



Since January 2020 Elsevier has created a COVID-19 resource centre with free information in English and Mandarin on the novel coronavirus COVID-19. The COVID-19 resource centre is hosted on Elsevier Connect, the company's public news and information website.

Elsevier hereby grants permission to make all its COVID-19-related research that is available on the COVID-19 resource centre - including this research content - immediately available in PubMed Central and other publicly funded repositories, such as the WHO COVID database with rights for unrestricted research re-use and analyses in any form or by any means with acknowledgement of the original source. These permissions are granted for free by Elsevier for as long as the COVID-19 resource centre remains active.



Free and perfectly safe but only partially effective vaccines can harm everyone [☆]



Eduard Talamàs ^{a,*}, Rakesh Vohra ^b

^a IESE Business School, Spain

^b University of Pennsylvania, United States of America

ARTICLE INFO

Article history:

Received 28 February 2019

Available online 8 May 2020

JEL classification:

C72

D85

I18

Keywords:

Epidemics

Vaccines

Risk compensation

Social structure

ABSTRACT

Risk compensation can undermine the ability of *partially-effective* vaccines to curb epidemics: Vaccinated agents may optimally choose to engage in more risky interactions 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: Free and perfectly safe but only partially effective vaccines can *reduce everyone's welfare*, and hence fail to satisfy—in a strong sense—the fundamental principle of “first, do no harm.” Our main departure from standard economic epidemiological models is that we allow agents to strategically choose their partners, which we show creates strategic complementarities in risky interactions. As a result, the introduction of a partially-effective vaccine can lead to a much denser interaction structure—whose negative welfare effects overwhelm the beneficial direct welfare effects of this intervention.

© 2020 Elsevier Inc. All rights reserved.

1. Introduction

Finding a safe and effective Covid-19 vaccine is an urgent priority. More than 50 different such vaccines are being developed, and billions of dollars are already being spent to produce the most promising ones (e.g., Aten [2020](#)). However, in the spirit of the wise motto “hope for the best, plan for the worst,” we have to plan for the possibility that none of these vaccines will be perfectly effective. In this case, we might be tempted to produce and distribute the best safe vaccine that we could find. Would this necessarily be a good idea?

In this paper, we show that a free and perfectly safe but only partially-effective vaccine can *reduce everyone's welfare*, and hence fail to satisfy—in a strong sense—the fundamental principle of “first, do no harm.” In particular, we uncover a novel mechanism that suggests that delivering a vaccine whose effectiveness is below a certain threshold is not necessarily a good idea, and it underscores the importance of developing models that help practitioners estimate this threshold in particular applications.

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 (as a

[☆] We thank Tim Roughgarden (the editor), the advisory editor and two reviewers for useful comments, as well as Simon Board, Craig Garthwaite, Edward Glaeser, Benjamin Golub, George Mailath and Pau Milán—as well as various audiences—for helpful feedback. This work was supported by the Warren Center for Network & Data Sciences, and The Rockefeller Foundation (#2017PRE301). All errors are our own.

* Corresponding author.

E-mail address: etalamas@iese.edu (E. Talamàs).



Fig. 1. Two illustrative partnership networks.

consequence of the increase in risky interactions), making them worse off. We show that—in contrast to the prediction of standard models—the second effect can uniformly dominate the first.

Our key departure from the standard economic epidemiological models (e.g., Kremer 1996 and Fenichel et al. 2011) is that we allow agents to strategically choose their partners—instead of only allowing each agent to choose her *number* of partners, and then having matches occur uniformly at random. This departure is crucial because it uncovers the existence—even in low-risk settings—of strategic complementarities in consensual risky partnerships, which are central to the mechanism that we uncover in this paper.

Let us illustrate the intuition behind the existence of these strategic complementarities with the simplest example. Suppose that there are two pairs of agents having risky partnerships to begin with: Ann and Bob are one pair, and Chloé and Dane the other (Network 1 in Fig. 1). Each individual has a fixed probability of contracting a given virus independently of her partnerships, and an infected individual transmits the virus in any given partnership with probability p . Infection and transmission are independent across agents and partnerships, respectively. To build intuition, consider first the extreme case in which each partnership transmits the virus with probability one—that is, $p = 1$. The partnership between Chloé and Bob is risky for each of them, since under some states of the world only one of them is infected, and hence their partnership would infect the other. We claim that Ann and Dane's partnership (that is, a switch from Network 1 to Network 2 in Fig. 1) increases Chloé and Bob's incentives to become partners (as long as their initial partnerships are sufficiently valuable so that neither of them wants to become isolated instead). 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 partnership is risk free. Remarkably, the same is true for *all* positive values of the transmission probability p : Ann and Dane's partnership decreases the probability that only one of Chloé and Bob is infected, hence increasing their incentives to become partners.

The introduction of a partially-effective vaccine has obvious positive direct welfare effects: Fixing the partnership network, it reduces everyone's probability of becoming infected, and hence makes everyone better off. But the introduction of a partially-effective vaccine can also have subtle negative indirect welfare effects: By reducing the (*ceteris-paribus*) cost of each risky interaction, it can destabilize the existing partnership network. We show that—because of the strategic complementarities in risky interactions illustrated above—the best stable partnership network after this intervention can be much denser than before such an intervention, and—as a consequence of the negative externalities of risky interactions—be worse for everyone.

More generally, our analysis highlights how *relatively high infection transmission probabilities can play a beneficial role* by preventing deviations from the efficient social structure. As a result, the beneficial effects of partially-effective vaccines—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 they can unleash.

In the context of a general game with externalities, Hoy and Polborn (2015) show that the combination of strategic complementarities and negative externalities implies that a safety technology improvement can be welfare reducing. From this perspective, the contribution of this paper is to (i) show how strategic complementarities in consensual risky partnerships arise naturally when the agents strategically choose their partners, and (ii) illustrate how this can imply that a free and perfectly-safe but only partially-effective vaccine can reduce *everyone's* welfare.

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. Indeed, a free and perfectly safe but only partially effective vaccine necessarily makes everyone better off in these models. Kremer (1996, page 555) explains the logic as follows¹:

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.

¹ 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*.

Hence, in these canonical models, everyone is better off after the adoption of a free and perfectly safe but only partially-effective vaccine. Indeed, since such a vaccine decreases the per-interaction probability of infection, everyone can choose the same amount of interaction as before its introduction, and in this way obtain the same benefits from her interactions with a reduced probability of infection. In order to clarify this point, and to highlight the role that the strategic choice of partners plays in the mechanism that we illustrate, in Appendix B we describe a random-matching version of our model, and we show that, as long as the infection risk is not extreme, strategic complementarities in risky partnerships do not arise in this model. As a result, a partially-effective vaccine *increases* everyone's welfare in this random-matching version of our model, which shows that strategic complementarities are central to the mechanism that we uncover in this paper.

The remainder of this paper is organized as follows. In section 2, we introduce the model that we use to illustrate our argument and, in section 3, we discuss how strategic complementarities in risky partnerships naturally arise in this model. In section 4, we characterize the set of stable networks and, in section 5, we show how a free and perfectly safe but only partially-effective vaccine can harm everyone. After discussing extensions of our model in section 6, we further discuss the contribution of this paper in the context of the related literature in section 7, and we conclude in section 8. In Appendix A, we derive the infection probabilities that we use to prove some of the statements in the main body of the paper. Finally, in Appendix B, we describe a random-matching version of our model, and we show how, in this case, (i) no strategic complementarities arise, and hence (ii) a free and perfectly safe but only partially-effective vaccine necessarily makes everyone better off.

2. Epidemiological model with strategic choice of partners

In order to capture the central tradeoff between the benefit and costs of partnership formation in the simplest possible way, we follow the modeling approach of Blume et al. (2011), which provides a tractable static reduced-form model for the inherently dynamic process of contagion (see Remark 2.1 and Remark 2.2 below).

For simplicity, we start by focusing on a relatively simple case featuring $n \geq 2$ men and n women, where n is even. We also assume that agents value only (up to two) partners of the opposite sex. This case represents a sexually-transmitted disease better than coronavirus, but is useful to illustrate the mechanism as transparently as possible. In section 6, we discuss how the results extend to the case in which homogeneous agents value an arbitrary amount of partnerships, which better represents infectious diseases like coronavirus.

The game consists in the following four stages:

- Stage 1: **Network Formation.** Each agent 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 form a partnership with the other.
- Stage 2: **Infection.** Each agent becomes exogenously infected with probability $q > 0$. Infection is independent across agents.
- Stage 3: **Contagion.** Each edge becomes *live* with probability $p > 0$. 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 (0 if no opposite-sex partners, s_1 if one opposite-sex partner, and $s_1 + s_2$ if two opposite-sex partners, with $s_2 \leq s_1$) minus the cost of infection (c if infected, and 0 otherwise).

Stage 1 is the only stage in which the agents take actions. We focus on situations in which having a risky partnership involves mutual consent. To capture this idea, we assume that the outcome in stage 1 is a *stable network*—in the sense that no agent can profitably drop any subset of his or her edges, and no pair of agents can both benefit by adding an edge between them (while possibly removing some of their existing edges). We discuss alternative notions of stability in Remark 4.2 below, and we show that our results are robust to stronger notions of stability in section 6.

Remark 2.1. This model can be seen as a network-based variation of the well-known epidemiological model in Philipson and Posner (1993, page 33), and it is similar to the one in Blume et al. (2011); the main differences are that Blume et al. (2011) assume that all edges provide the same benefit and that infected agents do not benefit from their edges, whereas we assume decreasing returns to scale in edges and that infected agents benefit from their edges 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 p and/or q —they focus on characterizing the structural differences between optimal and (pairwise) stable networks.

Remark 2.2. The contagion stage of the model (stage 3) can be seen as a reduced form version of the following dynamic model: Each partnership consists of interactions that occur periodically over time. There are two types of partnerships, risky and safe. Each interaction in a risky partnership has a positive probability of involving risky behavior that transmits the infection (if one of its members is infected). In contrast, the interactions in safe partnerships do not involve such risky behavior. Each partnership has a probability p of being risky. Agents are not able to tell whether a partnership is risky or safe, so forming a partnership is always risky. In this model, the infection must stop spreading at some point; the agents that are infected in stage 3 can be thought of as all the agents that are infected once the contagion stops spreading.




Fig. 2. The networks I , N and X .

Notation 2.1 reviews standard notation that we shall use throughout the text.

Notation 2.1. Given a network $G = (N, E)$ with node set N and edge set E , and given $N' \subseteq N$ and $E' \subseteq E$, we say that (N', E') is a *subnetwork* of G . We say that a subnetwork $G' = (N', E')$ of G is a *component* of G if E contains no edges between N' and $N \setminus N'$, and, for any two nodes i and j in N' , there is a *walk* (sequence of edges) in E' from i to j .

3. Strategic complementarities in risky partnerships

We start by showing how strategic complementarities in risky partnerships naturally arise in this model. For simplicity, we focus on the case in which the value s_1 of having one partner is high enough so that no network with an isolated agent is stable.

Fig. 2 depicts all the possible networks (up to isomorphism) of four non-isolated agents 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 brevity, we denote by μ_I and μ_X the infection probability of any given agent in the symmetric networks I and X , respectively. In Appendix A we derive the infection probability of the agent in each relevant network position.

Proposition 3.1 uses Definition 3.1 to formalize the idea that risky partnerships are strategic complements. For brevity, we denote the edge between nodes i and j by ij .

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 the edge ij is the same for agents i and j , we refer to it simply as the *risk of the edge ij* .


Proposition 3.1. The risk of the edge N_2N_3 in Network N is strictly 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 $\mu_X - \mu_{N_2} < \mu_{N_1} - \mu_I$ for all values of p and q . \square

Fig. 4 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 $q \in (0, 1)$. To understand this picture, it is useful to keep in mind that the transmission probability p affects agents' optimal actions in two different ways. On the one hand, p is the probability that infection is transmitted between two prospective partners. On the other hand, p influences the probability that the agents under consideration have already become infected through contagion from other agents in their networks. This explains why the risk of edge N_2N_3 is increasing in p for low values of p and decreasing in p for high values of p : 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 N_2 and N_3 are infected.

Remark 3.2. Philipson and Posner (1993) discuss how increases in the prevalence of a disease reduce the demand for risky partnerships for low levels of prevalence, but can increase it for high enough levels of prevalence. The intuition there is that, if prevalence is sufficiently high, then agents may have such a high probability of being already infected that they might rationally reduce their precautions. In contrast, the non-monotonicity depicted in Fig. 4 is present for all values of infection probability q , and hence for *all levels of prevalence*. To understand the difference, note that, even when q is very small, so the prevalence of the disease is very small, the edge N_2N_3 is essentially riskless when the transmission p is high—not because the probability that N_2 and N_3 are already infected is high, but because the probability that N_2 is infected *conditional on N_3 being infected* is high.

Remark 3.3. The key behind Proposition 3.1 is the observation that the risk of a partnership only comes from the situations in which one and only one of the parties is infected. In other words, the correlation between the health states of two agents

Fig. 3. The networks M and M' .

created by overlapping partnerships is important to understand the incentives behind strategic partnership formation. Relatedly, Toxvaerd (2017) discusses how overlapping partnerships can affect the speed of agents' learning about transmission risks via the correlation in health states that they induce, and that the riskiness of an activity does not necessarily correspond to how much it exposes one to a disease, but rather to how much it exposes one to a disease while not already infected.

4. Stable networks

Proposition 4.1 shows that only the *pair-complete* and the *cross-complete* networks (see Definition 4.1 below) can be stable. Every pair-complete network is isomorphic to $n/2$ copies of Network I , and every cross complete network is isomorphic to $n/2$ copies of Network X .

Definition 4.1. We say that a network is *pair complete* if each agent is part of exactly one edge (that is, each agent has exactly one partner). We say that a network is *cross complete* if each agent is in a component of the network that is isomorphic to Network X (see Fig. 2).

Proposition 4.1 shows that pair-complete networks are stable for intermediate values of the transmission probability p , which is intuitive: When the transmission probability p is small enough, Network I is unstable 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 4.1 also shows that cross-complete networks are stable with the only potential exception of a range of intermediate values of transmission probability, which is also 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 exactly one of them is infected in network N .

Proposition 4.1. Only pair-complete and cross-complete networks are ever stable. Moreover:

1. There exist $p^* \leq p^{**}$ such that any pair-complete network is stable if and only if p is in $[p^*, p^{**}]$.
2. There exist $p^* < \underline{p} \leq \bar{p}$ such that any cross-complete network is stable if and only if p is not in (\underline{p}, \bar{p}) .

Proof. A pair-complete network is stable if and only if (i) no agent wants to remove her existing edge (that is, the cost s_1 of removing this edge is greater than the associated benefit $c\mu_1$) and (ii) no two agents have incentives to form a partnership—that is, the cost $c(\mu_{N_1} - \mu_1)$ of an extra edge is greater than its benefit s_2 , or

$$\mu_{N_1} - \mu_1 \geq \frac{s_2}{c}. \quad (1)$$

Using equation (4) in Appendix A, it is easily verified that there exists p^{**} such that condition (i) holds for all transmission probabilities $p < p^{**}$. Using equation (6) in Appendix A, it is easily verified that there exists $p^* < p^{**}$ such that condition (ii) holds for all $p > p^*$.

A cross-complete network is stable if and only if (i) the benefit $c(\mu_X - \mu_{N_2})$ from deleting an edge is smaller than its cost s_2 , or

$$\mu_X - \mu_{N_2} \leq \frac{s_2}{c}, \quad (2)$$

and (ii) no two agents have incentives to remove one of their edges and create an edge between them; that is, $\mu_{M_3} \geq \mu_X$ (see Fig. 3). Using equation (7) in Appendix A, it is easily verified that condition (i) is satisfied for all p except possibly those in an intermediate range (\underline{p}, \bar{p}) . The fact that $p^* < \underline{p}$ follows from Proposition 3.1. Condition (ii) follows from the fact that combining equation (5) and equation (8) in Appendix A gives $\mu_{M_3} > \mu_X$.




Fig. 4. Illustration of Proposition 4.1 and when $q = \frac{1}{4}$ and $\frac{s_2}{c} = .11$. Assuming s_1 is large enough, a pair-complete network is stable if and only if $\frac{s_2}{c}$ is below $\mu_{N_1} - \mu_I$. A cross-complete network is stable if and only if $\frac{s_2}{c}$ is above $\mu_X - \mu_{N_2}$.

It only remains to be shown that a network that is neither pair complete nor cross complete is not stable. The fact that $\mu_{M_3} > \mu_X$ implies that no subnetwork of a stable network is isomorphic to network M . Note that no agent is part of more than two edges in any stable network, so a component of such a network contains a node that is part of two edges if and only if all of the nodes in this component are part of two edges (otherwise, because n is even, there must be two agents who are each part of only one edge who can profitably deviate by removing their one edge and matching to each other). Hence, we only have to show that no subnetwork of a stable network is isomorphic to network M' (see Fig. 3). This follows from the fact that $\mu_X < \mu_{M'}$, where $\mu_{M'}$ denotes the infection probability of each agent in M' .² This last inequality is intuitive: The probability that each of the agents M'_1, M'_2, M'_3 and M'_4 becomes infected goes down when the three edges $M'_3M'_6, M'_6M'_5$, and $M'_5M'_2$ are replaced by the edge $M'_3M'_2$, since this leaves each of them with the same number of edges but decreases the independent sources of infection. \square

Fig. 4 illustrates the determinants of the cutoff p^* when the infection probability is $q = \frac{1}{4}$ and $\frac{s_2}{c} = .11$. For clarity, in Fig. 4, we don't show the determinants of the cutoff p^{**} ; this cutoff is 1 if s_1 is large enough. Fig. 4 also illustrates the determinants of the cutoff values \underline{p} and \bar{p} for the same parameters $q = \frac{1}{4}$ and $\frac{s_2}{c} = .11$. For large enough values of $\frac{s_2}{c}$, every cross-complete network is stable for all transmission probabilities p .

Remark 4.2. Our notion of stability is stronger than the standard *pairwise stability* notion of Jackson and Wolinsky (1996). In particular, while pairwise stability considers one link at a time, we allow pairs of agents to deviate by forming an edge between them while, at the same time, dropping some of their existing edges. We regard this notion of stability as being just as reasonable, and it simplifies our arguments by allowing us to deem networks like M and M' as unstable, leaving only the pair-complete and cross-complete networks as the potentially stable ones. We show how our results are robust to stronger notions of group stability in section 6.

5. Partially-effective vaccines can harm everyone

Pair-complete and cross-complete networks—the only two (non-trivial) potentially stable types of networks—are fully symmetric. Hence, in every stable network, all agents' expected utilities are the same. Therefore, the total welfare in each network scales with the expected utility of a single agent in this network.

From Proposition 4.1, we have that there always exists a nonempty region (p^*, \underline{p}) of values of the transmission probability p in which (i) both pair-complete and cross-complete networks are stable and (ii) a reduction in p leads to only

² Indeed, $\mu_X < \mu_{M'}$ implies that M'_2 and M'_3 strictly benefit from partnering while dropping their partnerships with M'_5 and M'_6 , respectively.

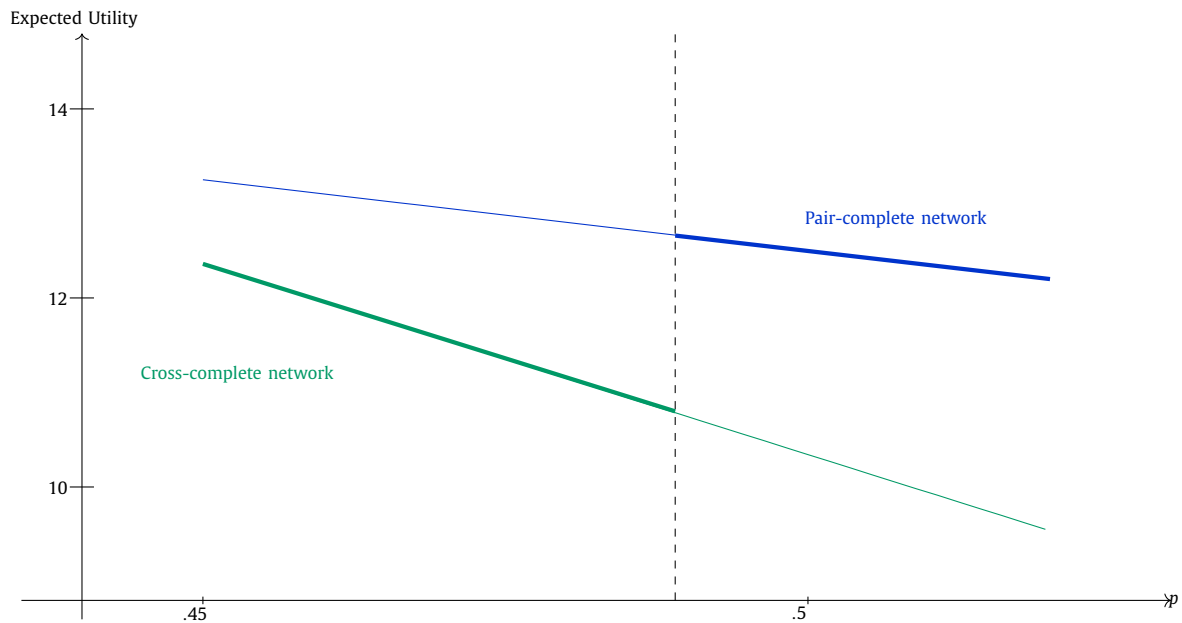


Fig. 5. Illustration of Theorem 5.1 when $s_1 = 40$, $c = 80$, $\frac{s_2}{c} = .11$ and $q = 1/4$. Each agent's expected utility under the most efficient stable network structure is in bold.

pair-complete networks being stable. Theorem 5.1 follows directly from this observation and the fact that, for all values of p close enough to p^* , everyone's expected utility in a pair-complete network is strictly greater than in a cross-complete network. 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, each agent's expected utility in a pair-complete network is strictly greater than that in a cross-complete network.

Theorem 5.1. *There exists $\Delta > 0$ such that each agent's expected utility in the most efficient stable network when $p \in (p^*, p^* + \Delta)$ is strictly greater than when $p \in (p^* - \Delta, p^*)$.*

Fig. 5 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.

Remark 5.2. Theorem 5.1 implies that there is always a non-empty range of transmission probabilities above p^* such that—assuming that we start with a transmission probability in this range and from an efficient stable network (a pair-complete network)—there is a non-empty range of reductions in the transmission probability p that necessarily harm everyone. A similar statement holds for interventions that either reduce q , or that reduce both p and q .

6. Discussion and extensions

We have shown how strategic complementarities in consensual risky partnerships naturally arise when agents strategically choose their partners. The intuition is that, when two agents become partners, they reduce the cost—in terms of infection risk—of partnerships among their partners. We have also shown how these strategic complementarities—combined with risk compensation—imply that an intervention that *reduces* the (ceteris-paribus) transmission probability of an infection can *reduce everyone's welfare*.

In order to make the analysis as transparent as possible, we have worked with a stylized model featuring a number of simplifying assumptions. But the mechanism that we have uncovered is a natural one, and it is thus robust to more general environments. For example, we have assumed that agents only value up to two partners, and that homosexual partnerships are not valuable. In this section, we show that the mechanism that we have illustrated in this paper does not rely on these assumptions. Also, so far we have focused on a stability notion that ensures that pairs that can form a beneficial partnership (while possibly dropping some of their other partnerships) do so. But this notion of stability does not ensure that larger groups of agents do not have incentives to deviate. In order to show that our results go through under stronger notions of group stability as well, we now focus on equilibria in which no group of agents (irrespective of its size) can arrange the partnerships among its members (while possibly dropping some partnerships with non-members) to make all of its members better off.




Fig. 6. The 3- and 4-complete networks when $n = 12$.

Consider the following modification of the game described in section 2, where agents are homogeneous (in particular, they are all of the same sex or, alternatively, they do not mind the sex of their partners) and value an arbitrary amount of partnerships. There are $n \geq 2$ agents. The expected utility of an agent with $m \in \{1, 2, \dots, n-1\}$ partnerships who has a probability μ of becoming infected is $\sum_{\ell=1}^m s_{\ell} - c\mu$, where s_{ℓ} denotes the benefit of the ℓ th-partnership, and the returns to partnerships are decreasing (i.e., $s_{\ell+1} \leq s_{\ell}$).

For the rest of this section, let $2 \leq m \leq n$ be such that n is divisible by both m and $m+1$. We say that a network is m -complete if it consists of n/m cliques of size m —that is, n/m components, each containing all possible links among its m members.³ Fig. 6 illustrates the 3-complete and the 4-complete network in the case $n = 12$. We restrict attention to situations in which the infection probability in the m -complete network is not extreme. In particular, we assume that the exogenous infection probability q is sufficiently small so that each agent's infection probability μ_m in the m -complete network is not larger than $1/2$.

Theorem 6.1 below is the analog of Theorem 5.1 in the more general framework described in this section. It establishes that a reduction in the transmission probability p can reduce everyone's welfare. For simplicity, we focus on *symmetric equilibria*—in which every agent has the same number of partnerships. To prove Theorem 6.1, we describe situations in which (i) only the m -complete and the $(m+1)$ -complete networks are stable, (ii) everyone's expected utility under the m -complete network is strictly higher than under the $(m+1)$ -complete network, and (iii) a reduction in the transmission probability p leaves the $(m+1)$ -complete network as the only stable network.

Theorem 6.1. *There exist $p_m > 0$, $\Delta > 0$ and $0 < \kappa_{m+1} < \kappa_m < \kappa_{m-1}$ such that, if $s_{m-1} > \kappa_{m-1}$, $s_m < \kappa_m$ and $s_{m+1} < \kappa_{m+1}$, then each agent's expected utility in the most efficient symmetric stable network when $p \in (p_m, p_m + \Delta)$ is strictly larger than when $p \in (p_m - \Delta, p_m)$.*

Proof. Let $s_{m-1} > \kappa_{m-1}$, where κ_{m-1} is sufficiently large so that networks in which agents have strictly less than $m-1$ partners are not stable. Also, let $s_{m+1} < \kappa_{m+1}$, where κ_{m+1} is sufficiently small so that networks in which agents have strictly more than m partners are not stable either. The only equilibrium network candidates are then the m -complete and the $(m+1)$ -complete networks. This is because—conditional on each agent having k partnerships—the probability of infection of a member of a clique is as low as it can be. As a result, in every network other than the m -complete network in which every agent has $m-1$ partnerships, a group of m agents can profitably deviate by dropping all of their edges with others and forming all possible links among themselves. Similarly, in every network other than the $(m+1)$ -complete network in which every agent has m partnerships, a group of $m+1$ agents can profitably deviate by dropping all of their edges with others and forming all possible links among themselves.

The m -complete network is stable if and only if an agent's cost (in terms of infection risk) of forming an additional edge is greater than its benefit. In other words, the m -complete network is stable if and only if⁴

$$s_m \leq (1 - \mu_m)\mu_m p c. \quad (3)$$

Since $\mu_m < 1/2$ is increasing in p , the right-hand side of (3) is also increasing in p . Hence, there exist unique values p_m and p_{m+1} that satisfy

$$p_m = \frac{s_m}{c(1 - \mu_m)\mu_m} \text{ and } p_{m+1} = \frac{s_{m+1}}{c(1 - \mu_{m+1})\mu_{m+1}},$$

³ For the purposes of illustration, we focus on the case in which n is divisible by both m and $m+1$, so that the m -complete network and the $(m+1)$ -complete network are well-defined. When this divisibility assumption is not satisfied, slight modifications of these networks (that include smaller "remainder components") can be constructed so that our main result still holds.

⁴ The increase in infection probability due to an added edge ab in the m -complete network is the probability $1 - \mu_m$ that a is not infected from another source times the probability μ_m that node b is infected times the probability p that the edge ab is active.




Fig. 7. The 3_{34} -network when $n = 12$.

respectively. These are the threshold values of the transmission probability p below which the m -complete and the $(m + 1)$ -complete network, respectively, are not stable. Given that $\mu_m < \mu_{m+1}$ and that $s_{m+1} \leq s_m$, we have that $p_{m+1} < p_m$. Hence, because the cost of an additional link $(1 - \mu_m)\mu_m pc$ increases with p , a reduction in p from p_m to a value between p_{m+1} and p_m renders the m -complete network unstable, and leaves the $(m + 1)$ -complete network as the only stable network.

We conclude by showing that there exists a threshold $\bar{p}_m > 0$ such that, if $p_m < \bar{p}_m$, then everyone's expected utility in the $(m + 1)$ -complete network when $p = p_m$ is strictly lower than in the m -complete network. Given that p_m is increasing in s_m , this implies that there exists κ_m such that, if $s_m < \kappa_m$, then $p_m \leq \bar{p}_m$, and hence the expected utility in the $(m + 1)$ -complete network when the transmission probability is at the threshold p_m at which the m -complete network becomes unstable is strictly lower than in the m -complete network.

Define the m_{ab} -network to be the m -complete network with an extra edge (between agents a and b). Fig. 7 illustrates the 3_{34} -network when $n = 12$. By definition of the threshold p_m , if the transmission probability p is equal to p_m , then everyone's expected utility in the m -complete network is the same as the expected utility of agents a and b in the m_{ab} -network. Since agent a has m edges in both the m_{ab} -network and the $(m + 1)$ -complete network, it is enough to show that, when p is small enough, her infection probability in the m_{ab} -network is smaller than in the $(m + 1)$ -complete network. For this, note that the number of different paths of length 2 that can transmit the infection to her in the $(m + 1)$ -complete network is larger than in the m_{ab} -network.⁵ Given that the number of different paths of length 1 that can transmit the infection to her in the $(m + 1)$ -complete network is the same as in the m_{ab} -network (namely, the number of neighbors she has, m), and that, when p is small enough, the probability that the network has a path of length strictly greater than 2 is second-order compared to the probability that the network has a path of length 2, we conclude that, when p is small enough, agent a 's probability of infection is larger in the $(m + 1)$ -complete network than in the m_{ab} -network. \square

An important part of the proof of Theorem 6.1 is to show that the increase in the density of the network triggered by the reduction in the transmission probability p makes everyone worse off. In the case that we have focused on throughout the main body of the paper (in which agents only value up to two heterosexual partnerships), this step is simpler because Network N is a strict subset of Network X , which directly implies that everyone's infection probabilities are strictly smaller in the former than in the latter. In the more general setting discussed in this section, the reasoning is a bit more subtle, because there are edges in the $(m + 1)$ -network that are not in the m_{ab} -network, and vice versa. As a result, it is not immediate to see that, when the transmission probability p is at the threshold p_m , the infection probability of agents a and b in the $(m + 1)$ -complete network is larger than in the m_{ab} network. As we have shown, however, this is true as long as the threshold transmission probability p_m is smaller than an upper bound $\bar{p}_m > 0$. As an illustration, the following table provides the approximate upper bounds on p_m in the case $m \in \{2, 3, 4, 5, 6\}$ and $q = 1/4$.

\bar{p}_2	\bar{p}_3	\bar{p}_4	\bar{p}_5	\bar{p}_6
.57	.51	.47	.43	.4

Hence, if $p_3 < .51$, for example, then a reduction in the transmission probability forces a society that starts with a 3-complete network and a transmission probability just above p_3 to switch to a different social structure (e.g., a 4-complete network), potentially making everyone worse off unless the reduction in the transmission probability is large enough to compensate for the reduction in welfare created by the induced change in social structure.

Remark 6.2. Beyond the specific modeling assumptions of this paper in terms of strategic network formation, the analysis in this section has shown that, as long as the threshold transmission probability p_m is below \bar{p}_m , when the transmission probability p is close to p_m , the $(m + 1)$ -complete network gives lower welfare to everyone than the m -complete network. A similar analysis holds when we define the m -complete and the $(m + 1)$ -complete networks to be networks with distinct components, with the components not necessarily being cliques, but with everyone having exactly $m - 1$ and m edges, respectively. Hence, the mechanism that we have illustrated in this paper is present in alternative models of network

⁵ A path of a network is a walk that contains any given node at most once. In the $(m + 1)$ -complete network, there are $m(m - 1)$ different paths of length 2 that end in node a . In the m_{ab} -network, there are $(m - 1)(m - 2) + m - 1 = (m - 1)^2$ different such paths.

formation in which the agents organize themselves into different components (not necessarily cliques), and in which no two agents can have incentives to form an additional edge (so that, when the transmission probability p moves below a threshold, the efficient network is no longer stable, while a denser network can still be stable).

7. Relation to the 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 example Adams 1879), it was popularized by Peltzman (1975), who controversially suggested that automobile safety regulations would not diminish automobile-related deaths. The importance of risk compensation in various settings has been a source of heated debates ever since. Most recently, Greenwood et al. (2019) calibrate a general equilibrium search model to quantitatively assess the extent to which behavioral reactions can reduce the effectiveness of several policy interventions.

In the context of HIV, the evidence on risk compensation is mixed.⁶ For example, on the one hand, Eaton and Kalichman (2007) (see also Chan et al. 2015, Delavande and Kohler 2015, and 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 mechanism that we illustrate in this paper is distinct from the one described in Kremer (1996), which relies on heterogeneities in agents’ preferences, and 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 a disease*, the composition of the pool of available partners worsens after such a change, which creates positive feedbacks. In contrast to our mechanism, however, the feedback effects in Kremer (1996) only make partially-effective vaccines more desirable. For example, for those with sufficiently many partners, the introduction of a partially-effective 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 for everyone. These feedback effects are absent in our analysis because—in order to illustrate our mechanism as simply as possible—we focus on the case of homogeneous preferences.

Pongou and Serrano (2013) provide a strategic model of network formation that shows how a higher optimal number of partners for men than for women can explain the fact that more women are infected from HIV/AIDS than men. For this, they describe two different dynamic processes of network formation whose long-run predictions refine the set of pairwise-stable networks. They assume that both the benefits and costs of each link are exogenous, so that each agent has a fixed optimal number of partners. In contrast, in the present setting, the costs of a link are endogenous—and this is the key for the strategic complementarities in risky interactions that we uncover.

This paper complements the growing body of theoretical literature that studies the effects of different interventions on epidemiological processes (e.g., 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—that is, the trade-off between changes in behavior and changes in infection rates—rather than just on the effects on *infection rates*. In a similar vein, Toxvaerd (2019) uses a dynamic version of a standard economic epidemiological model to show that reducing the infectiousness of a disease can—via negative welfare effects along the transition between steady states—reduce agents’ discounted lifetime welfare. However, as in related economic epidemiological models (e.g., Kremer 1996 and Fenichel et al. 2011), reductions in the infectiousness of a disease cannot reduce anyone’s steady state welfare. In particular, the steady state per-exposure probability of infection in Toxvaerd (2019) is independent of the infectiousness of the disease. Hence, each agent can—by choosing her exposure level exactly as before the change in infectiousness—be exactly as well off as before. 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.

8. 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.

In this paper, we show that an intervention that consists of distributing a free and perfectly safe but only partially effective vaccine can fail this fundamental principle of “first, do no harm” in a strong sense, since it can actually harm *everyone*. A key force in the mechanism that we uncover is that consensual risky partnerships can feature strategic complementarities, even in low-risk environments. These strategic complementarities can generate feedback effects that can lead to large

⁶ See Philipson and Posner (1993) for an interesting in-depth economic perspective on the HIV epidemic.

changes in the structure of risky partnerships after a relatively small intervention—overwhelming the direct effects of this intervention.

The result of this paper underscores the importance of taking into account the network of social interactions in theoretical and empirical epidemiological studies in order to understand the potential effects that different interventions have on social structure—and hence on behavior and welfare. 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 different interventions.

Appendix A. 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*

$$\mu_I = qp + (1 - qp)q, \quad (4)$$

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

$$\begin{aligned} \mu_X = & (1 - p)^2 \mu_I \\ & + p(1 - p) [q(2 - q) + (1 - q)^2 qp(2 - qp)] \\ & + p(1 - p) [q + (1 - q)p(pq + (1 - pq)q(2 - q))] \\ & + p^2 [q(2 - q) + (1 - q(2 - q))q(2 - q)p(2 - p)], \end{aligned} \quad (5)$$

the probability that N_1 is infected is

$$\mu_{N_1} = \mu_I + (1 - q)(1 - pq)\mu_I p, \quad (6)$$

the probability that N_2 is infected is


$$\mu_{N_2} = q + (1 - q)pq + (1 - q)^2 p^2 \mu_I, \quad (7)$$

and, finally, the probability that M_3 is infected is

$$\mu_{M_3} = \mu_{N_2} + (1 - \mu_{N_2})\mu_I p. \quad (8)$$

Proof. To see equation (4), 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 (5), 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 a 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

$$q(2 - q) + (1 - q)^2 qp(2 - qp) \quad (9)$$

To see this, note that the probability that X_1 is infected exogenously or from 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

$$q + (1 - q)p(pq + (1 - pq)q(2 - q)) \quad (10)$$

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

We conclude that each agent's expected probability of infection in this case is the average of expressions (9) and (10).

Case 2: Both edges X_1X_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 exogenously or from 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 (6), 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 either exogenously or from N_3 times the probability p that the edge N_1N_4 is live.

To see equation (7), 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 exogenously, from N_3 or from N_4 . That is, $\mu_{N_2} = q + (1 - q)p[q + (1 - q)p\mu_I]$, which is equivalent to equation (7).

To see equation (8), note that the probability that M_3 is infected is the probability μ_{N_2} that she becomes infected exogenously or from either M_1, M_2 or M_4 plus the complement probability times the probability that the edge M_3M_6 is live times the probability that M_6 is infected exogenously or from M_5, M_7 or M_8 . \square

Appendix B. Epidemiological model with random matching

In this section we describe a model similar to the one described in section 2 that features random matching instead of strategic choice of partners. We then discuss how both the strategic complementarities in risky partnerships and the negative welfare effects of partially-effective vaccines naturally vanish in this case. There are $n \geq 2$ agents, and four stages, listed below. For simplicity, we focus on the case in which n is even.

- Stage 1: **Network Formation.** Each agent announces how many partners he or she wants to have. Edges are then formed as follows: One pair of agents is selected uniformly at random, and if both of them have less edges than the number that they have announced, an edge is formed between them. This process continues until everyone but at most one agent has less edges than the number that she has announced.
- 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 directed 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 she derives from his partners (0 if no partners, s_1 if one partner, and $s_1 + s_2$ if two partners) less the cost of infection (c if infected, and 0 otherwise).

When everyone announces that she or she wants to have one partner, the outcome of the random network formation is necessarily a pair-complete network (Definition 4.1). Hence, it follows from our analysis in section 4 that this is an equilibrium if and only if $p \in (p^*, p^{**})$.

When the number of agents n is large, and each agent is part of at most two edges, the process of random matching implies that the event that an agent i transmits the infection to an agent j is approximately independent from the event that an agent $k \neq i$ transmits the infection to agent j . For simplicity, in what follows, we shall assume that n is large enough, so that these events can be safely treated as being independent (alternatively, we can assume, as is standard in the literature using epidemiological models with random matching, that there is a continuum of agents, so that these approximations are exact). Also, we focus on the natural case in which the equilibrium *per-partnership* infection probability is below $1/2$ (this is guaranteed, for example, by $c \geq 2s_1$).⁷

When an agent decides to have more partners, the probability of becoming infected from each partnership increases. Hence, everyone's incentives to have additional partners diminish after this deviation. In other words, *with random matching, there are no strategic complementarities in risky partnerships*. As a consequence, as we now discuss, *partially-effective vaccines cannot harm anyone in this random-matching version of our model*.

⁷ This rules out situations in which agents are *fatalistic*, in the sense that they are so likely to become infected by their first partnership that they might as well choose additional partnerships (e.g., Kremer 1996). Note that the mechanism that we illustrate in this paper does not rely on such extreme scenarios.

Suppose for contradiction that a partially-effective vaccine reduces someone's expected utility. Denote by μ and μ' the average equilibrium per-partnership infection probability before and after introducing the vaccine, respectively. We have that

$$\mu' > \mu. \quad (11)$$

Indeed, otherwise, each agent can have an expected utility at least as high after the intervention as she had before the intervention—by doing exactly as she was doing before.

Equation (11) implies that some agents must be choosing strictly more partners after the intervention. Moreover, the vaccine cannot reduce anyone's expected utility in equilibrium if either everyone is choosing zero partners or everyone is choosing two partners before the intervention, so we have that

$$s_1 \geq \underbrace{\mu pc}_{\text{expected infection cost of first partnership}} \quad \text{and} \quad s_2 \leq \underbrace{(1 - \mu p)\mu pc}_{\text{expected infection cost of second partnership}}$$

Hence, the following two cases are exhaustive. On the one hand, if no one chooses to have zero partners before the intervention, equation (11) implies that some agents choose to have two partners after the intervention, so $s_2 \geq (1 - \mu' p)\mu' pc > (1 - \mu p)\mu pc$, a contradiction.⁸ On the other hand, if some (but not all) agents choose to have zero partners before the intervention, we have that $\mu pc = s_1$, so that $\mu' pc > s_1$ (i.e., no one chooses to have any partners after the intervention), also a contradiction.

References

- Adams, C.F., 1879. Notes on Railroads Accidents. G.P. Putnam Sons.
- Aten, J., 2020. Bill Gates to Spend Billions on a Covid-19 Vaccine. Inc.com.
- 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., Easley, D., Kleinberg, J., Kleinberg, R., Tardos, É., 2011. Network formation in the presence of contagious risk. In: *Proceedings of the 12th ACM Conference on Electronic Commerce*. ACM, pp. 1–10.
- Blumenthal, J., Haubrich, R., 2017. Risk compensation in PrEP: an old debate emerges yet again. *Virtual Mentor*, 909–915.
- Chan, T.Y., Hamilton, B.H., Papageorge, N.W., 2015. Health, risky behaviour and the value of medical innovation for infectious disease. *Rev. Econ. Stud.* 83, 1465–1510.
- Chen, F., Toxvaerd, F., 2014. The economics of vaccination. *J. Theor. Biol.* 363, 105–117.
- Delavande, A., Kohler, H.-P., 2015. HIV/AIDS-related expectations and risky sexual behaviour in Malawi. *Rev. Econ. Stud.* 83, 118–164.
- Eaton, L.A., Kalichman, S.C., 2007. Risk compensation in HIV prevention: implications for vaccines, microbicides, and other biomedical HIV prevention technologies. *Curr. HIV/AIDS*, 165–172.
- Fenichel, E.P., Castillo-Chavez, C., Ceddia, M.G., Chowell, G., Parra, P.A.G., Hickling, G.J., Holloway, G., Horan, R., Morin, B., Perrings, C., et al., 2011. Adaptive human behavior in epidemiological models. *Proc. Natl. Acad. Sci.* 108, 6306–6311.
- Friedman, S.R., Curtis, R., Neaigus, A., Jose, B., Des Jarlais, D.C., 2006. *Social Networks, Drug Injectors' Lives, and HIV/AIDS*. Springer Science & Business Media.
- Galeotti, A., Rogers, B.W., 2013. Strategic immunization and group structure. *Am. Econ. J. Microecon.* 5, 1–32.
- Goyal, S., Jabbari, S., Kearns, M., Khanna, S., Morgenstern, J., 2016. Strategic network formation with attack and immunization. In: *International Conference on Web and Internet Economics*. Springer, pp. 429–443.
- Goyal, S., Vigier, A., 2015. Interaction, protection and epidemics. *J. Public Econ.* 125, 64–69.
- Greenwood, J., Kircher, P., Santos, C., Tertilt, M., 2019. An equilibrium model of the African HIV/AIDS epidemic. *Econometrica* 87, 1081–1113.
- Hoy, M., Polborn, M.K., 2015. The value of technology improvements in games with externalities: a fresh look at offsetting behavior. *J. Public Econ.* 131, 12–20.
- Jackson, M.O., Wolinsky, A., 1996. A strategic model of social and economic networks. *J. Econ. Theory* 71, 44–74.
- Jacquez, J.A., Simon, C.P., Koopman, J., Sattenspiel, L., Perry, T., 1988. Modeling and analyzing HIV transmission: the effect of contact patterns. *Math. Biosci.* 92, 119–199.
- Kremer, M., 1996. Integrating behavioral choice into epidemiological models of AIDS. *Q. J. Econ.* 111, 549–573.
- Marcus, J., Glidden, D., Mayer, K., Liu, A., Buchbinder, S., Amico, K., 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. *J. Polit. Econ.* 83, 677–725.
- Philipson, T., Posner, R., 1993. *Private Choices and Public Health: An Economic Interpretation of the AIDS Epidemic*. Harvard University Press.
- Pongou, R., Serrano, R., 2013. Fidelity networks and long-run trends in HIV/AIDS gender gaps. *Am. Econ. Rev. Pap. Proc.* 103, 298–302.
- Rowthorn, B.R., Toxvaerd, F., 2015. The optimal control of infectious diseases via prevention and treatment. Mimeo.
- Toxvaerd, F., 2017. On the Dynamics of Beliefs and Risky Sexual Behavior. Mimeo.
- Toxvaerd, F., 2019. Rational disinhibition and externalities in prevention. *Int. Econ. Rev.* 60, 1737–1755.

⁸ This inequality follows from the assumption that the equilibrium per-partnership probability $\mu'p$ is below 1/2.