret of approval and rejection decisions across all states are R_r and R_a respectively. Hence,

a minimax-regret decision is

(23) approve if $R_r \leq R_a$,

reject if $R_r \geq R_a$.

To obtain an explicit solution, observe that the maximum in (22a) occurs in the best-case scenario

for vaccination, where $p_s(\lambda v_d) = 0$. The maximum in (22b) occurs in the worst-case scenario where $p_s(\lambda v_d)$

$= p(0)$. Hence,

(24a) $R_a = p(0) - cv_d$,

(24b) $R_r = cv_d - \lambda p(0)v_d.$

Inserting these values into (23) yields the criterion

(25) approve if $2cv_d \leq p(0)(1 + \lambda v_d)$,

reject if $2cv_d \geq p(0)(1 + \lambda v_d).$

4.3.3. Numerical Example

A numerical example again is useful to illustrate findings. Let the direct effect of vaccination be λ

= 1. The planner observes that the illness rate in the absence of vaccination is $p(0) = 0.25$. He knows that if the vaccine is approved, decentralized decision making will yield vaccination rate $v_d = 0.5$. In this setting, inequality (19a) holds if $c \leq 0.25$ and (19b) holds if $c > 0.5$. Hence, the planner can conclude that approval is optimal if $c \leq 0.25$ and not optimal if $c > 0.5$. Optimal policy is indeterminate if $0.25 < c \leq 0.5$.

In cases of indeterminacy, policy choice depends on the decision criterion that the planner uses. The expected utility criterion can yield either approval or rejection of the vaccine depending on the planner's risk preference and the subjective distribution he places on the illness rate that would occur following approval. A planner using minimax criterion (21') approves the vaccine if $c \leq 0.25$ but not if $c > 0.25$. One using minimax-regret criterion (25) approve the vaccine if $c \leq 0.375$ but not if $c > 0.375$.

5. FDA Approval Reconsidered

This paper has expressed vaccine approval and mandate as planning problems, has characterized optimal decision making, and has suggested criteria for decision making when incomplete knowledge makes optimal policy indeterminate. Various assumptions—linearity of social cost in the illness rate and the cost of vaccination, abstraction from the dynamics of illness, uncertainty only about the indirect preventive effect of vaccination—enabled a simple analysis. Relaxing these assumptions may complicate policy determination but the principles would remain intact.

Perhaps the most basic virtues of formal decision analysis are that it makes explicit the considerations that affect policymaking and makes it possible to examine how variation in these considerations would affect decisions. I believed it well justified to suggest that performance of formal decision analysis, whether the version developed in this paper or some variant, can improve prevailing vaccine approval and mandate procedures. The FDA approval process particularly warrants reconsideration.

27

I observed in Section 2 that in 2012 Congress required the FDA to institute a "structured benefit-risk assessment framework." The subsequent FDA response rejected use of formal decision analysis in favor of what the FDA calls a "structured qualitative approach" (U. S. Food and Drug Administration, 2013). The nature of this approach is difficult for an external observer to discern, it being carried out mainly through the deliberations of FDA advisory panels. The FDA makes public only a "Summary of Basis for Approval" when a new vaccine or drug is approved. This document presents some information on the safety and efficacy of the product but it does not explain how the agency balances these considerations. The agency does not make public even a summary of its rationale when it rejects an application for drug or vaccine approval. Thus, the "structured qualitative approach" of the FDA is not transparent.

While much about the FDA approval process is obscure, two important features are clear. First, agency concern with negative consequences of vaccination focuses on the health risk to the recipient and does not extend to monetary costs or other burdens borne by society. Second, agency concern with the benefits of vaccination focuses on the direct preventive effect on the recipient and does not extend to the indirect preventive effects on other members of the population.

I find it difficult to rationalize why the outcomes considered by a public agency should have such limited scope. The qualitative decision process of the FDA shields the agency from scrutiny. Analysis of approval as a planning problem would make explicit the full social costs and benefits of vaccination.

References

Althouse, B., T. Bergstrom, and C. Bergstrom (2010), "A Public Choice Framework for Controlling Transmissible and Evolving Diseases," *Proceedings of the National Academy of Sciences*, 107, 1696-1701.

Anderson R. and R. May (1991), *Infectious Diseases of Humans: Dynamics and Control*, Oxford: Oxford University Press.

Ball F. and O. Lyne (2002), "Optimal Vaccination Policies for Stochastic Epidemics among a Population of Households," *Mathematical Biosciences*, 177&178, 333–354.

Becker, N. and D. Starczak (1997), "Optimal Vaccination Strategies for a Community of Households," *Mathematical Biosciences*, 139, 117-132.

Boulier, B. T. Datta, and R. Goldfarb (2007), "Vaccination Externalities," *The B.E. Journal of Economic Analysis & Policy* 7, issue 1, article 23. 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.

Cox, D. (1958), *Planning of Experiments*, New York: Wiley.

Fine, P. (1993), "Herd Immunity: History, Theory, Practice," *Epidemiological Reviews*, 15, 265-302.

Fine, P, K. Eames, and D. Heymann (2011), "'Herd Immunity': A Rough Guide," *Clinical Infectious Diseases*, 52, 911–916.

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

Hudgens, M. and E. Halloran (2008), "Toward Causal Inference with Interference," *Journal of the American Statistical Association*, 103, 832-842.

Institute of Medicine, (2011), *Clinical Practice Guidelines We Can Trust*, Committee on Standards for Developing Trustworthy Clinical Practice Guidelines, Board on Health Care Services; editors, R. Graham *et al.* Washington, DC: National Academies Press.

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

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

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

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

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

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

Manski, C. (2011), "Choosing Treatment Policies under Ambiguity," *Annual Review of Economics,* 3, 25-49.

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

Manski, C. (2013b), "Diagnostic Testing and Treatment under Ambiguity: Using Decision Analysis to Inform Clinical Practice," *Proceedings of the National Academy of Sciences*, 110, 2064-2069.

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

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.

Rubin, D. (1978), "Bayesian Inference for Causal Effects: The Role of Randomization," *Annals of Statistics*, 6, 34-58.

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.

U. S. Congress (2012), "Food and Drug Administration Safety and Innovation Act," Public Law 112–144, 112th Congress, Washington, DC: Government Printing Office.

U. S. Food and Drug Administration (1994), "Summary for Basis of Approval Reference Number: 93-0395 Merck and Co. Varicella Virus Vaccine Live VARIVAX," www.fda.gov/downloads/biologicsbloodvaccines/vaccines/approvedproducts/ucm142826.pdf, accessed August 3, 2014.

U. S. Food and Drug Administration (2009), "Vaccine Product Approval Process," www.fda.gov/biologicsbloodvaccines/developmentapprovalprocess/biologicslicenseapplicationsblaprocess/ucm133096.htm, accessed August 3, 2014.

U. S. Food and Drug Administration (2013), *Structured Approach to Benefit-Risk Assessment in Drug Regulatory Decision-Making, Draft PDUFA V Implementation Plan*, patientnetwork.fda.gov/sites/default/files/fda_benefit-risk_draft_plan_final_for_posting.pdf, accessed August 3, 2014.

Washington State Board of Health (2006), *Immunization Advisory Committee: Criteria for Reviewing Antigens for Potential Inclusion in Wac 246-100-166*, sboh.wa.gov/Portals/7/Doc/immunizations/FinalRecs_2006-06-14.pdf, accessed August 4, 2014.

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