PENN INSTITUTE for ECONOMIC RESEARCH
UNIVERSITY of PENNSYLVANIA

The Ronald O. Perelman Center for
Political Science and Economics (PCPSE)
133 South 36th Street
Philadelphia, PA 19104-6297

pier@econ.upenn.edu
http://economics.sas.upenn.edu/pier

PIER Working Paper
18-006

Go Big or Go Home: A Free and Perfectly Safe but Only Partially Effective Vaccine Can Make Everyone Worse Off

EDUARD TALAMAS
University of Pennsylvania
Department of Economics

RAKESH VOHRA
University of Pennsylvania
Department of Economics

January 15, 2018

https://ssrn.com/abstract=3102610

Go Big or Go Home:
A Free and Perfectly Safe but Only Partially Effective Vaccine Can Make Everyone Worse Off*

Eduard Talamàs & Rakesh Vohra†

Abstract

Vaccines are crucial to curb infectious-disease epidemics. Indeed, one of the highest priorities of the National Institutes of Health (NIH) on the HIV front is the development and delivery of a vaccine that is at least moderately effective. However, *risk compensation* could undermine the ability of partially-effective vaccines to curb epidemics: Since vaccines reduce the cost of risky interactions, vaccinated agents may optimally choose to engage in more of them and, as a result, may increase everyone's infection probability. We show that—in contrast to the prediction of standard models—things can be worse than that: *A free and perfectly safe but only partially effective vaccine can reduce everyone's welfare*. The reason is simple: By reducing the cost of risky interactions, a partially-effective vaccine can destabilize the existing interaction structure in favor of a less efficient one. Because of the strategic complementarities in risky interactions that we show arise when agents strategically choose their partners, the most efficient stable interaction structure after the introduction of a partially-effective vaccine can be much denser and—due to the negative externalities of risky interactions—worse for everyone. The result of this paper underscores the importance of taking into account the effects that different interventions have on social structure, and it suggests that the NIH might want to go big—i.e. deliver a highly-effective vaccine—or go home.

Date Printed: March 13, 2018

†Department of Economics & Department of Electrical and Systems Engineering, University of Pennsylvania. We thank Benjamin Golub for his feedback, encouragement and advice at many stages of this project. We also thank Edward Glaeser and Pau Milán for helpful discussions. This work was supported by the Warren Center for Network & Data Sciences, and the Rockefeller Foundation (#2017PRE301). The views expressed in this article do not reflect the views of the University of Pennsylvania, the Warren Center or the Rockefeller Foundation. All errors are our own.

1 Introduction

Infectious-disease epidemics like HIV are a major source of human suffering. According to the World Health Organization, about 35 million people have died from HIV, and roughly the same number are currently living with this virus.[1] The development of effective vaccines is crucial for preventing infectious-disease epidemics. Developing an HIV vaccine, for instance, is a high priority for the US National Institutes of Health. Anthony S. Fauci, the Director of the National Institute of Allergy and Infectious Diseases, recently observed[2]

> The development and delivery of a preventive HIV vaccine that is safe and *at least moderately effective* would help bring about a durable end to the HIV/AIDS pandemic. We are committed to pursuing multiple vaccine development strategies to achieve this goal.

In this paper we show that a free and perfectly safe but only partially effective vaccine can make everyone worse off. A partially-effective vaccine has two opposing effects on welfare. On the one hand, it allows agents to have more risky interactions, making them better off. On the other hand, it can increase the probability that agents become infected (because of the increase in risky interactions), making them worse off. We show that—in contrast to the prediction of existing economic epidemiological models—the second effect can dominate the first.

A key force in the mechanism is that there are strategic complementarities in risky interactions.[3] The reason is simple; we illustrate it here with an example. Suppose that there are two pairs of agents having risky interactions to start with: Ann and Bob are one pair, and Chloe and Dane the other (Network 1 in Figure 1). Each individual has a fixed probability of contracting a given virus independently of her interactions, and an infected individual transmits the virus in any given interaction with probability p. Infection and transmission are independent across agents and interactions, respectively.

To build intuition, consider first the extreme case in which each interaction transmits the virus with probability one—that is, $p = 1$. An interaction between Chloe and Bob is risky for each of them, since under some states of the world only one of them is infected, and

[1] See "Global Health Observatory (GHO) data" here.

[2] See "NIH and partners launch HIV vaccine efficacy study" here. Emphasis added.

[3] Friedman et al. (1987), Abdul-Quader et al. (1990) and Tross et al. (1992) provide evidence that is consistent with strategic complementarities in risky interactions: They document how—in the context of unprotected sex and needle sharing—partners' risk-reductions efforts are correlated with own risk-reduction efforts.

hence an interaction would infect the other. We claim that *an interaction between Ann and Dane* (a switch from Network 1 to Network 2 in Figure 1) *increases Chloe and Bob's incentives to interact*. Indeed, in Network 2, the states of the world where one catches the virus are the same as the states of the world where the other one catches it, so their interaction is risk free. A similar intuition holds when $p < 1$: Ann and Dane's interaction decreases the probability that only one of Chloe and Bob is infected, hence increasing their incentives to interact.

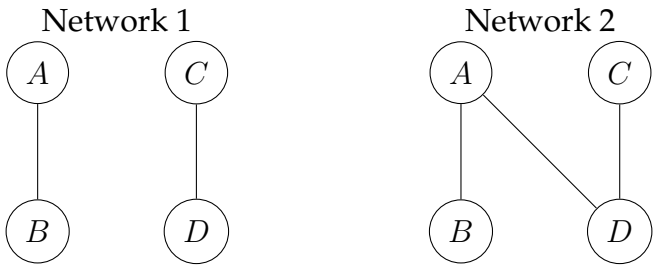

Figure 1: Two illustrative interaction networks.

This paper illustrates as simply as possible the mechanism by which a free and perfectly safe but only partially effective vaccine can make everyone worse off. Intuitively, *fixing the network of social interactions*, the introduction of a partially effective vaccine reduces everyone's probability of becoming infected, and hence makes everyone better off. However, by reducing the (ceteris-paribus) cost of each risky interaction, such an intervention can *destabilize the efficient network structure*. Because of the strategic complementarities in risky interactions just described, the next-best stable network structure can feature substantially more interactions and—as a consequence of the negative externalities that each interaction imposes on others via an increased infection probability—be worse for everyone. In other words, a relatively high transmission probability can play a beneficial role by preventing deviations from the efficient social structure. As a result, the beneficial effects of a partially-effective vaccine—in terms of decreased infection probability given any social structure—must be traded off against the welfare effects of the change in social structure that it unleashes.

Many social scientists have long realized that social networks play a central role in epidemiological processes (see for example Jacquez et al. 1988, Barnard 1993 and Friedman et al. 2006). Standard economic epidemiological models, however, abstract away from the structure of social interactions, so they are unable to capture the mechanism that we illustrate in this paper (see for example Kremer 1996 and Fenichel et al. 2011). Indeed, a free and perfectly safe but only partially effective vaccine necessarily makes everyone better off in

these models. The logic is simple; Kremer (1996, page 555) explains it as follows:[4]

> Adoption of an imperfectly effective vaccine could not cause the number of partners to increase so much that [the per-interaction probability of infection] increased, because people would not be willing to have more partners if the probability of infection from an additional partner increased.

Hence, in these canonical models, everyone is better off after the adoption of a free and perfectly safe imperfect vaccine. Indeed, since such a vaccine decreases the per-interaction probability of infection, everyone can choose the same amount of interaction as she was choosing before its introduction, and in this way obtain the same benefits from her interactions with a reduced probability of infection. From this perspective, the contribution of this paper is to show the existence of non-trivial tradeoffs in the distribution of free and perfectly safe but only partially effective vaccines: When agents strategically choose whom to interact with, there are strategic complementarities in risky interactions, which implies that the introduction of a perfectly safe and free but only partially-effective vaccine can make everyone worse off.

The result of this paper suggests that taking into account agents' strategic choice of partners is important in order to understand the potential effects that different interventions have on social structure—and hence on behavior and welfare. Moreover, it suggests that measuring the relevant interaction structure—and how it changes with different interventions—can be crucial for understanding which social groups are more likely to feature strategic complementarities in risky interactions, and hence which parts of a society are more vulnerable to the potentially-negative welfare effects of partially-effective vaccines and similar interventions.

The remainder of this paper is organized as follows. In section 2 we introduce the simple model that we use to illustrate our argument, and in section 3 we discuss how strategic complementarities in risky interactions naturally arise in this model. In section 4 we characterize the set of pairwise-stable networks in this simple model. In section 5 we present the main result of this paper: A free and perfectly safe but only partially-effective vaccine can make everyone worse off. We discuss the contribution of this paper in the context of the re-

[4] In this quote, we have substituted the symbol βY with its corresponding words: The per-interaction probability of infection. The sentence that follows the one in this quote is: "However, the combined costs of the increased prevalence, plus the expense and side effects of the vaccine, could outweigh the benefits of a reduced risk of infection per partner and so introduction of an imperfect vaccine could make everybody worse off." In this paper we show that an imperfect vaccine can reduce everyone's welfare *even if it is free and has no side effects.*

lated literature in section 6, and we conclude in section 7. Appendix A derives the infection probabilities that we use to prove some of the statements in the main body of the paper.

2 Simple Epidemiological Model

There are four agents (two men and two women) and four stages, listed below.

Stage 1: **Network Formation.** Each agent simultaneously announces which partners he or she wants to have. An edge between two agents is formed if and only if both of them have announced that they want to partner with the other.

Stage 2: **Infection.** Each agent becomes exogenously infected with probability q. Infection is independent across agents.

Stage 3: **Contagion.** Each edge becomes *live* with probability p. Each agent connected via a path of live edges to an infected agent becomes infected. Edges become live independently of each other.

Stage 4: **Utility Realized.** The utility of each agent is the benefit that he or she derives from his or her partners[5] (0 if no opposite-sex partners, s_1 if one opposite-sex partner, and s_1+s_2 if two opposite-sex partners) less the cost of infection (c if infected, and 0 otherwise).

Note 2.1. This model is similar to the one in Blume et al. 2011: The main difference is that Blume et al. 2011 assume that infected agents do not benefit from their links, whereas we assume that infected agents benefit from their links but pay a cost c when they become infected. More importantly, their objective is different: Whereas we focus on the effects of partially-effective vaccines—which we think of as reductions in the probabilities q and p—they focus on characterizing the structural differences between optimal and stable networks.

Stage 1 is the only stage in which agents take actions. We focus on situations in which having a risky interaction involves mutual consent. To capture this idea, we assume that the outcome in stage 1 is a *pairwise-stable* network. This solution concept—first proposed by Jackson and Wolinsky (1996)—is a natural refinement of Nash equilibrium in the network formation game (stage 1). Informally, a network is (pairwise) *stable* if no agent has an incentive to drop an existing edge, and no two agents have an incentive to form a new edge. To define it formally, let E be the edge set of a network and denote by $u_i(E)$ the utility that

[5]For simplicity, agents derive no benefit from same-sex partners.

agent i enjoys at edge set E. The network with edge set E is said to be (pairwise) stable if the following two conditions hold:[6]

1. For all $ij \in E$, $u_i(E) \geq u_i(E \setminus ij)$ and $u_j(E) \geq u_j(E \setminus ij)$.
2. For all $ij \notin E$, if $u_i(E \cup ij) > u_i(E)$, then, $u_j(E \cup ij) < u_j(E)$.

Note 2.2. We base our analysis on pairwise stability because it is both a well understood solution concept and natural for the applications that we focus on in this paper. Our analysis goes through if—in addition to pairwise stability—we require that no agent has incentives to sever any subset of her existing links.

For simplicity, we focus throughout on the case in which the utility s_1 of the first edge is high enough so that no network with an isolated agent is pairwise stable. Figure 2 depicts all the possible networks (up to isomorphism) that can emerge in the network formation stage. Let μ_i denote the probability that agent i becomes infected (exogenously—i.e. in stage 2—or endogenously—i.e. in stage 3); for simplicity we denote by μ_I and μ_X the infection probability of any given agent in the symmetric networks I and X, respectively. Appendix A describes the probability that the agent in each relevant network position is infected.

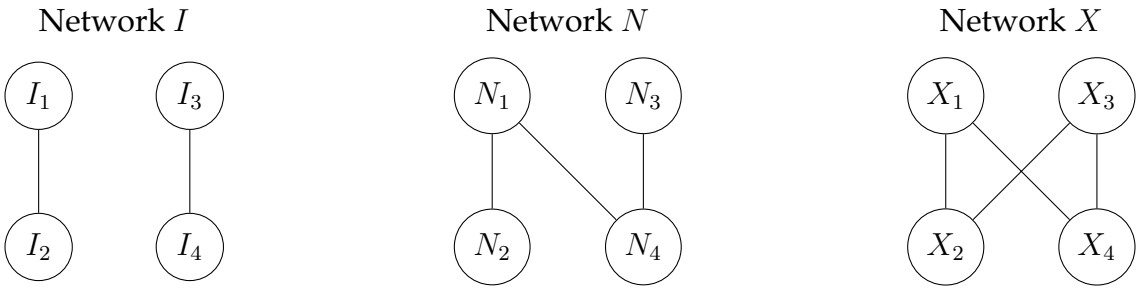

Figure 2: The Three Relevant Network Structures.

3 Strategic Complementarities in Risky Interactions

In this section we show how strategic complementarities in risky interactions naturally arise in the model described in section 2. Proposition 3.1 formalizes this idea using Definition 3.1.

Definition 3.1. Given a network G, the *risk of the edge ij for agent i* is the difference in i's infection probability in $G \cup ij$ and i's infection probability in G. When the risk of edge ij is the same for agents i and j, we refer to it simply by the *risk of the edge ij*.

[6]For brevity, we denote the edge between nodes i and j by ij.

Proposition 3.1. *The risk of the edge N_2N_3 in Network N is smaller than the risk of the edge I_2I_3 in Network I.*

Proof. The risk of edge N_2N_3 is $\mu_X - \mu_{N_2}$, and the risk of edge I_2I_3 is $\mu_{N_1} - \mu_I$. Using the expressions derived in Appendix A, it is easily verified that $\mu_X - \mu_{N_2} \geq \mu_{N_1} - \mu_I$ for all values of p and q. □

Note 3.1. Figure 3 depicts the risk of edge N_2N_3 and I_2I_3 as a function of the transmission probability p when the exogenous infection probability is $q = \frac{1}{4}$; the picture looks similar for all $q \in (0, 1)$. The risk of edge N_2N_3 is increasing for low values of p and decreasing for high values of p. Intuitively, the risk of edge N_2N_3 is highest when the transmission probability is high enough so that this edge has a substantial probability of transmitting an infection but low enough so that there is a substantial probability that only one of agents N_2 and N_3 are infected.

4 Stable Networks

Proposition 4.1 shows that requiring that the outcome in the network formation stage be stable reduces the candidate networks to I and X. This observation substantially simplifies the analysis, since these two networks are fully symmetric.

Proposition 4.1. *Network N is unstable for all p.*

Proof. Suppose for contradiction that network N is stable. This implies that the cost of the diagonal edge for N_2 is not greater than its benefit. By Proposition 3.1, the cost of adding the edge N_2N_3 for N_2 and N_3 is smaller than the cost of adding N_1N_4 for N_1 and N_4, while its benefit is exactly the same, so both N_1 and N_4 have incentives to remove the edge N_1N_4, a contradiction. □

Proposition 4.2 shows that network I is stable for intermediate values of the transmission probability p, which is intuitive: When the transmission probability p is small enough, network I is not stable because agents have incentives to form the diagonal links. In contrast, when the transmission probability p is high enough, network I is not stable because agents have incentives to remove their one link.

Proposition 4.2. *There exist p^*, p^{**} such that network I is stable if and only if $p \in [p^*, p^{**}]$.*

Proof. Network I is stable if and only if (i) no agent wants to delete her existing edge (that is, the cost s_1 of deleting this edge is greater than the associated benefit $c\mu_I$) and (ii) no two agents have incentives to partner up (that is, the cost $c(\mu_{N_1} - \mu_I)$ of an extra edge is greater than its benefit s_2). Using Equation 1 in Appendix A, condition (i) is easily verified to hold for all p small enough. Using Equation 3 in Appendix A, condition (ii) is easily verified to hold for all p large enough. □

Note 4.1. Figure 3 illustrates the determinants of the cutoff p^* when the infection probability is $q = \frac{1}{4}$ and $\frac{s_2}{c} = .11$. For simplicity, in Figure 3, we don't show the determinants of the cutoff p^{**}; this cutoff is 1 if s_1 large enough.

Proposition 4.3 shows that network X is stable with the only potential exception of a range of intermediate values of transmission probability, which is intuitive: The benefit from removing the edge X_2X_3 for its adjacent vertices is highest for intermediate values of p, when it is most likely that only one of them is infected in network N.

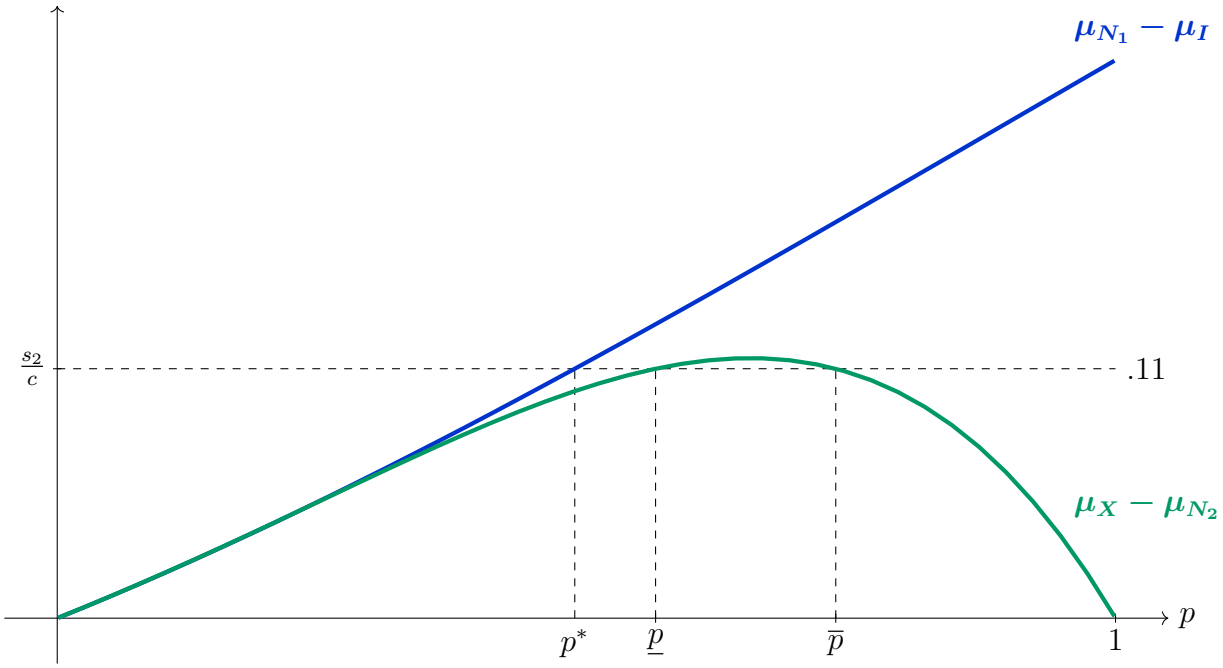

Figure 3: Illustration of Proposition 4.2 and Proposition 4.3 when $q = \frac{1}{4}$ and $\frac{s_2}{c} = .11$. Assuming s_1 is large enough, network I is stable if and only if $\frac{s_2}{c}$ is below $\mu_{N_1} - \mu_I$. Network X is stable if and only if $\frac{s_2}{c}$ is above $\mu_X - \mu_{N_2}$.

Proposition 4.3. *There exist $p^* < \underline{p} \leq \overline{p}$ such that network X is stable for all $p \notin (\underline{p}, \overline{p})$.*

Proof. Network X is stable if and only the benefit $c(\mu_X - \mu_{N_2})$ from deleting an edge is smaller than its cost s_2. Using Equation 4 in Appendix A, it is easily verified that this is satisfied for

all p except possibly those in an intermediate range $(\underline{p}, \overline{p})$. The fact that $p^* < \underline{p}$ follows from Proposition 3.1. □

Note 4.2. Figure 3 illustrates the determinants of the cutoff values $\underline{p}$ and $\overline{p}$ when the infection probability is $q = \frac{1}{4}$ and $\frac{s_2}{c} = .11$. For large enough values of $\frac{s_2}{c}$, network X is stable for all transmission probabilities p.

5 Partially-Effective Vaccines Can Make Everyone Worse Off

Network I and network X—the only two potentially stable networks—are fully symmetric, which implies that each agent's expected utility is the same in every stable network. Therefore, the welfare in each network scales with the expected utility of a single agent in this network.

From Proposition 4.2 and Proposition 4.3, we have that there always exists a nonempty region $(p^*, \underline{p})$ of values of the transmission probability p in which (i) both networks I and X are stable and (ii) a reduction in p leads to only network X being stable. Theorem 5.1 follows directly from this observation and the fact that for all values of p close enough to p^*, welfare in network I is greater than in network X.[7]

Theorem 5.1. *There exists $\Delta > 0$ s.t. each agent's expected utility in the most efficient stable network when $p \in (p^*, p^* + \Delta)$ is greater than when $p \in (p^* - \Delta, p^*)$.*

Figure 4 illustrates Theorem 5.1 for a particular utility function ($s_1 = 40$, $c = 80$ and $\frac{s_2}{c} = .11$) and exogenous infection probability $q = \frac{1}{4}$. In this case, p^* is approximately $.48$.

Note 5.1. Theorem 5.1 implies that there is always a range of transmission probabilities and a threshold $\Delta > 0$ such that—assuming that we start from the most efficient stable network (network I)—an intervention that reduces p by more than Δ but less than 2Δ necessarily harms everyone. A similar statement holds for interventions that reduce both p and q.

6 Relation to Existing Literature

The well-known phenomenon of *risk compensation* is an important element of the mechanism that we describe in this paper. Observed at least as early as the Victorian era (see for

[7] To see this last fact, note that, when $p = p^*$, agent 1's expected utility in network I is the same as that in network N, and hence, agent 1's expected utility in network I is strictly greater than that in network X.

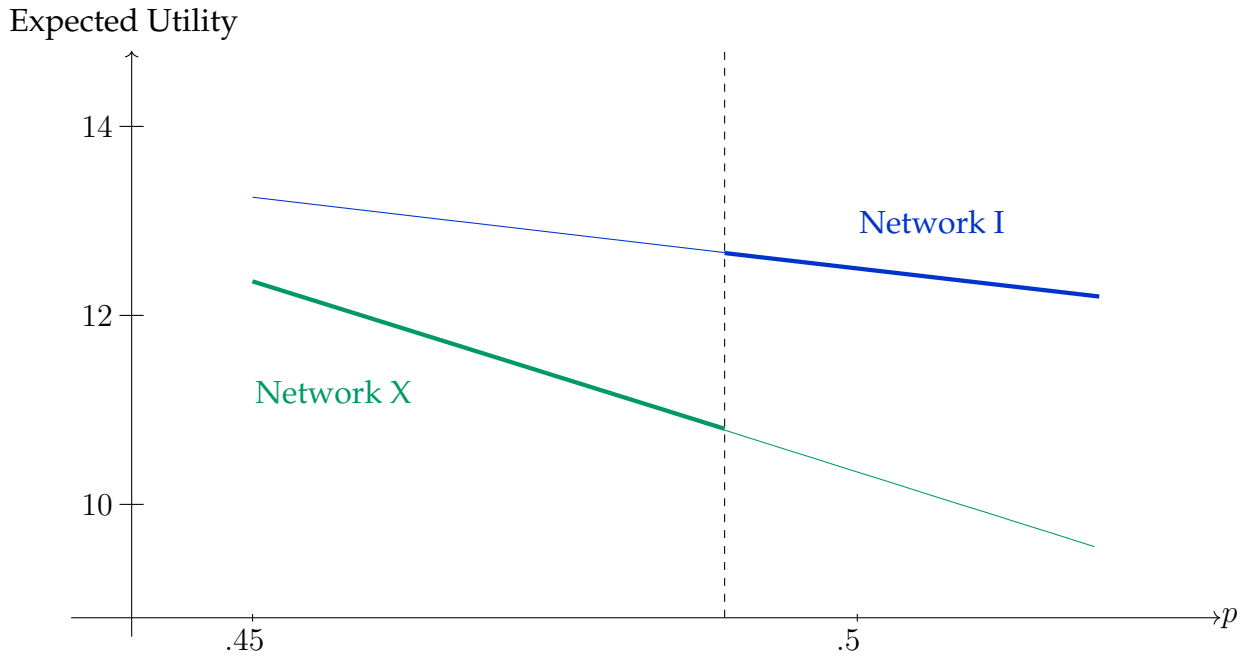

Figure 4: Illustration of Theorem 5.1 when $s_1 = 40, c = 80, \frac{s_2}{c} = .11$ and $q = 1/4$. Each agent's expected utility under the most efficient stable network structure is in bold.

example Adams 1879), it was popularized by Peltzman 1975, who controversially suggested that automobile safety regulations would not diminish automobile-related deaths. In the context of HIV, the evidence on risk compensation is mixed. For example, on the one hand, Eaton and Kalichman 2007 (see also Blumenthal and Haubrich 2017) review the empirical literature on risk compensation in HIV prevention and conclude that "risk compensation is evident in response to prevention technologies that are used in advance of HIV exposure and at minimal personal cost." On the other hand, Marcus et al. 2013 argue that there is no evidence of risk compensation in a recent trial of Daily Oral HIV Preexposure Prophylaxis (iPrEx).

The main contribution of this paper is to show that—as a result of risk compensation—a free and perfectly safe but only partially effective vaccine can make everyone worse off, which suggests that a non-trivial welfare trade-off must be considered when deciding whether or not to distribute partially-effective vaccines: While such an intervention increases the welfare associated with any given interaction structure, it can disrupt the existing interaction structure in favor of a more inefficient one. As already discussed in section 1, this contrasts with standard economic epidemiological models, which predict that free and perfectly safe vaccines—no matter how ineffective—necessarily make everyone better off.

Interestingly, using a dynamic version of a standard economic epidemiological model,

Toxvaerd 2017 argues that partially-effective vaccines can have negative welfare consequences in the transition between steady states. In contrast, we show—using a different model that allows agents to strategically choose whom to interact with—that the conclusion that a free and perfectly safe but only partially-effective vaccine necessarily makes everyone better off in steady state is an artifact of the anonymous-mixing assumption of the standard models.

The mechanism that we illustrate in this paper is related to—but distinct from—the one described in Kremer 1996, which can be summarized as follows: If low-activity people reduce their activity by a higher proportion than high-activity people in response to an increase in the prevalence of the disease, the composition of the pool of available partners worsens after such a change, which creates positive feedbacks. In stark contrast with our mechanism, however, the feedback effects in Kremer 1996 only make partially-effective vaccines more desirable. Indeed, in that model, the introduction of a vaccine reduces the marginal probability of infection for low-activity people more than for high-activity people.[8] This force is absent in our analysis because—in order to illustrate our mechanism as simply as possible—we focus on the case of homogeneous preferences.

The main ingredients of the mechanism that we illustrate in this paper are that risky interactions feature (i) strategic complementarities and (ii) negative externalities. Hoy and Polborn (2015) elegantly show how the combination of these two forces can imply that a safety technology improvement is welfare reducing. From this perspective, the contribution of this paper is to illustrate how strategic complementarities and negative externalities naturally arise in models of strategic risky interactions, and that a safety-technology improvement—a partially-effective vaccine in our application—can indeed reduce welfare in these models.

This paper complements the growing body of literature that studies the effects of different interventions on epidemiological processes (see for example Galeotti and Rogers 2013, Chen and Toxvaerd 2014, Rowthorn and Toxvaerd 2015, Goyal and Vigier 2015 and Goyal et al. 2016). The main difference between this paper and most of this literature is that we focus on the *welfare effects* of such interventions—rather than the effects on *infection rates*.

This paper is not the first to study epidemiological processes using the network formation model of Jackson and Wolinsky 1996. For example, Blume et al. 2011 use this approach to provide asymptotically tight bounds on the welfare of both optimal and stable networks. We use this simple and natural model to illustrate a simple mechanism that has important policy implications. While the particular model that we work with is useful to make our

[8] For those with sufficiently many partners, the introduction of a vaccine will actually increase the marginal risk of infection from an additional partner, reducing their optimal number of partners, and hence making the pool of available partners safer.

argument simply and precisely, the argument itself is general, so it should also apply to more general epidemiological models.

7 Conclusion

The capacity of infectious-disease epidemics to disrupt societies is comparable to that of wars and natural disasters. For this reason, considerable resources are expended to manage and ameliorate the effects of such epidemics. Because of risk compensation, however, the effects of different potential interventions are subtle. As a consequence, before deciding whether and how to intervene, we might wish to ensure that our interventions at least do no harm.

We show that—in contrast to what standard models predict—this fundamental principle is not necessarily satisfied by an intervention that consists of distributing a free and perfectly safe but only partially effective vaccine. In fact, we show that such an intervention can harm everyone. The reason is simple: Everything else equal, such a vaccine reduces the cost (in terms of infection probability) of having risky interactions, and hence it can destabilize the existing interaction structure in favor of a more inefficient one. We show how strategic complementarities—which arise naturally once we allow agents to strategically choose their partners—can generate feedback effects and, as a consequence, the next-best interaction structure can be much denser. Because of the negative externalities of risky interactions, this can make everyone worse off.

The result of this paper suggests that, on the HIV front, the National Institutes of Health might want to go big—e.g. deliver a highly effective vaccine—or go home. More generally, it underscores the importance of taking into account the network of social interactions in theoretical and empirical epidemiological studies: Changes in the structure of social interactions can have first-order effects on welfare, so understanding the forces that shape this structure is crucial for our understanding of infectious-disease epidemics and how to minimize their negative effects on human welfare.

A Appendix: Infection Probabilities

Lemma A.1 describes the probability that an agent becomes infected (exogenously—i.e. in stage 2—or endogenously—i.e. in stage 3) conditional on her network position.

Lemma A.1. *The probability that any given agent in network I is infected is*

$$\mu_I = qp + (1 - qp)q. \tag{1}$$

the probability that any given agent in network X is infected is

$$\begin{aligned} \mu_X \;=\;& (1-p)^2\mu_I \\ &+\; p(1-p)\left[q(2-q) + (1-q)^2 qp(2-qp)\right] \\ &+\; p(1-p)\left[q + (1-q)p(pq + (1-pq)q(2-q))\right] \\ &+\; p^2\left[q(2-q) + (1-q(2-q))q(2-q)p(2-p)\right], \end{aligned} \tag{2}$$

the probability that N_1 is infected is

$$\mu_{N_1} = \mu_I + (1-q)(1-pq)\mu_I p \tag{3}$$

and the probability that N_2 is infected is

$$\mu_{N_2} = q + (1-q)pq + (1-q)^2 p^2 \mu_I \tag{4}$$

Proof. To see Equation 1, consider for concreteness the probability that I_1 is infected. The probability that I_2 infects I_1 is qp and, conditional on not being infected by I_2, I_1 is infected with probability q.

To derive Equation 2, consider the three exhaustive and mutually exclusive cases depicted below, where thick edges correspond to live edges. We say that i *is infected from* j if j is exogenously infected and there is an live path from i to j.

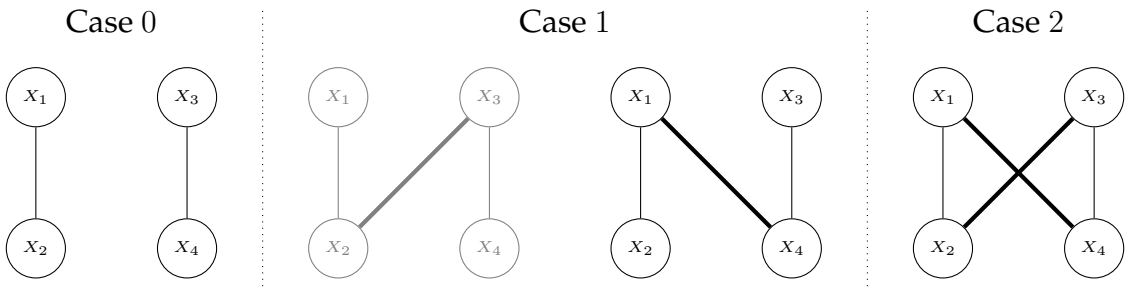

Case 0: *None of the edges X_1X_4 and X_2X_3 are live.* This happens with probability $(1-p)^2$. The probability that any given agent is infected is μ_I.

Case 1: *Exactly one of the edges X_1X_4 and X_2X_3 is live.* This happens with probability $2p(1-p)$. Assume without loss of generality that X_1X_4 is live (and hence X_2X_3 is not live). The probability that node X_1 is infected is $q(2-q)+(1-q(2-q))(qp+(1-qp)qp)$ or

$$q(2-q)+(1-q)^2qp(2-qp) \tag{5}$$

To see this, note that the probability that X_1 is infected from X_1 or X_4 is $1-(1-q)^2=q(2-q)$, and the probability that X_1 is infected from X_2 or X_3 is $qp+(1-qp)qp$.

The probability that node X_2 is infected is

$$q+(1-q)p(pq+(1-pq)q(2-q)) \tag{6}$$

To see this, note that the probability that X_2 is exogenously infected is q. Conditional on this not happening, the probability that X_2 is infected is p times the probability that X_1 is infected from X_3, or X_1 or X_4, which is $pq+(1-pq)q(2-q)$.

Hence, each agent's expected probability of infection in this case is the average of expressions (5) and (6).

Case 2: *Both edges X_1, X_4 and X_2X_3 are live.* This happens with probability p^2. The probability that X_1 is infected is

$$q(2-q)+(1-q(2-q))q(2-q)p(2-p).$$

To see this, note that the probability that X_1 is infected from X_1 or X_4 is $1-(1-q)^2=q(2-q)$, and the probability that X_1 is infected from X_3 or X_4 is the probability $q(2-q)$ that either of them is infected times the probability $p(2-p)$ that at least one of the edges X_1X_2 and X_3X_4 is live.

To see Equation 3, note that $\mu_{N_1}-\mu_I=(1-q)(1-pq)\mu_I p$, since the probability that N_1 is infected from either N_3 or N_4 and is not infected from either N_1 or N_2 is the probability $1-q$ that N_1 is not infected from N_1 times the probability $1-qp$ that N_1 is not infected from N_2 times the probability μ_I that N_4 is infected from either N_3 or N_4 times the probability p that the edge N_1N_4 is live.

Finally, to see Equation 4, note that the probability that N_2 is infected is the probability q that she becomes exogenously infected plus the probability $(1-q)p$ that she does not become infected and N_1N_2 is live times the probability $q+(1-q)p\mu_I$ that N_1 is infected from N_1, N_3 or N_4. That is, $\mu_{N_2}=q+(1-q)p\left[q+(1-q)p\mu_I\right]$, which is equivalent to Equation 4. □

References

ABDUL-QUADER, A. S., S. TROSS, S. R. FRIEDMAN, A. C. KOUZI, AND D. C. DES JARLAIS (1990): "Street-recruited intravenous drug users and sexual risk reduction in New York City." *Aids*, 4, 1075–1080.

ADAMS, C. F. (1879): "Notes on railroads accidents," *G.P. Putnam Sons*.

BARNARD, M. A. (1993): "Needle sharing in context: patterns of sharing among men and women injectors and HIV risks," *Addiction*, 88, 805–812.

BLUME, L., D. EASLEY, J. KLEINBERG, R. KLEINBERG, AND É. TARDOS (2011): "Network formation in the presence of contagious risk," in *Proceedings of the 12th ACM conference on Electronic commerce*, ACM, 1–10.

BLUMENTHAL, J. AND R. HAUBRICH (2017): "Risk compensation in PrEP: An old debate emerges yet again," *The Virtual Mentor*, 909–915.

CHEN, F. AND F. TOXVAERD (2014): "The economics of vaccination," *Journal of Theoretical Biology*, 363, 105–117.

EATON, L. A. AND S. C. KALICHMAN (2007): "Risk compensation in HIV prevention: Implications for vaccines, microbicides, and other biomedical HIV prevention technologies," *Current HIV/AIDS*, 165–172.

FENICHEL, E. P., C. CASTILLO-CHAVEZ, M. G. CEDDIA, G. CHOWELL, P. A. G. PARRA, G. J. HICKLING, G. HOLLOWAY, R. HORAN, B. MORIN, C. PERRINGS, ET AL. (2011): "Adaptive human behavior in epidemiological models," *Proceedings of the National Academy of Sciences*, 108, 6306–6311.

FRIEDMAN, S. R., R. CURTIS, A. NEAIGUS, B. JOSE, AND D. C. DES JARLAIS (2006): *Social networks, drug injectors' lives, and HIV/AIDS*, Springer Science & Business Media.

FRIEDMAN, S. R., D. C. DES JARLAIS, J. L. SOTHERAN, J. GARBER, H. COHEN, AND D. SMITH (1987): "AIDS and self-organization among intravenous drug users," *International Journal of the Addictions*, 22, 201–219.

GALEOTTI, A. AND B. W. ROGERS (2013): "Strategic immunization and group structure," *American Economic Journal: Microeconomics*, 5, 1–32.

GOYAL, S., S. JABBARI, M. KEARNS, S. KHANNA, AND J. MORGENSTERN (2016): "Strategic network formation with attack and immunization," in *International Conference on Web and Internet Economics*, Springer, 429–443.

GOYAL, S. AND A. VIGIER (2015): "Interaction, protection and epidemics," *Journal of Public Economics*, 125, 64–69.

HOY, M. AND M. K. POLBORN (2015): "The value of technology improvements in games with externalities: A fresh look at offsetting behavior," *Journal of Public Economics*, 131, 12–20.

JACKSON, M. O. AND A. WOLINSKY (1996): "A strategic model of social and economic networks," *Journal of Economic Theory*, 71, 44–74.

JACQUEZ, J. A., C. P. SIMON, J. KOOPMAN, L. SATTENSPIEL, AND T. PERRY (1988): "Modeling and analyzing HIV transmission: the effect of contact patterns," *Mathematical Biosciences*, 92, 119–199.

KREMER, M. (1996): "Integrating behavioral choice into epidemiological models of AIDS," *The Quarterly Journal of Economics*, 111, 549–573.

MARCUS, J., D. GLIDDEN, K. MAYER, A. LIU, S. BUCHBINDER, K. AMICO, ET AL. (2013): "No evidence of sexual risk compensation in the iPrEx trial of daily oral HIV preexposure prophylaxis," *PLoS ONE*, 8(12).

PELTZMAN, S. (1975): "The effects of automobile safety regulation," *Journal of Political Economy*, 83, 677–725.

ROWTHORN, B. R. AND F. TOXVAERD (2015): "The optimal control of infectious diseases via prevention and treatment," *Mimeo*.

TOXVAERD, F. (2017): "Rational disinhibition and externalities in prevention," *Mimeo*.

TROSS, S., A. ABDUL-QUADER, H. SILVERT, AND D. DES JARLAIS (1992): "Condom use among male injecting-drug users—New York City 1987-1990." *Morbidity and Mortality Weekly Report*, 41, 617–20.