Incentives for surveillance and reporting of infectious disease outbreaks

Anup Malani
University of Chicago

Incentives for surveillance and reporting of infectious disease outbreaks

Ramanan Laxminarayan and Anup Malani[†]

March 14, 2006

Abstract

In spite of the obvious global public goods nature of warnings about infectious disease outbreaks, international legal requirements for reporting outbreaks remain weak and disclosure of outbreaks depends on the self-interest of nations. Using a simple game-theoretic model, we explore the incentives of countries to invest in disease surveillance and to report outbreaks to international health authorities. We evaluate existing and potential policy instruments available to encourage countries to detect and report disease outbreaks, including medical assistance to control outbreaks and financial and technical assistance with disease surveillance. We examine the effect of "preemptive sanctions," which are triggered by fears of an outbreak rather than a confirmed report of an outbreak, and the role of "rumor surveillance," which is a form of involuntary disclosure based on rumors and leaks. Finally, we study the implications of specificity (or false positives) on the incentive to surveil and report outbreaks. Our discussion is relevant to policy design, not just in the context of singular events such as avian influenza, but also for more routine problems such as reporting hospital-acquired infections and crimes.

Prologue

In November 2002, health authorities in Guangdong Province reported a cluster of atypical pneumonia cases to China's National Ministry of Health in Beijing. In late February 2003, an infected medical doctor from Guangdong spent a single night on the ninth floor of a Hong Kong hotel and infected at least 16 other persons visiting his floor. The others

---

[†] Laxminarayan is a Fellow at the Resources for the Future; Malani is an associate professor at University of Virginia Law School and Medical School. Please send comments to amalani@virginia.edu. We thank Eili Klein for valuable research assistance, Arin Dutta, Doug Lichtman and Scott Baker for detailed comments, and participants in workshops at the University of Chicago Law School, the World Bank, and Resources for the Future for feedback on earlier drafts.

included a tourist from Toronto, a flight attendant from Singapore and a businessman who later traveled to Vietnam. From this single event, severe acute respiratory syndrome (SARS) spread internationally. By May 2003, there were 7,000 infected in 30 countries. During the height of the global epidemic, more than 200 new cases were being reported each day. By the time the contagion was brought under control in June 2003, over 600 people died (WHO May 20, 2003).

Despite the early report from Guangdong, China did not report the outbreak of atypical pneumonia to the World Health Organization (WHO) until early February 2003. And though Chinese scientists had uncovered evidence linking SARS to a new coronavirus that same month (Science July 18, 2003), they did not allow WHO teams to visit Guangdong until early April. As a result, WHO scientists were not able to established the coronavirus link for themselves until mid-April (WHO July 4, 2003). Had the Chinese government reported the outbreak and its likely cause earlier, many hundreds of lives could have been saved. Had the disease been more virulent than it turned out to be, the effects of secrecy and delayed reporting by China could have been much more severe.

This pattern of inadequate disclosure is being repeated in the emerging crisis over the H5N1-strain of avian influenza. Thirty-seven nations have discovered outbreaks in their domestic or wild bird populations. Of these 7 countries have also discovered the infection in humans. Of the 175 confirmed cases of human infection since 2003, 96 have resulted in death (WHO, Mar. 3, 2006). Although the H5N1 strain has mastered animal-to-animal and animal-to-human transmission, it has not, fortunately, mastered human-to-human transmission. If that occurs, WHO conservatively estimates that two to 7.4 million people might die (WHO Oct. 14, 2005). WHO's strategy for preventing an epidemic relies on rapid vaccination and quarantine of the immediate neighborhood where human-to-human transmission is detected. That strategy relies, however, on early detection of an outbreak. (Science, Jan. 20, 2006; but see Science Feb. 24, 2006) While most countries, alarmed by the possible impact of an epidemic, have cooperated with WHO (but see VOA News Mar. 3, 2006), many lack the basic surveillance infrastructure to detect an outbreak (Science Feb. 17, 2006). Moreover, China, which faces a high risk of an outbreak and has the capacity to detect an outbreak, is not fully cooperating with WHO. Although it has reported outbreaks with greater speed than in the case of SARS, China has failed to report outbreaks in certain provinces (NYT Feb. 1, 2006) and has delayed sharing the blood samples that are vital to the development of a vaccine targeted precisely at the current strain of bird flu (WSJ Dec. 23, 2005; WSJ Dec. 27, 2005).

1. Introduction

Although information on disease outbreaks is a global public good, international legal requirements for surveillance and disclosure of outbreaks remain weak. The first international convention ever to be globally adopted was a 1851 accord to contain cholera. But since then, there has been very slow progress in strengthening the accord or expanding its scope. Until May 2005, International Health Regulations only required countries to report outbreaks of cholera, plague and yellow fever. Current regulations cover a broader array of diseases and provide standards for surveillance systems and timely disclosure of outbreaks (WHO May 23, 2005). Unfortunately, because non-reporting carries no penalties, warnings continue to depend on the self-interest of nations.

This paper explores the incentives of countries to surveil their own populations for infectious diseases and report outbreaks to international health authorities such as WHO. At first blush, countries face conflicting incentives to disclose outbreaks. On the one hand, reporting may trigger trade sanctions that can impose large economic costs. For example, when Peru reported an outbreak of cholera in 1991, its South American neighbors imposed an immediate ban on Peruvian food products. The subsequent loss of $790 million in food sales and tourism revenues far exceeded the domestic health and productivity costs of the epidemic. As the Peruvian Minister of Health noted, "...nothing compares to the loss of markets [other countries] took away from us in a difficult time" (Panisset 2000, p. 150).[1] On the other hand, countries may also report an outbreak in order to obtain international assistance for containing the outbreak before it develops into a full blown epidemic. This could be an important reason why Vietnam and Turkey have been quick to report incidents of avian flu. Incentives to report an outbreak, however, tell only half the story. A disease must be detected to be reported and countries control not only reporting but also investments in surveillance and detection. If a country does not want to report information about an outbreak, it may have an incentive to limit surveillance. This paper presents a simple game-theoretic model to capture these basic dynamics.

The paper also extends our understanding of disclosure incentives in five ways. First, we explore some other policy instruments available to

[1] Ironically, before China itself became a serious source of avian flu outbreaks, it and Vietnam imposed trade sanctions on other countries that reported outbreaks. This prompted a visiting U.S. official to urge caution lest sanctions discourage nations from reporting outbreaks (WSJ Nov. 11, 2003).

WHO to encourage countries to detect and report disease outbreaks. These include punitive sanctions for non-reporting and subsidies for surveillance. Because WHO can only verify that an outbreak was not reported after it escalates into an epidemic, punitive sanctions can only be imposed on non-reporting countries that have already suffered grave social and economic losses. But at that point a sanction will have little additional marginal cost; therefore, punitive sanction are unlikely to have significant incentive effects. As for surveillance subsidies, we find that they mainly displace private investment in surveillance. If subsidies are conditioned on reporting, countries will have a further incentive to curtail their own surveillance. If a country does not surveil, it cannot report an outbreak.

Second, we question whether trade sanctions are unambiguously harmful to reporting and disease control. It is often forgotten that sanctions—especially restrictions on travel—can limit the spread of a disease. The more important observation, however, is that while some sanctions ("*ex post* sanctions") follow the report of a disease outbreak, others ("preemptive sanctions") are triggered by the fear or risk of an outbreak. Preemptive sanctions can be informal, as when demand for American beef fell 80% after initial reports that a U.S. cattle was infected with mad cow disease (Blayney 2005), or formal, as when the U.S. banned poultry imports from all countries that had birds infected with avian flu even (WSJ Nov. 21, 2005). In neither case were trade-restrictions limited to countries reporting human infections, which is what really concerned the sanctioning countries. We find that preemptive sanctions encourage countries to report disease outbreaks. If a country is hit with preemptive sanctions even if it does not report, the relative marginal cost of reporting will fall.

Indeed, we suggest that WHO piggy-back on preemptive sanctions in order to encourage even more reporting. Preemptive sanctions are often based on beliefs that are imperfect. The problem is that countries who feel that such sanctions unfairly punish them cannot credibly demonstrate to the rest of the world that they are less likely to experience an outbreak that is feared. After all, it is very easy to fake negative diagnostic test results. WHO could solve this problem by offering to serve as an audit agency that confirms negative test results. This would trigger a "reverse-lemons" effect or an "unraveling" of non-disclosure (Grossman 1981; Milgrom 1981). Once some low-risk countries are audited, the rest of the world will revise upward its threat assessment for remaining non-audited countries. As a result, new countries will fall into the category of low-risk, non-audited countries. These countries may then find it cost-effective themselves to be audited. The process will repeat until only the most recalcitrant, high-risk countries refuse to admit WHO inspectors.

Third, conventional wisdom in the public health community is that investments in surveillance technology ought to focus on sensitivity, i.e., on increasing the probability of a positive test result given the patient has an infection. The logic is that if one does not detect a disease, one cannot control its spread. We question the resulting near-exclusive attention given to lowering the number of false negatives. Because tests that lack specificity and incorrectly signal an outbreak trigger *ex post* trade sanctions and bring medical assistance that is of little value, they discourage reporting. Therefore, emphasizing a more balanced investment in sensitivity and specificity is likely to increase public information on disease outbreaks. Another way to put this is that improving the technological specificity of testing will improve the behavioral sensitivity of reporting.

Fourth, we discuss the phenomenon of "rumor surveillance." This is the process by which the international community obtains information about the risk or probability of an outbreak in a country that does not report its outbreaks. Rumor surveillance relies on media leaks, rumors, and other indirect, unofficial indicators of an outbreak (Samaan et al. 2005). One of the criticisms of rumor surveillance is that it is prone to false positives (Harris 2006). In our view, however, this may be a blessing in disguise. There are two ways that information about outbreaks can spread: unofficial rumor surveillance or official surveillance and reporting. Either way, this information triggers sanctions: rumor surveillance informs preemptive sanctions and official reports trigger *ex post* sanctions. One advantage of official reporting is that it brings medical assistance. If that is not sufficient to induce reporting, then inaccurate rumor surveillance might help. False positives in rumors inflate preemptive sanctions, causing countries to want to admit WHO auditors in order to limit sanctions to instances where there are actual outbreaks.

This paper relates to the economics literature on principal-agent problems. One contribution is to consider the case where the agent's cost of effort is falling in her ability or quality. In our application there are multiple agents and each agent's cost of disclosure (effort) takes the form of sanctions from fellow agents proportional to the risk of infection (lower quality). The generalizable insight of the paper is that if sanctions among agents are triggered by beliefs about type rather than by revelation of type, the inter-agent dynamic can reduce the principal's agency costs. This result has many obvious applications, such as encouraging hospitals to report infection rates to national authorities, incentivizing cities to disclose their crime rates, inducing employees to "blow-the-whistle" on other employees and encouraging citizens to report local crimes to police.[2] A second contribution of the paper is to examine the value of different types of audit technologies. In particular, it examines how false positives versus false negatives affect an agent's willingness to accept an audit in the case where the principal does not have the right to audit at-will. Our findings generalizes to other contexts such as International Atomic Energy Agency audits of countries' nuclear weapons programs.

The remainder of the paper is organized as follows. Section 2 presents the basic model and demonstrates the conventional wisdom regarding incentives to report outbreaks. Section 3 explores the various instruments available to WHO to encourage surveillance and reporting. Section 4 demonstrates how preemptive sanctions can encourage reporting. Section 5 examines the value of surveillance technology that targets false negatives versus false positives. It also examines the role of rumor surveillance. The appendix offers derivations of the results in the main text.

2. Basic model

Consider a country with a small, non-zero probability of experiencing an outbreak of infectious disease. It must decide how much money to invest in surveillance and, if it discovers an outbreak, whether to report this information to WHO. These decisions are made over time. First, the country must decide how much (θ) to invest in surveillance or disease testing. Second, the country experiences an outbreak with probability p_0. An outbreak is defined as a geographically localized infection affecting a

[2] One application that has been the subject of prior work is inducing firms to disclose criminal violations by their employees. Articles by Arlen (1994) and Kaplow and Shavell (1994) have considered the problem of holding firms liable for the illegal behavior of their employees. Arlen argues that holding firms liable may have the perverse effect of discouraging reporting because it is reporting that triggers criminal punishment for the firm. This is identical to the notion that *ex post* trade sanction discourage the reporting of diseases. Kaplow and Shavell note that one can negate the perverse effect if criminal sanctions are lower for firms that turn their employees in to the authorities. In other words, Kaplow and Shavel argue that if inter-agent sanctions separately penalize a bad-type agent and a non-reporting agent, agents can be induced to report. This is similar to the role that punitive sanctions for non-reporting play in our model. Neither model considers the value of preemptive sanctions, i.e., inter-agent sanctions triggered by beliefs rather than disclosures about type. This is particularly valuable when the principal cannot coordinate its strategy towards one agent with other agents.

relatively small number of humans.[3] Third, the country observes the results of its surveillance program. If the country experienced an outbreak, its surveillance system will identify the outbreak with probability $q(\theta)$. A larger investment in surveillance increases the probability of detection or sensitivity, though at a diminishing rate: $q'(\theta) > 0$, $q''(\theta) < 0$. If there were no outbreak in stage two, the surveillance system would not report an outbreak. In other words, for now, it is assumed there may be false negatives $(1 - q(\theta))$ but no false positives.

Fourth, if its surveillance network detected an outbreak, the country must decide whether to report the outbreak to WHO. If it reports, it will receive medical assistance c from WHO, but will also suffer *ex post* trade sanctions from other countries – we'll call them the "rest of the world" or ROW – with cost S_0. The purpose of the sanctions is to stop the spread of the epidemic to the ROW. Fifth, the outbreak becomes an country-wide epidemic with probability $p_1(m + c)$, where m is the amount of the country's own medical resources. Because medical assistance from WHO may help control the outbreak, the probability of an epidemic falls (though at a diminishing rate) with assistance: $p_1'(m + c) < 0$, $p_1''(m + c) > 0$. An epidemic is far more harmful than an outbreak. We assume it will kill y_1 residents. The country values this at $\beta(1 - \alpha)y_1$, where β is the country's discount factor and $\alpha \in [0, 1]$ measures how *in*senstive the country's government is to its population's welfare. Because the country cannot hide an epidemic, it will trigger sanctions S_1 from the ROW, with cost βS_1. For simplicity, we define $E = \beta[(1 - \alpha)y_1 + S_1]$ to be the total domestic cost of an epidemic.

The country only observes whether there is a positive test result, not whether it suffered an outbreak. Therefore, payoffs depend on two states – positive or negative test results – and a country's decision whether to report if there is a positive test result. Table 1 describes the payoffs. If the country observes a positive test result but does not report, it faces a higher risk that the outbreak will grow into an epidemic. It if does report, it will face immediate sanctions, but a smaller risk of an epidemic due to emergency medical assistance from WHO. If the country observes a negative test result, it has nothing to report and will receive no medical assistance to prevent an epidemic even if there actually is an outbreak.[4]

[3] If there is an outbreak, a certain number of residents will die. The country cares about this loss, but because it cannot avoid this loss, we ignore it in our theoretical analysis.

[4] We could relax the model to allow the country to report a positive test result even if it observed a negative test result. In this case, the payoff to a positive report would be $-\theta - S_0 - d(\theta, p_0)\, p_1(c)\, E$, where $d(\theta, p_0) = \{[1 - q(\theta)]\, p_0\}/\{1 - q(\theta)\, p_0\}$. A country would not report so long as the probability of detection is sufficiently high ($q(\theta) > 1 - [S_0$ /

WHO and the ROW observe neither a country's investment in surveillance nor an outbreak. They do not even observe test results unless the country reports them to WHO. We assume, however, that WHO and the ROW can observe full-blown epidemics. These are difficult to conceal, even under authoritarian regimes.

Reporting. Because we employ backward induction to solve the game, we consider the reporting decision before the surveillance decision. A country will report a positive test result if the cost of reporting (sanctions) is lower than the benefit (lower likelihood of an epidemic):

$$S_0 \leq \Delta p_1(m, m + c)\, E = \Delta p_1(m, m + c)\, \beta((1 - \alpha)y_1 + S_1) \tag{1}$$

where $\Delta p_1(a, b) = p_1(a) - p_1(b)$ is the benefit of medical assistance in terms of reducing the probability of an epidemic. (This benefit rises in the level of medical assistance: $\Delta p_1'(m + c) > 0$.) A country is more likely to report the smaller are *ex post* sanctions and the greater is the amount of medical assistance. They are also more likely to report the greater the cost of an epidemic, i.e., the speed with which an epidemic might erupt, the number of lives that would be lost, and how much the government values those lives.

There are two, more subtle conclusions that follow from (1). First, if a country has its own medical resources, WHO medical assistance does not provide as strong an incentive to report an outbreak. Conversely, if the country does not rely on foreign trade or tourism for its economic wellbeing, the sanctions-related costs of reporting are likely to be negligible Second, the probability with which a country's surveillance network detects an outbreak does not affect the reporting decision. This probability does not differ whether the country does or does not report because we are conditioning on a positive test result.

Surveillance. Suppose that a country has decided it will report a positive test result. The country will choose the amount to invest in surveillance so as to maximize the benefit of surveillance given its commitment to report positive test results:

$$\max_{\theta}\ -\theta - p_0 q(\theta)[S_0 + p_1(m + c)E] - p_0[1 - q(\theta)]p_1(0)E \tag{2}$$

$p_0 p_1(0,c)E$]) or the probability of an outbreak is sufficiently low ($p_0 < S$ / {[$1 - q(\theta)$] $p_1(0,c)E$}). In reality, WHO is likely to require that a positive test result is independently validated. Otherwise it would be wasting scarce medical resources on a country that is unlikely to have an outbreak. In this case, the payoff to a positive report would be $-\theta - S_0 - d(\theta, p_0)\, p_1(0)\, E$ and not reporting would strictly dominate reporting.

The optimality condition[5] is

$$q'(\theta) = \frac{1}{p_0\left[\Delta p_1(0, m+c)E - S_0\right]} \tag{3}$$

Because the marginal benefit of surveillance is a higher marginal probability of detection $q'(\theta)$ times the net benefit of reporting (the denominator in (2)), it follows that the conditions that encourage reporting also encourage investment in surveillance.

3. WHO's policy levers

WHO's central objective is to lower the probability of global epidemic. The primary instrument in WHO's toolkit is the promise of medical expertise and resources in the event of a reported outbreak. In the previous section we demonstrated that such assistance, by reducing the risk of an epidemic, increase the incentive to report a positive test result. In this section we explore two other instruments available to WHO. One instrument is the use of WHO's financial and technical resources to build up a country's domestic surveillance system. In our model, this takes the form of a surveillance subsidy θ_w. This subsidy may be given without strings or conditioned on allowing WHO access to information generated by a country's surveillance network. We assume that WHO still cannot observe a country's own investment θ in surveillance. Another instrument is punitive sanctions against countries that suffer an outbreak but delay reporting it. This instrument, which is not presently employed by WHO, takes the form of a punitive sanction S_p imposed when a country does not report an outbreak but experiences an epidemic.

Surveillance subsidy. Surveillance subsidies take the form of monetary or in-kind assistance such as provision of epidemiological training, diagnostic tests, and software to track disease. There are two types of surveillance subsidy policy we consider. One is an unconditional subsidy and the other is a subsidy conditioned on WHO audits. In general, each dollar of unconditional surveillance subsidy θ_w displaces a dollar of private investment in surveillance because it does not change either the cost of private investment or the benefit to the total amount of WHO and private investment in surveillance. This is evident from the fact that the optimality condition for total investment with the subsidy

[5] The second order condition is satisfied by the condition that the country chooses to report and the assumption that $q''(\theta) < 0$.

$$q'(\theta + \theta_w) = \frac{1}{p_0[\Delta p_1(0, m + c)E - S_0]} \quad (4)$$

is the same as the optimality condition for private investment without the subsidy (3). The only exception to this result is the case where the subsidy is greater than the level of private investment without the subsidy, i.e., $\theta_w > \theta^*$, where θ^* is the solution to (3). A second problem with the unconditional subsidy is that it does not change the incentives to report. The reporting decision is made after detection. At that point, all that matters are the cost of sanctions and the benefit of medical assistance.

A better idea is a surveillance subsidy conditioned on WHO audits. A convenient way to implement this is to require that the country allow WHO to send its doctors to do the testing or the lab work on patient samples. This would allow WHO to know everything the country's surveillance system finds without the need for reporting. In essence, the condition is a pre-commitment to report. Obviously, if a country were planning to report in the absence of the subsidy, then it will report in the presence of a conditional subsidy. For such countries, a conditional subsidy operates just like an unconditional subsidy. Therefore, the benefit of a conditional subsidy should be judged by its effect on countries that would otherwise prefer not to report.

A perverse effect of a conditional subsidy on such countries is that, in general, the subsidy will reduce total investment in surveillance. The reason is that a country that does not otherwise want to report wants a way to reduce the cost of reporting. One approach is to reduce the amount of surveillance conducted within its borders. If WHO does not detect a disease, it cannot report anything that triggers an *ex post* sanction.[6] As before, the exception to this result is the case where the surveillance subsidy is greater than private surveillance in the absence of a subsidy, i.e., $\theta_w > \theta^*$ where θ^* solves $q'(\theta) = [p_0 \Delta p_1(0, m)]^{-1}$.

Nevertheless, a conditional subsidy may encourage reporting. A country will change its mind and report if the benefits outweigh the costs:

[6] Formally, the optimality condition for investment with a subsidy is $q'(\theta^{**} + \theta_w) = [p_0\{\Delta p_1(0, m+c)E - S_0\}]^{-1}$. The optimality condition without a subsidy is $q'(\theta^*) = [p_0 \Delta p_1(0, m)]^{-1}$. However, because the country does not want to report, we know $S_0 \geq \Delta p_1(m, m+c)E$. This implies $\Delta p_1(0, m+c)E - S_0 = \Delta p_1(0, m)E + \Delta p_1(m, m+c)E - S_0 \leq \Delta p_1(0, m)$. Therefore, the optimality conditions and the assumption that q is increasing and concave imply $\theta^{**} + \theta_w \geq \theta^*$.

$$(\theta^* - \theta^{**}) + p_0\,[q(\theta^{**} + \theta_w)\,\Delta p_1(0, m{+}c) - q(\theta^*)\,\Delta p_1(0, m)]\,E$$
$$> p_0\,q(\theta^{**} + \theta_w)\,S_0 \qquad (5)$$

where θ^* is the optimal level of private investment without the subsidy and mandatory reporting and $\theta^{**} \leq \theta^*$ is the optimal level of private investment with the subsidy. A benefit of accepting the subsidy is reduced private investment in surveillance. A cost is that a positive test result will be reported and trigger an *ex post* sanction. The big unknown is whether accepting a conditional subsidy will increase or decrease the likelihood of an epidemic. One the one hand, it will lower the probability of detection because it discourages overall surveillance: $q(\theta^{**} + \theta_w) \leq q(\theta^*)$. On the other hand, reporting does bring medical assistance to contain an outbreak. If the effect of medical assistance is larger, then more countries are likely to report. If the moral hazard with respect of surveillance is severe, then fewer countries will report. In either case, a conditional subsidy will yield more reporting than no subsidy because countries that would have reported without a subsidy will continue to do so with a subsidy. Of course, the moral hazard from even unconditional subsidies means that WHO will be wasting money on countries that were going to report even without the subsidy.

Punitive sanctions. Because WHO only observes a country's report of an outbreak or an epidemic, it can only impose a punitive sanction on countries that experience an epidemic without first reporting an outbreak. Such a sanction would alter a country's positive-test/no-report payoff as follows: $-\theta - p_1(0)\,[E + S_p]$. This expands the range of *ex post* trade sanctions over which a country will report:

$$S_0 \leq \Delta p_1(0, c)\,E + p_1(0)S_p, \qquad (6)$$

which is just another way of saying the country is more likely to report.

In a reporting equilibrium, punitive sanctions are only triggered in the state where there is an outbreak but no positive test result. Reporting countries have an incentive to reduce the probability of this state by investing in surveillance. In other words, punitive sanctions encourage surveillance. This is evident from the optimality condition for investment,

$$q'(\theta) = \frac{1}{p_0\left[\Delta p_1(0, m + c)E - S_0 + p_1(0)S_p\right]} \qquad (7)$$

and the assumption that $q(\theta)$ is increasing and concave.

Punitive sanctions can take two forms. One is a trade restriction: WHO convinces the ROW to limit trade with an offending country even after pathogens may have spread to other countries. This is unlikely to be an effective deterrent. Put aside the problem that it may not be politically feasible to sanction a country while its population is suffering an epidemic. Prior to the punitive sanctions, the ROW is likely to have adopted *ex post* trade sanctions to limit the spread of the epidemic. Further sanctions will have little marginal impact on the country's economic condition. In technical terms, a country suffering an epidemic is already rubbing up against its participation constraint. Further penalties will simply cause the country to terminate relations with WHO.

A solution is to offer *ex ante* developmental assistance. The punitive sanction can then take the form of withdrawing the offer of assistance. Since external assistance is particularly valuable after a country has suffered an epidemic, the possible loss of support would be a strong inducement to participation. Assuming the participation constraint is a non-negative level of wealth, the amount of assistance required is $A = (S_0 + S_p) - W$, where W is the country's initial level of wealth. Poorer countries, which are less likely to meet the participation constraint but are also likely to have poor surveillance systems, will have to be offered a relatively greater amount of developmental assistance to participate.

The difficulty with withholding assistance is that it may not be politically and morally feasible to penalize a country by withdrawing assistance when a country has suffered an epidemic.[7] A difficulty with offering assistance to satisfy the participation constraint is that it is very expensive. Because WHO does not know whether the country is in the outbreak or no outbreak state, it must offer assistance in both states. It cannot even offer lower levels of assistance to countries with lower *ex ante* probabilities of an outbreak p_0. If it did, countries with a lower probability would lose less assistance and therefore have less incentive to report a positive test result than a country with a higher probability. However, given a positive test result, the value to WHO and the ROW of a report is the same both low and high probability countries.

Comparing policies. WHO has a choice of policy levers. It would be useful to know which lever is the most productive at inducing reporting and surveillance. The second and third columns of Table 2, respectively, summarize the answers. In general, one cannot endorse one

[7] A related concern is that the country may suffer an outbreak and thus an epidemic even without a positive test result. Because the punitive sanction is triggered even in this case, the sanction may appear unfair. This too may make it hard to follow through on the punitive sanction.

lever as better than the rest. The appropriate lever to push at any given moment will depend on factors such as expected mortality from an epidemic and the expected *ex post* sanctions for reporting, as well the pressure currently applied on other levers. Nevertheless, there are four lessons to keep in mind.

First, conditional surveillance subsidies only increase the total level of surveillance if the subsidy is greater than the amount of private investment in surveillance with no subsidy. That said, conditional subsidies always increases the level of publicly-available surveillance because surveillance conducted by a country in the absence of conditional subsidies is not reported to WHO. Second, if conditional surveillance subsidies are greater than private investment in surveillance without the subsidies and investments in surveillance are very productive (i.e., elasticity is greater than one), then these subsidies may reduce reporting. In that case, the cost for reporting countries (an increase in the frequency of positive test results and thus *ex post* sanctions) may be greater than the benefit to those countries (WHO medical assistance). Third, from the perspective of WHO, the productivity of punitive sanctions will depend on the cost ϕ to WHO from punitive sanctions. If the ROW primarily bears the cost of punitive sanctions, then, from the perspective of WHO, the productivity of this lever will be high.[8] Finally, a punitive sanction is only productive if the target country is wealthy enough to suffer under the sanction. Where the participation constraint binds, WHO will have to promise developmental aid that it can withdraw as punishment. The cost of this aid is greater than one because WHO must provide it even in cases where there is no outbreak.

4. Preemptive sanctions

It is well known that countries face sanctions once they report that they have suffered an outbreak. These sanctions may be official – imposed by policies of economic partners – or unofficial – imposed by domestic and foreign consumers of the country’s goods and services. Informal sanctions are just another way of stating that consumer demand for a country’s products and amenities are a function of how safe the latter are. Whether formal or informal, sanctions may limit a country’s trade,

[8] From the perspective of the ROW, punitive trade sanctions will have a negative efficiency cost, but a positive distributive effect. According to the conventional economic theory of international trade, the net cost including the country targeted by the sanctions is negative. But excluding costs imposed on the target country, the sign is ambiguous.

tourism, or inward foreign investment, i.e., they limit the mobility of goods, labor and capital.

What is less well known is that countries may face sanctions even if they do not report an outbreak. Such sanctions are based economic partners' and consumers' suspicions or fears of an outbreak even when one is not reported. Table 3 provides examples of such sanctions. In order to distinguish the two types of sanctions, we label sanctions that follow a report of an outbreak as "*ex post*" sanctions, and sanctions that are triggered merely by expectations of an outbreak as "preemptive" sanctions. Both economic partners and consumers recognize that countries have imperfect detection technology and imperfect incentives to report outbreaks. Therefore it is natural that they curtail interactions with countries based not just on reports of an outbreak but also on expectations of an outbreak. In this section we examine how sanctions triggered by the *probability* of an outbreak rather than a *report* of an outbreak affect the incentive to surveil and report outbreaks. Our main finding is that, unlike ordinary *ex post* sanctions, preemptive sanctions actually encourage investigation and disclosure.

We begin with the model as set forth in Section 2. There the rest of the world (ROW) imposed sanctions on a country only if it reported a positive test result. Now we assume that, even in the absence of a report, the ROW imposes sanctions S_0 with probability λ, where λ is the ROW's assessment of the probability of an outbreak given that the country does not report an outbreak:

$$\lambda(p_0) = \frac{\Pr\{outbreak\ and\ no\ report\}}{\Pr\{no\ report\}} = \frac{[1 - q(\theta(p_0))]p_0}{1 - q(\theta(p_0))p_0} \qquad (8)$$

The probability of a sanction depends on the ROW's belief about the probability of an outbreak p_0 and about how much the country invests in surveillance To keep things simple, we assume that the probability of an outbreak and any additional information (c, E, S_0) required to deduce a country's investment in surveillance is public knowledge. The preemptive sanction is only imposed in the states of the world without a report of an outbreak. The reason is that, when there is a report, there will be certain *ex post* sanctions and these make preemptive sanctions redundant.[9] Table 4 describes how preemptive sanctions modify the country's payoffs in each privately observed state of the world.

---

[9] In any case, it is easy to verify that the probability of an outbreak given a report is $1/q(\theta(p_0)) > 1$.

Reporting and surveillance. It is immediately apparent that preemptive sanctions increase reporting. Because a country might now be sanctioned even if it does not report an outbreak, the incremental cost of reporting falls. The same logic also encourages more surveillance. The marginal benefit of surveillance is an increase in the probability of detecting an outbreak times the net benefit of reporting. If the relative cost of reporting declines, the net benefit of reporting and thus marginal benefit of the surveillance rises. More formally, the reporting condition is now

$$S_0 \leq \frac{\Delta p_1(m, m+c)E}{1-\lambda(p_0)} \tag{9}$$

The right-hand side is clearly larger than in the reporting condition (1) for the basic model. The new surveillance condition is

$$q'(\theta) = \frac{1}{p_0\left[\Delta p_1(0, m+c)E - (1-\lambda(p_0))S_0\right]} \tag{10}$$

The right-hand side is smaller than in the analogous condition (3) for the basic model. The lesson is that, whereas *ex post* sanctions discourage reporting and surveillance, preemptive sanctions increase both activities.

Static policy levers. Preemptive sanctions also affect the productivity of WHO's various policy levers. These sanctions add to the power of medical assistance, punitive sanctions and conditional surveillance to encourage reporting in the sense that smaller amounts of these incentives are now required to induce any given level of reporting. These levers continue to encourage more surveillance, though at a slightly lower rate because preemptive sanctions are now doing some of the heavy lifting.[10]

Dynamic policy lever: WHO-as-auditor. WHO simultaneously plays the surveil-and-report game against many countries. These countries likely differ in the probability that they will suffer an outbreak. If countries have private information on these probabilities, WHO can piggy-back on preemptive sanctions to get all but the most recalcitrant countries to surveil and report outbreaks. The critical policy instrument is

---

[10] The only caveat to these findings is that, although preemptive sanctions increase the productivity of punitive sanctions, they also increase their cost to WHO. The reason is that, for countries whose participation constraint is not very slack, preemptive sanctions may lower wealth levels and make the participation constraint bind. This will increase the foreign aid that WHO must promise to ensure punitive sanctions are productive.

that WHO offers to audit a country's surveillance system. Specifically, WHO would review a country's surveillance system and verify a country's *negative* test results. WHO’s goal would be to make a country’s private information on its probability of an outbreak verifiable. When that happens, Grossman (1981) and Milgrom (1981), among others, have demonstrated that higher quality agents (here, low-risk countries) will voluntarily disclose their private information. As the principal (or, in our case, the rest of the world) updates it beliefs about non-disclosing agents, more agents will disclose. Ultimately, all but the worst quality agents will reveal their types.

The driving force behind this result in the disease outbreak context is that the ROW imposes preemptive sanctions on what they believe is the average probability of an outbreak across countries (or classes of countries). It is difficult for lower-than-average-risk countries to prove they are relatively less dangerous and reduce their risk of suffering preemptive sanctions. If they report only negative test results, the ROW may remain skeptical because negative results are easy to fabricate: have only known uninfected persons take the test. If they report positive test results, the ROW will impose full sanctions immediately.

To formally demonstrate how preemptive sanctions and WHO audits can facilitate the unraveling of non-disclosure, let countries be indexed by i and let p_{0i} be country i’s probability of an outbreak. We assume that there is heterogeneity in this probability across countries and that p_{0i} is private information to country i, but that the ROW knows the average probability of an outbreak:

$$\bar{p}_0 = \frac{\sum_i p_{0i}}{N} \tag{11}$$

among the N non-audited countries. To simplify, assume that probability of detection is one (so we can ignore the investment decision) and that WHO employs no other policy levers – not even medical assistance – to encourage reporting (so that no country is in a report equilibrium).

Because a non-audited country never reports, the ROW estimates the probability of an outbreak without a report as simply the probability of an outbreak: $\lambda = \bar{p}_0$. Also because it does not report, the non-audited country suffers preemptive sanctions of λS_0 whether or not it suffers an outbreak. If a country submits to an audit, it will suffer no preemptive sanctions when there is no outbreak because WHO vouches for negative test results and the probability of detection is assumed to be one. If there is a positive test result, however, the audit will disclose this to the ROW, triggering an *ex post* sanction of S_0.

A country prefers an audit only if the probability of *ex post* sanctions with an audit is lower than the incidence of preemptive sanctions without an audit: $p_{0i}S_0 \leq \lambda S_0$. But this condition is identical to $p_{0i} \leq \bar{p}_0$. In other words, all countries with a lower-than-average probability of outbreak will submit to an audit.

This has important dynamic effects. The ROW will revise upward its estimate of the probability of an outbreak among non-audited countries to be

$$\bar{p}_0' = \sum_{p_{0i} > \bar{p}_0} p_{0i} / N' > \bar{p}_0 \tag{12}$$

where N' is the number of countries with probabilities of outbreak above $\bar{p}_0$. Of the countries that did not submit to an audit in the first round, which will now submit to an audit? Those for whom the probability of *ex post* sanctions with an audit is less than the incidence of preemptive sanctions without an audit, or countries with $p_{0i} \leq \bar{p}_0'$. This second round of exit from the ranks of the non-audited will cause ROW to again revise upward its estimate of outbreak among non-audited countries. This process will continue until countries with the maximum p_{0i} remain. At that point an audit provides the ROW no new information because the ROW knows the precise probability of outbreak among non-auditors: max $\{p_{0i}\}$.

Introducing medical assistance only hastens this dynamic because it relaxes the auditing condition. Even some higher-than-average risk countries will now audit because an audit brings medical assistance after an outbreak and medical assistance can lower the changes of an epidemic. The relevant condition is that $p_{0i} \leq g(S_0, E, m, c)\, \bar{p}_0$ where

$$g(S_0, E, m, c) = \frac{S_0}{S - \Delta p_1(m, m+c)E} > 1 \tag{14}$$

At each iteration, a greater proportion of non-audited countries decide to be audited.

Introducing imperfect detection, by which we mean false negatives, has mixed effects. One the one hand, it may slow unraveling because even audited countries face preemptive sanctions. With imperfect detection, the ROW knows that audited countries may have an outbreak even though WHO tests find no outbreak. It will use preemptive sanctions to address this risk. The preemptive sanctions against audited countries will be lower than preemptive sanctions against non-audited countries because lower-risk countries tend to submit to audits and WHO audits are not completely uninformative. Nevertheless, preemptive sanctions against audited countries reduces the relative benefit of getting audited. On the other hand, countries that submit to an audit will surveil less in order to reduce the risk of a positive test result and thus *ex post* sanctions. Lower surveillance expenses makes getting audited more attractive. The net effect is such that only countries with probabilities of outbreak such that $p_{0i} \leq h(S_0, E, m, c, \bar{p}_0)$ get audited, where

$$h = \frac{\Delta\theta - \lambda^{**}S_0 + \bar{p}_0 S_0}{q^{**}\left[(1-\lambda^{**})S_0 - \Delta p_1(m, m+c)E\right] + (q^* - q^{**})\Delta p_1(0, m)E} \geq 0, \quad (15)$$

θ^{**} and θ^* are surveillance expenses with and without an audit, respectively, $\Delta\theta = \theta^* - \theta^{**} \geq 0$ is the decrease in surveillance expenses to avoid finding a outbreak, λ^{**} is the probability of a preemptive sanction even with a audit that finds no outbreak, $q^{**} = q(\theta^{**})$, and $q^* = q(\theta^*)$.

5. Sensitivity versus specificity

The public health community places a great deal of emphasis on the sensitivity of diagnostic testing for disease, i.e., on the probability of detecting disease in an infected patient. For example, both WHO Manual on Animal Influenza Detection and Surveillance (2002) and the Bush Administration's National Strategy for a Pandemic Influenza (2005) repeatedly stress sensitivity but never once mention specificity, or the probability of not detecting disease in an uninfected person, as an objective of surveillance. This focus on sensitivity makes a great deal of sense. One cannot stop an epidemic if one does not detect an outbreak. Increasing sensitivity and its corollary – reducing false negatives – ensure that the infected do no go without treatment and spread a contagion.

False negatives, however, have two sources. The obvious one is technological – the inability of a diagnostic test to identify an infected person. The less obvious source is behavioral – the failure of countries to surveil their populations and report infections to international authorities with the capacity to contain their spread. Ironically, an important cause of behavioral false negatives are diagnostic tests that lack specificity. Low specificity and its corollary – false positives – discourage surveillance and reporting because they increase the cost of these activities. False positives trigger sanctions but do not offer any benefits from medical assistance because there is no outbreak and thus epidemic to stop.[11] The lesson is that investments in diagnostic testing should not neglect the problem of technological specificity lest behavioral false negatives offset advances in technological sensitivity.

To illustrate our logic, let q be (as before) sensitivity, or the probability of a positive test result given an outbreak, and let r be specificity, or the probability of a negative test result given no outbreak or specificity. This implies that $(1 - q)$ and $(1 - r)$ are the probabilities of technological false negatives and false positives, respectively. Ignoring heterogeneity in the risk of outbreaks, Table 5 describes how imperfect testing modifies the joint probabilities of an outbreak and different test results. The main change from the basic model is that the lower left hand cell is no longer zero. Finally, let $\theta(q, r)$ be the monetary cost of investment in sensitive and specific surveillance. As is usual, we will assume this function is increasing and convex in its arguments.

Sensitivity and specificity alters the payoffs in the model with preemptive sanctions in two ways. First, the rest of the world (ROW) will reduce the frequency with which it imposes *ex post* sanctions because not all positive test results indicate an outbreak and sanctions hurt both the target and the sanctioning country. Specifically the probability of an *ex post* sanction falls to $(1 - \bar{\gamma})$ where $\bar{\gamma}$ is the ROW's assessment of the probability of no outbreak even though the country reports a positive test result. Note that this probability is based on the ROW's priors because the ROW does not observe (q, r). Second, the country alters its own assessment of the probability $(1 - \gamma(q, r))$ that it suffered an outbreak given a positive test result and the probability $\lambda(q, r)$ given a negative test result as follows:

[11] A natural question is whether repeated testing can overcome false positives. That will not always be the case. First, early in their development, diagnostic tests might employ indicators for multiple ailments, including the disease being targeted. This will imply strong positive correlation across test results for a patient who has one of those ailments, but not the disease. Second, early on in an outbreak, there may be a great deal of confusion. This confusion can lead to false positives that are not reversed for some time (Zamiska Oct. 18, 2005). Third, false positives may trigger sanctions before a second test is conducted (Canadien Press Sept. 9, 2003). Indeed, it may even trigger preemptive sanctions afterwards as the ROW might rationally believe there is a greater chance of an outbreak despite a second negative test result because there is a risk of a false negative with the second test.

$$1-\gamma(q,r)=\frac{\Pr\{outbreak, positive\ test\}}{\Pr\{positive\ test\}}=\frac{p_0 q}{p_0 q+(1-p_0)(1-r)} \quad (16)$$

$$\lambda(q,r)=\frac{\Pr\{outbreak, negative\ test\}}{\Pr\{negative\ test\}}=\frac{p_0(1-q)}{p_0(1-q)+(1-p_0)r} \quad (17)$$

The resulting payoffs are summarized in Table 6.

Reporting. Imperfect detection will have conflicting effects on reporting. On the one hand, because positive test results may be false, the value of medical assistance a country receives for reporting them falls. If there is no outbreak, there is no epidemic to prevent. On the other hand, false positives also cause the ROW not to sanction every report of a positive test result. This effect will reduce the cost of reporting. Formally, the reporting condition will be

$$S_0 \leq \frac{(1-\gamma(q,r))\Delta p_1(m,m+c)E}{1-\bar{\gamma}-\bar{\lambda}} \quad (18)$$

Ignoring preemptive and less-than-automatic *ex post* sanctions, only the first effect matters and reporting will fall relative to the basic case (see equation (1)). Relative to reporting under preemptive sanctions alone (see equation (9)), however, one cannot say whether reporting will rise or fall because it is unclear which effect is greater.

The more interesting finding is that sensitivity and, importantly, specificity increase the incentive to report. Sensitivity reduces the possibility that there is an outbreak that is not detected and therefore not controlled. Specificity rules out cases where there are sanctions because there is a positive test result but no benefits from medical assistance because there is no actual outbreak.

Surveillance. Before turning to surveillance, we should clarify who controls the rate of testing error. In reality, private firms and research institutions develop diagnostic tests. Their objectives, in turn, are set by a combination of WHO priorities and actual demand by target countries. If WHO highlights sensitivity, some countries will follow and tests that reduce false negatives will be supplied. Other countries may prefer specificity and tests that target false positives will be supplied. Early in the process the price of sensitivity and specificity will be high because research and development will still be required to develop tests and set up mass-production facilities. So when we assume countries must choose investments in sensitivity and specificity, we are really speaking of investments to generate demand for development and production of tests by organizations that may not even reside in those countries.

It is unclear whether imperfect detection raises or lowers the level of private investment in surveillance. The direction of the effect will depend on whether there are trade-offs between enhanced sensitivity and specificity. To see this formally, observe that the optimality conditions for sensitivity and specificity in the present model are

$$\theta_1(q, r) = p_0[\Delta p_1(0, m + c)E - \left(1 - \bar{\lambda} - \bar{\gamma}\right)S_0] \qquad (19)$$

$$\theta_2(q, r) = (1 - p_0)\left(1 - \bar{\lambda} - \bar{\gamma}\right)S_0 \qquad (20)$$

In order to compare these to prior models, we must change the control variable in prior models from the level of investment to the sensitivity q of testing and interpret $\theta(q)$ as the cost of sensitivity. The optimality condition under, e.g., the preemptive sanctions model is

$$\theta'(q) = p_0[\Delta p_1(0, m + c)E - \left(1 - \bar{\lambda}\right)S_0] \qquad (21)$$

Because the basic model, on which the model with preemptive sanctions is based, permitted false negatives, it is not surprising that the optimality condition under preemptive sanctions resembles that for sensitivity in the present model. What we are observing is actually the change in investment once we introduce specificity. The question is: does total investment rise or fall? The answer will depend on the form of $\theta(q, r)$. If there are no trade-offs between sensitivity and specificity, such as when $\theta(q, r) = \theta(q) + \theta(r)$, then total investment will rise. Condition (19) suggests that the investment in sensitivity is higher. Sensitivity is set at a higher level of marginal cost and θ is convex in q. Moreover, condition (20) implies positive investment in specificity. If there are trade-offs such that sensitivity comes at the expense of specificity or vice verse, i.e., $\theta_{12}(q, r) > 0$ for either technological or budgetary reasons, then it is possible that there will be less overall investment. The investment in specificity may cause investment in sensitivity to fall because the former increases the marginal cost of the latter, and the amount invested in specificity may not overcome that loss.

Rumor surveillance. The ROW does not solely rely on a formal announcement of a positive test result from a country to determine whether a country has an outbreak. It also relies on an informal channel that is called rumor surveillance. Rumor surveillance relies on rumor, local newspaper reports, public health online chat rooms, dissident reports, and indirect indicators such as the importation of large quantities of body bags or medicine to determine whether a country has suffered an outbreak (Samaan et al. 2005). WHO plays an important role in organizing rumor surveillance. Specifically it coordinates a private global surveillance network called the Global Outbreak Alert and Response Network (GOARN). GOARN has 120 member institutions (mainly universities and labs) around the world that track disease outbreaks. During the SARS crisis for example, although China did not report atypical pneumonia cases from mid-November 2002 until early-February 2003, GOARN members in the U.S. and Canada began to pick up local media reports of influenza cases in rural China as early as the end of November (Heymann and Rodier 2004).

One of the problems with rumor surveillance is that it is subject to a high rate of false positives. For example, in June 2005, a Chinese-language website, Boxum.com, maintained by Chinese dissidents claimed that 120 people died from H5N1 avian flu in Gangcha country in Qinghai province. It said that the Chinese government had blocked media and WHO access to the outbreak site and had quarantined 1,300 people who may have been infected with avian flu. The rumor turned out to be false, but not before it had spread well beyond Boxun and triggered a official denial from China. Ultimately, a WHO delegation to Qinghai was unable to confirm the deaths and the rumor turned out to be false (Harris 2006). Table 7 provides other examples of false rumors in recent years. In order to determine the veracity of rumors, the Western Pacific Regional Office of WHO assigned an officer to survey traditional and internet media sources as well as public health chat rooms for rumors. Every time the officer finds a rumor, WHO contacts a country's health authorities to verify the rumor. WHO has found that 31 of 40 (or 77%) of rumors identified between January 20 and February 26, 2004 were false (Samaan et al. 2005).

This subsection examines the effect of rumor surveillance, and in particular, error in rumor surveillance, on a country's incentive to surveil and report disease outbreaks. Our primary finding is that false positives in rumor surveillance have the unintended benefit of encouraging countries to disclose outbreaks through official channels. The reason is that false positive rumors raise preemptive sanctions in the absence of reporting. WHO can (and to a limited extent does) capitalize on these errors to encourage countries to submit to WHO audits. Disclosure through official channels is better for social welfare than disclosure through rumors for two reasons. First, only official disclosure brings with it international access to local specimens of an infectious agent. These samples are critical for finding a vaccine to combat contagion throughout the world. Second, disclosure through rumors does not give WHO the access it needs to treat local victims of a disease and control an outbreak. Only official disclosure grants WHO doctors the immigration rights required properly to deliver medical assistance.

Rumor surveillance modifies the basic model in two ways. First, it changes the strategy set of the country from {report, do not report} to {voluntarily report through official channels, involuntarily report via rumor}. Second, it increases the number of objective states of the world to four: positive or negative test result × positive or negative rumor. Each state will bring different probabilities of a sanction and outbreak. To keep matters simple, however, we shall assume that the country does not know whether rumor surveillance will report an outbreak or not before it must decide whether to report through official channels. Let the subscripts "*o*" and "*u*" on *q* and *r* indicate the sensitivity of official tests and rumors, respectively. Let subscripts "+" and "–" on γ and λ indicate the probabilities of an outbreak conditional on official test results and a positive rumor and negative rumor, respectively. Table 8 describes how rumor surveillance changes the joint probabilities of an outbreak and different official test results and rumors. Table 9 describes how rumor surveillance modifies payoffs in the game.

It should be obvious from the second and third rows of Table 9 that increasing the specificity r_u of rumors increases the payoff from not reporting by $(1 - p_0)\left(\bar{\lambda}_+ - \bar{\lambda}_-\right) > 0$ because it lowers preemptive sanctions. It also lowers the payoff to reporting by $(1 - p_0)\left(\bar{\gamma}_- - \bar{\gamma}_+\right) > 0$ because it reduces the relative downward adjustment the ROW makes to *ex post* sanctions to account for the possibility of a false positive in official test results.[12] Increasing the sensitivity of rumors has the opposite effect: it makes reporting more rewarding.

Increasing the specificity of rumors affects private investment in surveillance in interesting ways. If a country remains in a report-equilibrium, it lowers investment in sensitivity and increases investment in specificity. Increasing the sensitivity of rumors has the opposite effects. In effect, changes in the quality of rumors trigger analogous changes in the

[12] So long as rumors are informative, it will be the case that a positive rumor will increase the probability of a preemptive sanction, i.e., $\left(\bar{\lambda}_+ - \bar{\lambda}_-\right) > 0$, and a negative rumor will increase the downward adjustment of *ex post* sanctions to account for false positives, i.e., $\left(\bar{\gamma}_- - \bar{\gamma}_+\right) > 0$.

quality of official surveillance so as to reduce the costs of reporting. To see this, note that the optimality conditions for investment in sensitivity and specificity of official surveillance are similar to those without rumor surveillance,

$$\theta_1(q_o, r_o) = p_0[\Delta p_1(0, m + c)E - \left(1 - \bar{\lambda} - \bar{\gamma}\right)S_0] \tag{22}$$

$$\theta_2(q_o, r_o) = (1 - p_0)\left(1 - \bar{\lambda} - \bar{\gamma}\right)S_0 \tag{23}$$

except that now

$$\bar{\lambda} = \left[p_0 q_u + (1 - p_0)(1 - r_u)\right]\bar{\lambda}_+ + \left[p_0(1 - q_u) + (1 - p_0)r_u\right]\bar{\lambda}_- \tag{24}$$

$$\bar{\gamma} = \left[p_0 q_u + (1 - p_0)(1 - r_u)\right]\bar{\gamma}_+ + \left[p_0(1 - q_u) + (1 - p_0)r_u\right]\bar{\gamma}_- \tag{25}$$

These are the average probabilities of a preemptive sanction or of no outbreak taking into account the probabilities of positive and negative rumors. As before, increasing specificity (sensitivity) of rumors lowers (raises) the probability of preemptive sanctions and also the downward adjustment in *ex post* sanctions that the ROW makes to account for false positives. Therefore, increasing specificity (sensitivity) of rumors lowers (raises) the marginal cost at which investment in sensitivity is set and raises (lowers) the marginal cost at which investment in specificity is set. Because θ is increasing and convex, this implies lower (higher) investment in sensitivity and higher (lower) investment in specificity. Whether overall investment in official surveillance rises or falls depends, as before on the form of θ. If investments in sensitivity and specificity are independent, total investment will be unaffected.

6. Conclusion

The goal of this paper was to explore some complexities concerning the incentives countries have to disclose disease outbreaks. Our analysis, however, is still relatively simplistic. It introduces some twists because of imperfect information and gaming, but it does not account for other turns that may be found in real life. One example is that there may be serious constraints on the ability of WHO to audit countries' surveillance systems. Wide-ranging verification of test results requires a great deal of human resources and funding. The supply of these resources, like those for other international organizations, is limited by collective action problems among member states. Another example is that the rest of the world may be quick to impose sanctions on a target because it benefits politically-powerful domestic industries. This may increase the power of preemptive sanctions by reducing the relative cost of reporting. But it may also cause the rest of the world to increase the frequency of *ex post* sanctions despite false positives and those sanctions would increase the cost of reporting. Future work must explore these and other considerations to ensure that the public good value of disease surveillance and reporting is fully realized.

References

Arlen, Jennifer. 1994. "The Potentially Perverse Effects of Corporate Criminal Liability." J. Legal Studies 23(2): 833-867.

Blayney, Don P. Aug. 2005. "Disease-Related Trade Restrictions Shaped Animal Product Markets in 2004 and Stamp Imprints on 2005 Forecasts." Electronic Outlook Report from Economic Research Service, United States Department of Agriculture LDP-M-133-01.

Canadien Press. Sept. 9, 2003. "Health Canada working on SARS surveillance plan." Accessed at http://www.ctv.ca/servlet/ArticleNews/story/CTVNews/1063061919889_58471119/.

Science. July 18, 2003. "China's Missed Chance." 301: 294-296.

Harris, Shane. 2006. "The bug bloggers." Bulletin of Atomic Scientists. 62(1): 38-43.

Heymann, David L., and Guenael Rodier. 2004. "Global Surveillance, National Surveillance, and SARS." Emerging Infectious Diseases 10(2): 173-175.

Kahn, Joseph. 2003. "Hong Kong stirs, Beijing frets Across the border, Chinese ask: Could it happen here?" Intl. Herald Trib. July 23: 1.

Kaplow, Louis, and Steven Shavell. 1994. "Optimal Law Enforcement with Self-Reporting of Behavior." J. Political Economy 102(3): 583-606.

New York Times. Feb. 1, 2006. "Bird Flu Case in Hong Kong Raises Questions About China."

Panisset, U. 2000. International Health Statecraft: Foreign Policy and Public Health in Peru's Cholera Epidemic. Lanham: University Press of America.

Samaan, Gina, Mahomed Patel, Babatunde Olowokure, Maria C. Roces, Hitoshi Oshitani, and the World Health Organization Outbreak Response Team. 2005. "Rumor Surveillance and Avian Influenza H5N1." Emerging Infectious Diseases 11(3): 436-466.

Science. Jan. 20, 2006. "WHO Proposes Plan to Stop Pandemic in Its Tracks." 311: 315-16.

Science. Feb. 17, 2006. "Combating the Bird Flu Menace, Down on the Farm." 311: 944-946.

Science. Feb. 24, 2006. "New Study Casts Doubt on Plans for Pandemic Containment." 311: 1084.

TTG Asia. Jan. 27 - Feb. 6, 2006. "Tourism takes a hit." Accessed at http://www.ttgasia.com/index.php?option=com_content&task=view&id=12071&Itemid=48.

Voice of America News. Mar. 3, 2006. "UN Contractor: Countries Refuse to Cooperate in Bird Flu Investigation." Accessed at: http://www.voanews.com/english/2006-03-03-voa45.cfm."

Wall Street Journal. Nov. 11, 2003. "Beijing steps up bird-flu fight, enacts import ban." A10.

Wall Street Journal. Nov. 21, 2005. "Poultry From British Columbia Is Banned in U.S. Due to Bird Flu." A5.

Wall Street Journal. Dec. 23, 2005. "WHO Official Urges China To Share Bird-Flu Samples." Accessed at http://online.wsj.com/article_print/SB113534754713630527.html.

Wall Street Journal. Dec. 27, 2005. "Gaps seen in bird-flu alerts." A12.

World Health Organization. May 20, 2003. "Severe acute respiratory syndrome (SARS): Status of the outbreak and lessons for the immediate future." http://www.who.int/csr/sars/resources/en/index.html.

World Health Organization. July 4, 2003. "Update 95 - SARS: Chronology of a serial killer." Accessed at http://www.who.int/csr/don/2003_07_04/en/.

World Health Organization. May 23, 2005. "World Health Assembly adopts new International Health Regulations." Accessed at http://www.who.int/mediacentre/news/releases/2005/pr_wha03/en/index.html.

World Health Organization. Oct. 14, 2005. "Ten things you need to know about pandemic influenza." Accessed at: http://www.who.int/csr/disease/influenza/pandemic10things/en/index.html.

World Health Organization. Mar. 3, 2006. "H5N1 avian influenza: timeline." Accessed at: http://www.who.int/entity/csr/disease/avian_influenza/Timeline_3_03.pdf.

Zamiska, Nicholas. Oct. 18, 2005. "Inside UN Agency, Flu Data Sparked a Tense Debate." Wall Street Journal.

Appendix

Derivation of Table 2. The first row is derived by taking the derivative of (1) and (2), respectively, with respect to *c*. The first result in the second row is derived by taking the derivative of (3) with respect to θ_w, taking into account the moral hazard result that each dollar increase in θ_w is offset by a dollar reduction in θ so long as $\theta_w \leq \theta^*$, where θ^* is private investment in the absence of the subsidy. The second result in the second row follows from this moral hazard. The third row is derived by taking the derivative of (4) and (5), respectively, with respect to S_p.

Derivation of (15). Let ** and * indicate the audit and no audit states. Because, the audit option does not affect the decision of countries who will report even without an audit, we shall focus on a country that will not report in the absence of the audit option. By its decision not to report, we know that

$$(1 - \lambda)S_0 > \Delta p_1(m, m + c) \qquad \text{(A1)}$$

whether or not the country ultimately submits to an audit. The audit and no audit payoffs of the country are

$$\begin{aligned} &-\theta^{**} - p_{0i}q(\theta^{**})S_0 - (1 - p_{0i}q(\theta^{**}))\lambda^{**}S \\ &- p_{0i}q(\theta^{**})p_1(m + c)E - (1 - q(\theta^{**}))p_{0i}p_1(0)E \end{aligned} \qquad \text{(A2)}$$

$$-\theta^* - \lambda^* S_0 - p_{0i}q(\theta^*)p_1(m)E - (1 - q(\theta^*))p_{0i}p_1(0)E \qquad \text{(A3)}$$

respectively. A country will submit to an audit if (A2) ≥ (A3). Using the result $\Delta p_1(0, m + c) = \Delta p_1(0, m) + \Delta p_1(m, m + c)$, the audit condition can be re-written

$$p_{0i} \le \frac{(\theta^* - \theta^{**}) + (\lambda^* - \lambda^{**})S_0}{q(\theta^{**})[(1-\lambda^{**})S_0 - \Delta p_1(m, m+c)E] + \Delta q \Delta p_1(0, m)E} \quad \text{(A4)}$$

where $\Delta q = q(\theta^*) - q(\theta^{**})$. If a country neither reports nor is audited, the rest of the world (ROW) assumes the probability of an outbreak is the average across all countries, i.e., $\lambda^* = \bar{p}_0$. Substituting this into (A4) yield the condition $p_{0i} \ge h(S_0, E, m, c, \bar{p}_0)$ where h is defined by (15). What remains is to show that h (or equivalently, the right hand side of (A4)) is non-negative.

The first term in the numerator and the last term in the denominator are non-negative. The optimality conditions for investment with and without an audit are

$$\begin{aligned} q'(\theta^{**}) &= \{p_{0i}[\Delta p_1(0, m + c)E - (1 - \lambda^{**})S_0]\}^{-1} \\ &= \{p_{0i}[\Delta p_1(0, m)E + \Delta p_1(m, m + c)E - (1 - \lambda^{**})S_0]\}^{-1} \end{aligned} \quad \text{(A5)}$$

$$q'(\theta^*) = \{p_{0i}[\Delta p_1(0, m)E]\}^{-1} \quad \text{(A6)}$$

respectively. Given the no-report condition (A1), it is evident that the term in the square brackets in (A5) is smaller than the term in the term in the square brackets in (A6). Given that q is increasing and concave, this implies $\theta^* - \theta^{**} \ge 0$ and $\Delta q \ge 0$.

The second term in the numerator is also non-negative. The probability of preemptive sanctions with an audit is

$$\lambda^{**} = \Pr(\text{outbreak} \mid \text{no report}) = \frac{(1 - E[q^{**}])\bar{p}_0}{1 - E[q^{**}]\bar{p}_0 \; {}_0} \quad \text{(A7)}$$

where $E[q^{**}]$ is the ROW's expectation about the level of surveillance by the country. It is easily verified that (A7) $\le \bar{p}_0 = \lambda^*$. This may be surprising given that, for the reasons in the last paragraph, $E[q^{**}] \le E[q^*]$. The two results can be reconciled by the fact that, although the non-audited country does more surveillance, it does not report any test results to the ROW. Therefore, there is more disclosed surveillance with an audit.

Finally, the first term in the denominator is positive because of the no-reporting condition (A1).

Tables and Figures

Table 1. Payoffs in basic model.

Test result	Report positive result?	Payoff
Positive ($p_0 q(\theta)$)	No	$-\theta - p_1(m)E$
	Yes	$-\theta - S_0 - p_1(m + c)E$
Negative ($1 - p_0 q(\theta)$)	No	$-\theta - \dfrac{(1-q(\theta))p_O}{1-q(\theta)p_O} p_1(0)E$

Note. Probability of each test result are listed in its respective cell in column 1.

Table 2. Productivity of different policy levers.

	Marginal productivity for:	
Lever	Reporting	Total surveillance
c	$p_1'(c)\, E$	$q'(\theta)\, p_0\, p_1'(m+c)\, E / H > 0$, where $H < 0$ is the second order condition
Conditional θ_w	If $\theta_w \le \theta^*$ where θ^* is private investment without subsidy: $[p_0\, q(\theta^{**} + \theta_w)]^{-1} > 0$	If $\theta_w \le \theta^*$ where θ^* is private investment without subsidy: 0
S_p	If participation constraint is not binding: $p_1(0)/\phi$, where ϕ is marginal cost of S_p to WHO	$-q'(\theta)\, p_0\, p_1(0) / H > 0$, where $H < 0$ is the second order condition

Table 3. Examples of preemptive sanctions.

Date	Disease/Location	Sanction
Sanctions imposed before animal outbreaks		
2005	HPAI/not specific	Vietnam bans imports of poultry from 16 countries
2006	HPAI/France	Poultry consumption fell 20% in France before HPAI discovered
2006	HPAI/Bulgaria	Poultry sales fell 60% in Bulgaria before HPAI discovered in swans
Sanctions imposed after animal outbreak, before human outbreak		
1997	HPAI/Hong Kong	Hong Kong kills 1.5 mil. chickens
2001	FMD/UK	Cost UK tourism and beef industries £3 billion even before human casualties
2003	AI/ US	US poultry exports may have fallen 3% (Blayney 2005)
2003	BSE/US	US beef exports fell 80% without a single human infection (Blayney 2005)
2003-2005	HPAI/SE Asia	SE Asian economies loss $12 bil. in output (Thailand alone $1 bil., Vietnam up to 1.8% GDP; outside SE Asia, poultry prices up 20%, volume down 8%)
2005	HPAI/not specific	US ban poultry imports from all countries reporting animal outbreaks (WSJ Nov. 21, 2005)
2006	HPAI/Italy	Poultry consumption fell 70% in Italy

Notes. HPAI = highly pathogenic avian influenza, AI = avian influenza (not highly pathogenic), FMD = foot and mouth disease, BSE = mad cow disease.

Table 4. Payoffs with preemptive sanctions.

Test result	Report positive result?	Payoff
Positive ($p_0 q(\theta)$)	No	$-\theta - \lambda S_0 - p_1(m)E$
	Yes	$-\theta - S_0 - p_1(m + c)E$
Negative ($1 - p_0 q(\theta)$)	No	$-\theta - \lambda S_0 - \frac{(1-q(\theta))p_O}{1-q(\theta)p_O} p_1(0)E$

Note. Probability of each test result are listed in its respective cell in column 1.

Table 5. Probability of an outbreak conditional on test results.

	Positive test	Negative test
Outbreak	p_0q	$p_0(1 - q)$
No outbreak	$(1 - p_0)(1 - r)$	$(1 - p_0)r$

Table 6. Payoffs with false positives and false negatives.

Test result	Report	Payoff
Positive ($p_0q + (1 - p_0)(1 - r)$)	No	$-\theta(q, r) - \bar{\lambda} S_0 - (1 - \gamma(q, r))p_1(m)E$
	Yes	$-\theta(q, r) - (1 - \bar{\gamma})S_0 - (1 - \gamma(q, r))p_1(m + c)E$
Negative ($p_0(1 - q) + (1 - p_0)r$)	No	$-\theta(q, r) - \bar{\lambda} S_0 - \lambda(q, r)p_1(0)E$

Note. Probability of each test result are listed in its respective cell in column 1.

Table 7. Examples of false rumors.

Rumor (source, date)	Public health response
14-year old boy, tracked to Guangdong, China, died in Hong Kong (Wenhui Newspaper, 8/2/04)	Hong Kong investigation reveals report in incorrect
Four pigs tested positive for bird flu, Vietnam (Reuters Health Online, 2/6/04)	WHO and Food and Agricultural Organization (FAO) investigation finds report is incorrect
Bird flu in German tourist returning from Asia (Washington Times, 1/22/04)	WHO verifies that report is incorrect and issues a press release that there is no need to shift to higher state of readiness (Influenza Pandemic Plan Phase 1)
48 children with respiratory illness, Nam Dinh Province, Vietnam (WHO Network, 8/2/04)	WHO verified the report was incorrect.

Note. Table is reproduced from table in Samaan et al. (2005).

Table 8. Probability of an outbreak conditional on official test results and rumors.

	Positive test, positive rumor	Positive test, negative rumor	Negative test, positive rumor	Negative test, negative rumor
Outbreak	$p_0 q_o q_u$	$p_0 q_o$	$p_0(1 - q_o)q_u$	$p_0(1 - q_o)$
No outbreak	$(1 - p_0) \times (1 - r_o)(1 - r_u)$	$(1 - p_0) \times (1 - r_o)r_u$	$(1 - p_0) \times r_o(1 - r_u)$	$(1 - p_0) \times r_o r_u$

Table 9. Payoffs with rumor surveillance.

Test result	Report	Payoff
Positive ($p_0 q_o + (1 - p_0)(1 - r_o)$)	No	$-\theta(q_o, r_o) - \left(\left[p_0 q_u + (1-p_0)(1-r_u)\right]\bar{\lambda}_+ + \left[p_0(1-q_u) + (1-p_0)r_u\right]\bar{\lambda}_-\right)S_0 - (1 - \gamma(q_o, r_o))p_1(m)E$
	Yes	$-\theta(q_o, r_o) - \left(1 - \left[p_0 q_u + (1-p_0)(1-r_u)\right]\bar{\gamma}_+ + \left[p_0(1-q_u) + (1-p_0)r_u\right]\bar{\gamma}_-\right)S_0 - (1 - \gamma(q_o, r_o))p_1(m + c)E$
Negative ($p_0(1 - q_o) + (1 - p_0)r_o$)	No	$-\theta(q_o, r_o) - \left(\left[p_0 q_u + (1-p_0)(1-r_u)\right]\bar{\lambda}_+ + \left[p_0(1-q_u) + (1-p_0)r_u\right]\bar{\lambda}_-\right)S_0 - \lambda(q_o, r_o)p_1(0)E$

Note. Probability of each test result are listed in its respective cell in column 1.