UNIVERSITY OF CAMBRIDGE
Faculty of Economics

Institute for New Economic Thinking

Cambridge-INET Institute

Cambridge-INET Working Paper Series No: 2014/09

Cambridge Working Paper in Economics: 1458

INTERACTION, PROTECTION AND EPIDEMICS

Sanjeev Goyal (University of Cambridge) Adrien Vigier (University of Oslo)

ABSTRACT

Individuals respond to the risk of infectious diseases by restricting interaction and by investing in protection. We develop a model that examines the trade-off between these two actions and the consequences for disease prevalence.

There exists a unique equilibrium: individuals who invest in protection choose to interact more relative to those who do not invest in protection. Changes in the contagiousness of the disease have non-monotonic effects: as a result interaction initially falls and then rises, while disease prevalence too may initial increase and then decline.

We then consider a society with two communities that differ in their returns from interaction - High and Low. Individuals in isolated communities exhibit different behavior: the High community has a higher rate of protection and interaction and a lower rate of infection. Integration amplifies these differences.

Interaction, Protection and Epidemics

Sanjeev Goyal * Adrien Vigier†

June 11, 2014

Abstract

Individuals respond to the risk of infectious diseases by restricting interaction and by investing in protection. We develop a model that examines the trade-off between these two actions and the consequences for disease prevalence.

There exists a unique equilibrium: individuals who invest in protection choose to interact more relative to those who do not invest in protection. Changes in the contagiousness of the disease have non-monotonic effects: as a result interaction initially falls and then rises, while disease prevalence too may initial increase and then decline.

We then consider a society with two communities that differ in their returns from interaction – High and Low. Individuals in isolated communities exhibit different behavior: the High community has a higher rate of protection and interaction and a lower rate of infection. Integration amplifies these differences.

*Faculty of Economics & Christ's College, University of Cambridge. Email: sg472@cam.ac.uk

†Oekonomisk Institutt, University of Oslo. Email: adrien.vigier@uio.econ.no

We are grateful to the editor, Erzo Luttmer, and two anonymous referees for comments that have significantly improved the paper. We thank Diego Cerdiero, Julien Gagnan, and Anja Prummer for detailed comments on an earlier draft. We thank Eddie Dekel, Matthew Elliot, Andrea Galeotti, Mark Gersovitz, Chris Gilligan, Jesper Rudiger, Brian Rogers, Pauline Rutsaert, Larry Samuelson, Flavio Toxvaerd, and seminar participants at Cambridge, UCL (Louvain), Oslo and Oxford for useful discussions. Sanjeev Goyal thanks the Keynes Fellowship and the Cambridge-INET Institute for financial support.

1 Introduction

Interactions between individuals generate benefits but also facilitate the spread of viruses and infections. This tension is salient in infectious diseases such as influenza, HIV and tuberculosis.[1] Infectious diseases spread primarily through contact between human beings: so prevalence can be reduced through restricting interaction and by investing in protection. This paper develops a model that examines the trade-off between these two courses of action and its consequences for the prevalence of diseases.

In our model, a population of individuals faces a disease. Every individual chooses how much to interact with others and whether or not to protect himself. Interactions generate benefits but increase the risk of infection from infected others. There is a fixed cost to protection; for simplicity it is assumed that protection ensures immunity from disease. Individuals who protect themselves thus choose the maximal feasible interaction level. The protection rate, the interaction level, and the contagiousness of the disease determine disease prevalence (i.e., the infection rate).

We first establish that a (Nash) equilibrium exists and is unique. For a broad class of circumstances, equilibrium protection rates are interior: only a part of the population protects. Individuals who invest in protection interact more than those who do not. Restricted interaction and protection are substitutes. This relation is consistent with empirical observation. For example, in their well known study on British sexual attitudes and behavior, Wellings et al. (1994) report a positive correlation between the frequency of new partners and the use of condoms.

The contagiousness of a disease is a key parameter in the study of epidemics.[2] We find that equilibrium response to contagiousness is non-monotonic. There exists a threshold level of contagiousness: below this value, protection rates are zero, and the response to higher contagiousness is through reduced interaction. This is because below a threshold incurring

[1] There are 3 to 5 million cases of acute influenza and between 250,000 and 500,000 deaths are attributed to this infection, annually. In 2012, over 8.5 million people were infected with tuberculosis and 1.3 million deaths were attributed to it. In the same year, 2.3 million new cases of AIDS were reported and over 1.5 million people died due to the disease; over 36 million people have died due to HIV/AIDS so far (WHO (2013, 2014a, 2014b)).
The Internet reflects a similar tension: on-line interactions generate rewards but may serve as a conduit for the spread of viruses and worms which compromise user value. As energy, communication, travel, consumer interaction increasingly adopt digital networks, cybersecurity has emerged as a major concern. We discuss the relevance of our analysis for cybersecurity later in the introduction.

[2] For a classical exposition of the theory of epidemiology, see Anderson and May (1991). For a recent survey on epidemics, see Gersovitz (2011).

the fixed cost of protection is not worthwhile. Above the threshold, the returns from protection outweigh the fixed costs of protection. Greater contagiousness now induces greater protection and a corresponding increase in interaction. Infection rates too may vary non-monotonically – initially increasing and then declining – in contagiousness.

In our basic model individuals are homogenous. We then turn to a society with two communities that differ in their returns from interaction – High and Low. Individuals in isolated communities exhibit different behavior: the High community has a higher rate of protection and interaction. As communities integrate, protection and interaction further increase in the High community while they fall in the Low community. Integration thus leads to a fall (increase) in disease prevalence in the High (Low) community.

The theoretical prediction on the relation between returns and equilibrium behavior is broadly consistent with empirical observation. Wellings et al. (1994) report that single people have more partners and are much more likely to use condoms as compared to cohabiting couples. Philipson and Posner (1993) report a negative correlation between education/income and HIV infection: they surmise that higher income raises the returns from the future and thereby leads to greater investments in protection (the use of condoms). This in turn lowers the rate of infection.

Our model and its predictions are also related to cybersecurity.[3] The equilibrium property of positive correlation between protection and interaction is consistent with the findings of Anderson et al. (2007) and Moore, Clayton and Anderson (2011) on the positive relation between investments in security and Internet use. The model predicts that the High community will have higher protection and interaction: this is consistent with the fact that larger firms are more active in securing themselves as compared to smaller firms (Anderson et al (2007)).

Our paper is a contribution to the economic study of epidemics and cybersecurity. It is useful to separate the existing research in economic epidemiology into two strands. The first strand of work takes interaction as given and explores the response in protection rates. This work includes Brito, Sheshinski and Intriligator (1991), Geoffard and Philipson (1996, 1997), Francis (1997), Goldman and Lightwood (2002), Gersovitz and Hammer (2004), Galeotti and Rogers (2013), and Chen and Toxvaerd (2014). A second (and complementary) group of

[3]Estimates of the costs of cyber crime vary greatly. A recent study estimates the costs to be in the range of 300 billion USD to 1 trillion USD; this is between 0.4% and 1.4% of global GDP. A recent study for the UK Cabinet Office reported that the cost to the UK economy is over 27 billion USD per annum (Detica, (2011)). In 2009, roughly 10 million computers were infected with malware designed to steal online credentials. The annual damages caused by malware is of the order of 9.3 billion Euros in Europe, while in the US the annual costs of identity theft are estimated at 2.8 billion USD (Moore, Clayton and Anderson (2009)).

papers assumes that protection is absent and studies the response in interaction. This work includes Philipson and Posner (1993) and Kremer (1996). To the best of our knowledge, the present paper is the first attempt to provide a unified treatment of interaction and protection. The analysis yields a number of new insights; we highlight two of them via a comparison with the benchmark models.

Compared to the 'pure' protection benchmark our model yields lower rates of protection. This is because part of the population foregoes protection and responds instead by adapting interaction. But compared to that benchmark, infection rates are higher in our model. This tells us that differences in protection are 'insufficiently' compensated for by restricted interaction.

Next, consider the benchmark where interaction is the only choice variable (and there exists no protection). The more an unprotected individual interacts, the greater the chances that he becomes infected and, in turn, transmits the disease to others around him. Thus an increase in returns from interaction unambiguously raises the infection rate (Kremer (1996)). In our setting, on the other hand, the more individuals value interaction the less inclined they are to respond to an epidemic by reducing interaction. This implies that higher returns from interaction lead to higher protection rates and – in sharp contrast to the pure interaction benchmark – to lower infection.

Our results have potential policy implications. A first order implication is that demand for protection will be lower in a model where interaction levels are a choice variable.[4] An important insight from the economic models of epidemiology is the externality in individual protection. In our model, choosing protection creates an additional externality: *protected individuals interact more and this alters the pool of contacts.* We show that this expands the scope for policy intervention, as compared to the benchmark model of pure protection. Finally, our work suggests that subsidies on protection should target those valuing social interaction least, as doing so minimizes crowding-out effects.

The problem of computer network security has been extensively studied in electrical engineering and computer science; for an overview of this work see Alpcan and Basar (2011) and Anderson (2011). Aspnes, Chang and Yampolskiy (2006) (and the literature that follows them) study protection choices by nodes faced with a viral infection that spreads through a

[4]The experience with swine flu vaccines is worth mentioning in this regard. Most OECD countries have large stocks of swine flu vaccines; for instance, in England, the NHS stock is estimated to have around 40 million vaccines in stock. This large stock of vaccines has provoked much discussion in recent years, see e.g., Times (2010). Our theoretical result points to one relatively unexplored reason for this large stock: lowered international travel and interaction in response to public measures on quarantine and the fears of epidemic.

given network. Our paper contributes to this literature by proposing a general framework in which interaction (network) and security investments are both endogenous.

The rest of the paper is organized as follows. The model is presented in Section 2. Section 3 contains the analysis of the basic model. Section 4 studies the model with heterogenous individuals. Section 5 concludes. All proofs are presented in the Appendix.

2 Model

Social interaction. We consider a continuum of individuals with measure one, labeled by $[0,1]$. A typical individual, denoted by i, interacts with others in the population according to $k_i \in \mathbf{R}_+$. We can think of k_i as the number of hours which individual i spends drinking in a bar (or in on-line activity) every day.[5] Each hour, i strikes up conversation with another person, who he picks uniformly at random.[6] For an arbitrary subset A of individuals in the population and at any hour of the day, individuals in A constitute fraction $(\int_{i\in A} k_i di)(\int_{i\in[0,1]} k_i di)^{-1}$ of all individuals present in the bar. The probability that a new meeting of i is with an individual in subset A is thus given by:

$$\text{probability of new meeting within } A = \frac{\int_{i\in A} k_i di}{\int_{i\in[0,1]} k_i di}.$$

We suppose that k_i is a choice variable for individual i. Let $r_i(k_i) \geq 0$ denote the payoff of i from interacting according to k_i. We let

$$r_i(k_i) := \theta_i g(k_i), \tag{1}$$

where θ_i denotes the type of individual i. We will assume that $g(0) = 0$, $g'(.) > 0$ and $g''(.) \leq 0$ if $k_i < \overline{k}$, and $g(k) = g(\overline{k})$, for all $k > \overline{k}$. All individuals thus maximize interaction returns at $\overline{k}$. Higher θ_i may reflect greater returns from interaction due to differences in human, physical or financial capital.

The example of linear returns from interaction is a simple special case of our model:

[5] We borrow this example from Kremer (1996).

[6] We relax this assumption in Section 4, when considering heterogenous populations.

Example 1 *Linear returns.*

$$g(k) = \begin{cases} k & \text{if } 0 \leq k \leq \overline{k} \\ \overline{k} & \text{if } k > \overline{k} \end{cases} \tag{2}$$

A population is (ex ante) homogenous if all individuals have the same type, i.e. $\theta_i = \theta$ for all i, and heterogenous otherwise. Section 3 studies homogenous populations; section 4 takes up the case of heterogenous populations.

Epidemics. The population faces an infectious disease. Individuals become infected exogenously with probability ϵ, where $\epsilon > 0$. They can also be infected via interaction with other infected individuals. We let α denote the *contagiousness* of the disease, i.e. the probability of contracting the disease from interacting with an infected individual. The (total) probability that individual i becomes infected is given by

$$p_i = \epsilon + (1-\epsilon)\big(1 - (1-\alpha q_i)^{k_i}\big), \tag{3}$$

where q_i denotes the probability of meeting an infected individual. We assume that α is small so that we can approximate p_i as

$$p_i = \epsilon + (1-\epsilon)\alpha q_i k_i. \tag{4}$$

Protection. In addition to choosing how much to interact, each individual chooses whether or not to protect himself. Protection costs $\gamma > 0$, and immunizes the individual to the disease. The cost may be financial (as in the purchase of a condom, a vaccine or a computer security software) or it may reflect direct disutility (as in the case of condoms, or possible side-effects in the case of a vaccine). The binary variable v_i records the protection status of individual i: $v_i = 1$ if individual i protects, and $v_i = 0$ otherwise. We will say that an individual is unprotected or susceptible if $v_i = 0$. The *protection rate* V denotes the fraction of the population which protects. We let I denote the *infection rate*, i.e. the fraction of individuals in the population who are infected:

$$I := \int_{i \in [0,1]} (1 - v_i) p_i di. \tag{5}$$

Remark: We have opted for a simple formulation of protection: the costs are invariant with respect to level of interaction. This is a good model for a vaccine and for computer security software, but appears to be inappropriate for condoms where protection costs vary with frequency of interaction. We note that our main results continue to hold with variable costs of protection so long as the choice for protection is made ex-ante and applies to all interactions.

Payoffs. Infected individuals receive zero payoffs. Uninfected/healthy individuals enjoy returns from interaction given by (1). An unprotected individual's expected payoff is thus[7]

$$\Pi_i(k_i, v_i = 0) = \big(1 - p_i\big) r_i(k_i), \tag{6}$$

with p_i given by (4). A protected individual's expected payoff is

$$\Pi_i(k_i, v_i = 1) = r_i(k_i) - \gamma. \tag{7}$$

Equilibrium. The structure of the game as well as the distribution of types in the population are common knowledge. The model defines a game of complete information between individuals. We study the Nash equilibria of this game; an equilibrium is said to be interior if the protection rate $V \in (0, 1)$.

3 Analysis

This section presents our analysis of the choice of adaptive interaction and protection and its implications for infection in homogenous populations. We establish that there exists a unique equilibrium. In an interior equilibrium, some individuals protect and opt for maximal interaction, while others remain unprotected and economize on interaction. A rise in returns from interaction leads to greater interaction and protection and to lower infection. A rise in contagiousness of the disease has a non-monotonic effect on interaction: it initially falls but eventually rises. Protection is monotonically increasing; in an interior equilibrium, a rise in contagiousness lowers infection rates.

[7] A simple way to account for risk aversion is to replace (6) by $\Pi_i(k_i, v_i = 0) = \big(1 - p_i\big) r_i(k_i) - cp_i$, where c represents the cost of uncertainty for susceptible individuals. All results and proofs in this paper apply *mutatis mutandis* under this alternative specification. We are grateful to a referee for drawing our attention to risk aversion.

Our first result establishes existence and uniqueness of equilibrium.

Proposition 1 *Consider a homogenous population. An equilibrium exists and is unique.*

It is useful to briefly describe the ways in which interaction and protection shape the payoffs of the unprotected individual. The negative externality exerted by unprotected individuals is governed by the incidence q of the disease – the probability that an individual with whom an individual interacts is infected. The incidence of the disease can be written as the product of two terms: (i) the probability z of meeting a susceptible individual, and (ii) the probability p that a susceptible is in fact infected.

Since the incidence faced by all individuals is the same, and the payoff function is strictly concave in interaction, then all unprotected individuals must in any equilibrium choose the same interaction level. Let k_u denote the interaction of unprotected individuals in a given equilibrium. Since $\overline{k}$ maximizes protected individuals' payoffs, the probability of meeting a susceptible is:

$$z = \frac{k_u(1-V)}{k_u(1-V)+\overline{k}V}. \tag{8}$$

Substituting $q = zp$ in (4) and solving for p yields

$$p = \frac{\epsilon}{1-(1-\epsilon)\alpha z k_u}. \tag{9}$$

A pair (V, k_u) constitutes an (interior) equilibrium if and only if protected and unprotected individuals have equal payoffs, and

$$k_u = \arg\max_k (1-\epsilon)\big(1-\alpha z p k\big) r(k) \tag{10}$$

with z and p satisfying, respectively, equations (8) and (9).

Now the proof of the Proposition follows in two steps. We first consider the pure interaction game given a fixed protection rate V – this is a game played among unprotected individuals choosing how much to interact when the remaining fraction V of the population protects. We show that this is a game of strategic substitutes, with a unique equilibrium. The 'equilibrium' payoffs attained in the pure interaction game are strictly increasing and continuous in the protection rate V. As payoffs from protection are invariant, existence and uniqueness of equilibrium in the overall game of protection and interaction follows.

The trade-off between restricted interaction and protection plays a central role in our model. To illustrate its role, we now compare our equilibrium with (a) the equilibrium of the 'pure' protection model with interaction fixed exogenously at the no-disease optimum $\overline{k}$ (benchmark 1) and (b) the equilibrium of the 'pure' interaction model with no protection (benchmark 2), respectively. We consider here the contrast between our model and benchmark model 1; the contrast with benchmark model 2 is developed after the statement of Proposition 3 below.

An important insight from the economic models of epidemiology is that a part of the benefits from protection is the reduced risk of infecting others: something that agents fail to internalize. In our model, choosing protection creates an additional externality compared to the benchmark model 1: *protected individuals interact more and this alters the pool of contacts.* This remark suggests that the scope for policy intervention may be greater when individuals can adapt social interaction as compared to the benchmark of pure protection. Our next result illustrates this point in two ways. First, we show that public protection programs may prove superfluous in benchmark 1, and yet be socially desirable in our framework of adaptive interaction. Second, we show that adaptive interaction ultimately induces higher rates of infection. Let (V^*, I^*) and (V^1, I^1) denote the equilibrium protection and infection rates in our model and benchmark model 1, respectively.

Proposition 2 *Consider a homogenous population.*

1. $V^* \leq V^1$, *with strict inequality if* V^* *and* $V^1 \in (0,1)$.

2. V^* *and* $V^1 \in (0,1) \Rightarrow I^* > I^1$.

The proof of the first part builds on the observation that, for any given rate of protection, the resulting incidence of the disease is higher with fixed interaction as compared to the case where unprotected individuals adapt (and hence lower) interaction. This pushes up the returns from protection, inducing higher protection rates. The second part is more delicate: protection is higher in the benchmark model, but adaptive interaction in our model may compensate for the difference in protection rates. Our proof shows that while interaction adapts downward, it does so 'insufficiently'. As a result, incidence and infection are both higher in our model, as compared to the benchmark model 1.

We turn next to the comparative static implications of the choice between reduced interaction and protection. An increase in returns from interaction makes the former option less

attractive and this in turn enhances the appeal of protection. We will say that interaction undergoes a cross-sectional increase/decrease if *all* individuals interact more/less.

Definition 1 *A profile of interaction* $\left(k_i'\right)_{i\in[0,1]}$ *constitutes a* cross-sectional increase *of* $\left(k_i\right)_{i\in[0,1]}$ *if and only if, up to relabeling of the individuals:* $k_i' \geq k_i$ *for all* i*, with strict inequality for some subset of individuals with measure strictly greater than zero.*

Let $V(\theta)$ denote the equilibrium protection rate in homogeneous populations with type θ.

Proposition 3 *Suppose* $V(\theta) \in (0,1)$*. An increase in returns to interaction,* θ *leads to a cross-sectional increase in interaction and raises protection. If the returns function* $g(k)$ *is given by (2) an increase in* θ *lowers infection rates.*

At an interior equilibrium, changes in θ must leave individuals indifferent between protection and no protection. As protected individuals interact according to $\overline{k}$, their payoffs increase at the rate $g(\overline{k})$ (viz. (1)). Susceptibles' payoffs on the other hand increases at the lesser rate $(1-p)g(k_u)$, and must therefore be compensated by a decrease in the incidence $q = zp$. But the optimal interaction k_u must then increase in θ.

Next consider protection. If the protection rate were decreasing in θ, the probability z of meeting a susceptible would increase in θ. The infection probability p would increase too (viz. (9)), leading in turn to an increase in the incidence $q = zp$ and contradicting our previous observation that incidence is falling in θ. The protection rate thus increases in θ. Since susceptible interaction increases, an increase in θ leads to a cross-sectional increase in interaction.

Finally, consider the infection rate. The protection rate is increasing, thus setting a decreasing upper bound on the infection rate. However, interaction of the unprotected is increasing too. The key issue is therefore the elasticity of interaction. If this elasticity is sufficiently low then the higher protection rate will prevail. Observe that the infection rate is given by $I(\theta) = p(\theta)(1 - V(\theta))$, where $p(\theta)$ denotes the equilibrium probability of infection in a population with type θ. We show by direct computation that if payoffs are given by (2), then $p(\theta)$ is constant.[8]

Proposition 3 is helpful in bringing out the differences between our model and the 'pure' interaction benchmark model 2. In the absence of protection, the more individuals value

[8]In fact, the result developed here holds more generally – provided the elasticity of interaction with respect to contagiousness is no more than 1.

interaction, the greater the negative externality each agent imposes on others. This in turn implies that in the pure interaction benchmark an increase in θ unambiguously raises the incidence and hence the infection rate (Kremer (1996)). As Proposition 3 demonstrates, in our model the effects of increasing returns to interaction θ are very different. The more individuals value interaction, the less inclined they are to respond to an epidemic by reducing interaction: this pushes up the protection rate and leads to falling infection rates for a broad class of models.

We turn finally to the effects of contagiousness on behavior and infection rates. Our analysis establishes that there are two phases in the equilibrium response to contagiousness. In phase I, at low α, the response to higher contagiousness is entirely through reduced interaction. In phase II, at high α, greater contagiousness induces more protection but there is no reduction of social interaction. Hence, somewhat surprisingly, social interaction is non-monotonic in contagiousness: it first decreases but eventually rises.

Proposition 4 *Consider a homogenous population. Equilibrium protection increases with contagiousness. There exist* $\underline{\alpha}$ *and* $\overline{\alpha}$, $\underline{\alpha} \leq \overline{\alpha}$, *such that:*

1. *If* $\alpha < \underline{\alpha}$ *then an increase in* α *leads to a cross-sectional decrease in interaction. Infection rate changes depend on the elasticity of interaction with respect to* α.

2. *If* $\underline{\alpha} < \alpha < \overline{\alpha}$ *then an increase in* α *leads to a cross-sectional increase in interaction and a fall in the infection rate.*

We start by observing that if $V = 0$ then the negative externality exerted by unprotected individuals is strictly increasing in contagiousness, α. This implies that interaction of unprotected individuals, k_u, decreases in α and completes the arguments for part 1.

Next consider an interior equilibrium: an individual must be indifferent between protection and no protection. Since payoffs from protection are unaffected by α, the payoffs earned by susceptibles must be unchanged as well. This implies that the externality $\alpha q(\alpha)$ remains constant in α. As α is rising, the incidence $q(.)$ must be falling. Now consider protection. Suppose that protection is falling in α. Then the only way to reduce the incidence $q(\alpha)$ would be to reduce the interaction k_u. But we have already proven that the externality $\alpha q(\alpha)$ is constant in α, and so is k_u. This contradiction establishes that protection rates must be increasing in α. As susceptible interaction is constant and the fraction of protected individuals rises, it follows that an increase in contagiousness leads to a cross-sectional increase in interaction. This completes the argument for part 2.

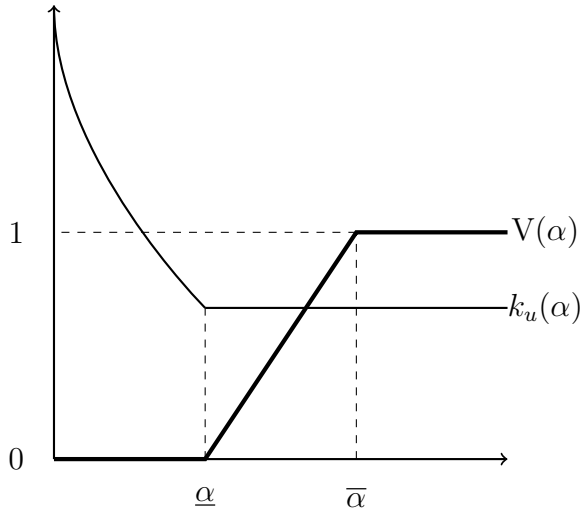

Figure 1

We conclude with a discussion on the effect of contagiousness on infection rates. Kremer (1996) observes that (in the absence of protection) higher contagiousness may lead to lower infection rates if the elasticity of interaction with respect to contagiousness is greater than one. The remark on infection rates in the first part follows similar logic.

The contrast is, however, sharp if we look at infection rates in an interior equilibrium (when $\alpha \in (\underline{\alpha}, \overline{\alpha})$: in this case higher contagiousness *always* lowers infection rates. The idea is as follows. The externality $\alpha q(\alpha)$ remains constant over this range. The probability of infection of an unprotected individual must therefore remain unchanged as well. The infection rate then falls, since the protection rate is rising in α.

4 Heterogenous populations

This section studies the effect of population heterogeneity on behavior and infection rates. We model heterogeneity in terms of differences in returns from interaction. Our goal is to examine the interplay between this heterogeneity and the choice between restricted interaction and protection. Observe that the case where the different 'types' of individuals are completely segregated is covered by Proposition 3 above. Our analysis reveals that the integration level between different types has powerful effects on behavior and infection rates.

We shall suppose that a heterogenous population consists of two equal size communities of individuals with types θ_H and θ_L respectively, $\theta_H > \theta_L$. A parameter β measures integration between the two communities. With probability β (resp., $1-\beta$), social interaction takes place across communities (resp., within one's own community). Thus, for $\beta = 0$ the two communities of individuals are completely segregated, while for $\beta = 1$ all interaction takes place across communities.[9] The probability that a new meeting of an individual in community J is with an individual in subset $A \subset J'$ is thus:

$$\text{probability of new meeting within } A \subset J' = \big((1-\beta)I_{\{J=J'\}} + \beta I_{\{J\neq J'\}}\big)\frac{\int_{i\in A} k_i di}{\int_{i\in J'} k_i di}.$$

where, as usual, $I_{\{.\}}$ denotes an indicator variable.

We first study the effects from integrating communities moderately. We then explore, through an example, the multiplicity of equilibria arising when most interaction takes place across communities. Throughout, we focus the analysis on the more interesting case where the equilibrium protection rate in a homogenous population with type θ_J is interior, i.e. $V(\theta_J) \in (0,1)$, $J \in \{H, L\}$.[10]

Consider first $\beta = 0$. When $\beta = 0$, results from Proposition 3, concerning homogenous populations, apply. Equilibrium is unique and entails strictly higher incidence in community L than in community H.

Consider now a small increase in β. Other things equal, an individual in community H faces higher incidence compared to $\beta = 0$, since he will interact with some probability with an individual from community L where the incidence is higher. This raises incentives to protect in community H. By symmetry, an individual in community L now faces lower incidence since he will interact, with some probability, with an individual in community H where the incidence is lower. This lowers incentives to protect in community L. Integrating communities moderately therefore benefits L types: As integration grows, the burden of protection tends

[9]For consistency, and to abstract from issues regarding community size, we assume that if interaction in community J is greater than it is in community J' then only a fraction of individuals in community J interact across communities. This ensures in a simple way that total interaction from community J to J' equals total interaction from J' to J. Formally, if $\int_{i\in J} k_i di > \int_{i\in J'} k_i di$ then partition community J as $J_1 \bigcup J_2$ such that $\int_{i\in J_1} k_i di = \int_{i\in J'} k_i di$, and individuals in J_2 interact within J_2 only.

[10]The analysis of the other cases is similar, but the results are plainer. For instance if $V(\theta_H) = V(\theta_L) = 0$ then for any β, equilibrium is unique and entails all individuals in both communities remaining unprotected. Similarly, if $V(\theta_H) = V(\theta_L) = 1$ then for any β, equilibrium is unique and entails all individuals in both communities protecting.

to fall on H types, while L types free-ride. The next proposition summarizes these insights, and explores their implications for interaction and infection in the two communities.

Proposition 5 *Consider a heterogenous population and suppose that* $V(\theta_J) \in (0,1)$, $J \in \{H, L\}$. *Then an equilibrium exists for all* β, *and is unique over a range* $[0, \overline{\beta}]$, *where* $\overline{\beta} \in (0, 1/2]$. *Moreover, over this range, increasing* β *induces (i) higher (lower) protection in community* H *(in community* L*), (ii) lower (higher) infection rate in community* H *(in community* L*) and (iii) a cross-sectional increase (decrease) in interaction within community* H *(within community* L*).*

We explained Part (i) above. We now develop the intuition underlying parts (ii) and (iii). By continuity, for low values of β the unique equilibrium is interior in both communities. So, for small values of β, individuals in each community must remain indifferent between protection and no protection. The payoffs of unprotected individuals therefore remain unchanged as we vary β. Since these payoffs are strictly falling in the incidence of the disease, this implies that the incidence *faced* by individuals in both communities must remain unaffected by changes in β. This in turn implies that, in each community, susceptible interaction must be unchanged, and so must their probability of infection. However, since the protection rate grows (falls) in H (L), the fraction of infected individuals must in fact decrease (increase) in community H (L). Finally, the arguments above establish that more integration induces higher protection in community H but no reduction of interaction by unprotected individuals in that community. Hence, an increase in β induces a cross-sectional increase in interaction by community H. Similar reasoning establishes the claim for community L.

Proposition 5 has interesting implications. To illustrate, suppose that prior to the stage at which individuals make a decision a social planner is able to allocate vaccines to a small subset of the population. If communities are largely segregated (β small), the unique equilibrium is interior in both communities. Any allocation of the social planner therefore induces complete crowding-out of private incentives, leaving the resulting equilibrium unaffected. Suppose, on the other hand, that communities are sufficiently integrated so that in the unique equilibrium all individuals in community H protect while all individuals in community L remain unprotected. In this case, allocating protection to high types once again crowds-out private incentives, whereas allocating vaccines to low types strictly improves welfare in the resulting equilibrium, due to the positive externalities from protection within community L. More generally, targeting individuals valuing interaction least tends to minimize crowding-out in

our model and as such, somewhat counter-intuitively, makes it attractive for public policy to protect individuals valuing social interaction least.

We have so far focused on the case of limited integration between communities. For greater integration there will typically exist multiple equilibria. To bring out this point in the simplest way, we consider an example where $\beta = 1$. We concentrate on the range of parameters where the isolated communities would have interior equilibria. Then in the case of integration, there exist three equilibria. The first equilibrium is interior in both communities. The other two equilibria are extremal: all individuals protect in one of the two communities, while all individuals remain unprotected in the other. In these equilibria, a protected individual interacts only with unprotected individuals, and vice versa. Protection is thus attractive for those who do protect, while remaining unprotected is optimal under the guarantee never to meet an infected individual. In this setting, it is possible to show that social welfare is highest in the corner equilibrium where the H type protect and the L type do not protect. The details of the computations are provided in the appendix.

5 Conclusion

Social and economic interaction generate rewards but also facilitate the spread of infections and diseases. Individuals can respond to this risk by restricting interaction and by investing in protection. This paper develops a model that examines the trade-off between these two actions and the implications for the prevalence of infections.

We study a setting in which a disease may be contracted either through direct infection or through contact with an infected individual. Individuals decide on how much to interact with others and whether to protect themselves. We study the equilibrium levels of protection and interaction and the population rate of infection.

We establish that there exists a unique equilibrium. We then derive a number of interesting implications of the co-determination of interaction and protection. There is a threshold property with regard to the effects of contagiousness of the disease. Below a threshold level, interaction falls in contagiousness, while it rises above the threshold. Infection rates too may move non-monotonically: below a threshold level they rise in contagiousness and above the threshold they fall with contagiousness.

In the basic model, all individuals are homogenous. We extend the model to study a society composed of two communities that differ in their returns from interaction – High and

Low. Individuals in isolated communities exhibit different behavior: the High community has a higher rate of protection and interaction. As communities integrate, protection and interaction further increase in the High community while they fall in the Low community. Integration thus leads to lower (higher) disease prevalence in the High (Low) community.

6 Appendix

Consider a homogenous population. Define the *pure interaction game given protection rate* V (PIGGV) as the game played among unprotected individuals choosing how much to interact when the remaining fraction V of the population protects. Let $q_i(k_u; V)$ denote the incidence of the disease facing individual i in the PIGGV, when other unprotected individuals interact according to k_u. Note that $q_i(k_u; V) = p_{-i}.z_i$, where z_i is the probability that an arbitrary meeting of i is with an unprotected individual and p_{-i} is the probability that an unprotected individual is in fact infected. Since the population is composed of a continuum of individuals, notice that $p_{-i} = p$ and $z_i = z$, where $z = k_u(1-V)/(k_u(1-V) + \overline{k}V)$. Moreover, using (4) yields $p = \epsilon + (1-\epsilon)\alpha p z k_u$ and, upon rearrangement

$$p = \frac{\epsilon}{1 - (1-\epsilon)\alpha z k_u}. \tag{11}$$

This shows, in particular, that $q_i(k_u; V) = p.z$ strictly increases in k_u (z increases in k_u, while p is increasing in each of z and k_u).

The payoffs of individual i in the PIGGV as a function of k_i as well as other players' actions k_{-i} can be written

$$\Pi_i(k_i, k_{-i}) = \big(1 - p_i(k_i, q_i(k_{-i}; V))\big) r(k_i) \tag{12}$$

where $p_i(k_i, q_i(k_{-i}; V)) = \epsilon + (1-\epsilon)\alpha q(k_{-i}; V) k_i$. Taking cross derivatives yields

$$\frac{\partial \Pi_i}{\partial k_i \partial k_{-i}} = -\alpha(1-\epsilon)\frac{\partial q_i}{\partial k_{-i}}\big(r(k_i) + k_i r'(k_i)\big) \tag{13}$$

The PIGGV is thus a game of strategic substitutes. Standard considerations then establish that an equilibrium of the PIGGV exists, is unique, and symmetric (uniqueness and symmetry follow the fact that the incidence of the disease is the same for all individuals, together with the strict concavity of payoffs with respect to interaction). Let in what follows $\Pi(V)$ denote

the equilibrium payoff achieved by unprotected individuals in the PIGGV. The next lemma is useful:

Lemma 1 *$\Pi(V)$ is continuously differentiable.*

Proof: First we show that $k_u(V)$ – the equilibrium interaction of unprotected individuals in the PIGGV – is continuously differentiable. Using first order conditions yields a continuously differentiable $G(.,.)$ implicitly defining $k_u(V)$ through the equation $G(k_u(V), V) = k_u(V)$. Hence, using the Implicit Function Theorem shows that $k_u(V)$ is continuously differentiable.

Next, note that $q_i(k_u; V)$ (the incidence of the disease facing individual i in the PIGGV, when other unprotected individuals interact according to k_u) is continuously differentiable. The equilibrium value of the incidence of the disease in the PIGGV, $q(V)$, is thus continuously differentiable too since $q(V) = q_i(k_u(V); V)$.

The proof is concluded by noting that the probability of infection is continuously differentiable in the incidence q, while the returns from interaction are continuously differentiable in k_u.

■

Proof of Proposition 1: Our proof uses the analysis of the PIGGV given above. Recall that $\Pi(V)$ denotes the equilibrium payoff achieved by unprotected individuals in the PIGGV. By vaccinating, an individual achieves payoff $r(\overline{k}) - \gamma$. Thus, an equilibrium of the overall game with protection rate V and in which unprotected individuals interact according to k_u exists if and only if k_u is the equilibrium interaction of unprotected individuals in the PIGGV and:

1. if $V = 0$ then $\Pi(0) > r(\overline{k}) - \gamma$;
2. if $V = 1$ then $\Pi(1) < r(\overline{k}) - \gamma$;
3. if $V \in (0, 1)$ then $\Pi(V) = r(\overline{k}) - \gamma$.

Next, we claim that $\Pi(V)$ is strictly increasing in V. To establish the claim, we first show that $q(V)$ – the equilibrium value of the incidence of the disease in the PIGGV – strictly decreases in V. Suppose, for a contradiction, that $q(V)$ increases in V. Each unprotected individual would then find it in his best interest to interact (weakly) less following a rise in V, i.e. k_u must decrease in V. Since $z = k_u(1 - V)/(k_u(1 - V) + \overline{k}V)$, this also implies that z must strictly decrease in V (more individuals protect, *and* unprotected individuals interact

less). But then, by (11) (giving p as a function of z and k_u), each unprotected individual must exhibit strictly lower probability of contracting the disease following a rise in V. We have thus established that if $q(V)$ increases in V, then both z and p must strictly decrease in V. But $q = zp$, so that $q(V)$ must strictly decrease in V too – an obvious contradiction. Hence, $q(V)$ – the equilibrium value of the incidence of the disease in the PIGGV – strictly decreases in V. The maximum payoff achieved by unprotected individuals in the PIGGV thus strictly increases in V. This establishes that $\Pi(V)$, the *equilibrium* payoff achieved by unprotected individuals in the PIGGV, strictly increases in V.

Combining the first part of the proof with the fact that $\Pi(V)$ strictly increases in V establishes existence and uniqueness of an equilibrium. Specifically:

1. if $\Pi(0) > r(\overline{k}) - \gamma$ then $V = 0$ in the unique equilibrium;
2. if $\Pi(1) < r(\overline{k}) - \gamma$ then $V = 1$ in the unique equilibrium;
3. if $\Pi(0) < r(\overline{k}) - \gamma < \Pi(1)$ then the equilibrium protection rate is uniquely determined by
$$\Pi(V) = r(\overline{k}) - \gamma. \tag{14}$$

■

Proof of Proposition 2: We begin with Part 1. Let $\Pi^1(V)$ denote the payoffs of unprotected individuals with fixed interaction $\overline{k}$ when a fraction V of the population protects. Fixing V, note that the probability z of meeting an unprotected individual is strictly greater with fixed interaction than in the PIGGV. Since, by (11), p is increasing in each of z and k_u this also implies that the probability of infection is strictly greater with fixed interaction than in the PIGGV as is ultimately, the incidence $q = zp$. This establishes that $\Pi(V)$, the maximum payoff achieved by unprotected individuals (in equilibrium) in the PIGGV, must be greater than $\Pi^1(V)$. That $V^* \leq V^1$, where $V^*/\ V^1$ denote the equilibrium protection rate under adaptive/fixed interaction, now follows immediately ($V^*/\ V^1$ are uniquely determined by the intersection of the strictly increasing functions $\Pi(V)/\ \Pi^1(V)$ with $r(\overline{k}) - \gamma$).

We next prove Part 2. V^* and $V^1 \in (0, 1)$ implies that in both cases (fixed and adaptive) the equilibrium payoffs achieved by unprotected individuals equal $r(\overline{k}) - \gamma$. Since for any given value of the incidence the option to adapt interaction raises the maximum payoffs attainable

by unprotected individuals, we must have $q^* > q^1$. The proof is concluded by noting that

$$\begin{aligned} I^* - I^1 &= p^*(1-V^*) - p^1(1-V^1) \\ &= p^*[(1-V^*) + z^* - z^*] - p^1[(1-V^1) + z^1 - z^1] \\ &= [q^* - q^1] + p^*[(1-V^*) - z^*] \\ &> 0. \end{aligned}$$

The last equality follows from noting that $(1 - V^1) = z^1$. The inequality is due to the remark above that $q^* > q^1$, and the observation that $(1 - V^*) - z^* > 0$.

■

Proof of Proposition 3: The proof is in three parts. We first establish $dq(\theta)/d\theta < 0$. We then show that $dV(\theta)/d\theta > 0$. Finally, we show that with payoffs given by (2) then $dI(\theta)/d\theta < 0$.

Part 1: Let $p(k_i, q_i)$ denote the infection probability of i interacting according to k_i and facing incidence q_i (notice that, fixing k_i, the probability of infection is independent of θ_i). Observe that $\partial p/\partial q_i > 0$.

The equilibrium indifference condition between vaccinating or not can be written as

$$\theta g(\overline{k}) - \gamma = \Big(1 - p(k_u(\theta), q(\theta))\Big)\theta g(k_u(\theta)) \tag{15}$$

where $k_u(\theta)$ and $q(\theta)$ denote equilibrium values for unprotected interaction and incidence, as a function of type θ. Differentiating with respect to θ yields

$$g(\overline{k}) = \Big(1 - p(k_u(\theta), q(\theta))\Big)g(k_u(\theta)) - \frac{\partial p}{\partial q}\frac{dq}{d\theta}\theta g(k_u(\theta)) \tag{16}$$

where we have used the fact that $k_u(\theta)$ is optimal for unprotected individuals, given $q(\theta)$. Hence

$$\frac{dq}{d\theta} = -\frac{g(\overline{k}) - \Big(1 - p(k_u(\theta), q(\theta))\Big)g(k_u(\theta))}{\frac{\partial p}{\partial q}\theta g(k_u(\theta))}. \tag{17}$$

Step 1 is concluded by noting that both numerator and denominator in the expression above are strictly positive.

Part 2: First notice using (10) that the first Part of the proof implies $dk_u(\theta)/d\theta > 0$ (lower

incidence reduces the marginal cost of interaction). Suppose now for a contradiction that $dV(\theta)/d\theta \leq 0$. Since $z = k_u(1-V)/(k_u(1-V)+\overline{k}V)$, we obtain $dz(\theta)/d\theta > 0$. This in turn yields $dp(\theta)/d\theta > 0$, where we have used (11). Finally, since $q = zp$, we obtain $dq(\theta)/d\theta > 0$, contradicting step 1. This establishes $dV(\theta)/d\theta > 0$.

Part 3: Notice that $I(\theta) = p(\theta)(1-V(\theta))$, where $p(\theta) = p(k_u(\theta), q(\theta))$. Since by Part 2 we have $dV(\theta)/d\theta > 0$ then $dp(\theta)/d\theta \leq 0 \Rightarrow dI(\theta)/d\theta < 0$. Straightforward computations based on (2) establish $k_u(\theta) = (2\alpha q(\theta))^{-1}$. The proof is concluded by noting that in this case $p(\theta) = \epsilon + (1-\epsilon)\alpha q(\theta) k_u(\theta) = (1+\epsilon)/2$.

■

Proof of Proposition 4: Our proof exploits the impact of α on the equilibrium of the PIGGV. Specifically, we first show that a rise in α induces a rise in the equilibrium value of the externality $\alpha.q(V;\alpha)$ in the PIGGV. We then go on to show that $\Pi(V;\alpha)$, the equilibrium payoff achieved by unprotected individuals in the PIGGV, strictly decreases in the contagiousness α. Finally we show that a (uniform) fall of the equilibrium payoff achieved by unprotected individuals in the PIGGV implies a rise in the equilibrium protection rate of the overall game.

Step 1: Let $\alpha_2 > \alpha_1$, and fix V. Let $q(V;\alpha)$ denote the equilibrium incidence in the PIGGV under contagiousness α. Suppose for a contradiction that $\alpha_2.q(V;\alpha_2) \leq \alpha_1.q(V;\alpha_1)$. An unprotected individual would then find it optimal to interact at least as much under α_2 as under α_1 (due to lower marginal cost of interaction). Since V is fixed, we must then have $z_2 \geq z_1$. But then, from the probability equation in (11), each unprotected individual has higher probability of infection under α_2. Given that $q = zp$, this implies $q(V;\alpha_2) \geq q(V;\alpha_1)$, and ultimately $\alpha_2.q(V;\alpha_2) > \alpha_1.q(V;\alpha_1)$ – contradicting our original assumption. Thus a rise in α induces a rise in the equilibrium externality $\alpha.q(V,\alpha)$ in the PIGGV.

Step 2: From step 1, the (PIGGV) equilibrium value of $\alpha.q(V;\alpha)$ strictly increases in α. In equilibrium, the maximum payoff achieved by unprotected individuals in the PIGGV thus strictly decreases in α. This implies that $\Pi(V;\alpha)$, the *equilibrium* payoff achieved by unprotected individuals in the PIGGV, strictly decreases in α.

Step 3: In the proof of Proposition 1 it was shown that the equilibrium protection rate V is determined by the intersection of (the strictly increasing function of V) $\Pi(V;\alpha)$ and $r(\overline{k}) - \gamma$. From step 2, we know that $\Pi(V,\alpha)$ is strictly decreasing in α. For equilibrium, it then follows that the protection rate V must increase in α. Specifically, there exist $\underline{\alpha}$ and $\overline{\alpha}$, $\underline{\alpha} \leq \overline{\alpha}$, such that in equilibrium: $V = 0$ if $\alpha < \underline{\alpha}$, $V \in (0,1)$ if $\alpha \in [\underline{\alpha}, \overline{\alpha}]$, and $V = 1$ if $\alpha > \overline{\alpha}$. (Notice that in our setting, $V = 1$ is not sustainable in equilibrium, and so $\overline{\alpha} = 1$).

We can now conclude the proof of the proposition. Let $\underline{\alpha}$ and $\overline{\alpha}$ be as defined in step 3. Consider first $\alpha < \underline{\alpha}$. The equilibrium protection rate is $V = 0$. It follows that (in the overall game) the equilibrium value of the externality is $\alpha.q(0;\alpha)$. From step 1, $\alpha.q(0;\alpha)$ strictly increases in α. From the optimization problem (10), the equilibrium interaction k_u of unprotected individuals thus strictly decreases in α. Consider next $\alpha \in [\underline{\alpha}, \overline{\alpha}]$. We have over this range $V \in (0,1)$; so individuals are indifferent between vaccinating or not vaccinating. The maximum payoff attainable by unprotected individuals is thus $r(\overline{k}) - \gamma$ over the entire range – the payoff they could achieve by vaccinating. Since the maximum payoff attainable by unprotected individual is strictly decreasing in the externality $\alpha.q$, the equilibrium value of $\alpha.q$ must remain constant over the whole range $\alpha \in [\underline{\alpha}, \overline{\alpha}]$. This in turn implies that over this range (i) the equilibrium incidence q strictly decreases, (ii) the equilibrium interaction k_u of unprotected individuals remains constant, and (iii) the equilibrium probability of infection p remains constant too. Since $q = zp$, the equilibrium value of z thus strictly falls over this range. The identity

$$V = \frac{k_u(1-z)}{k_u + z(\overline{k} - k_u)} \tag{18}$$

then shows that V strictly rises over this range. Finally, for $\alpha > \overline{\alpha}$ all individuals protect in equilibrium, and each individual interacts according to $\overline{k}$.

■

Proof of Proposition 5: The proof of existence follows along the lines of Proposition 1, and is therefore omitted. We next show that an equilibrium is unique (and interior in both communities) for β positive but small.

The equilibrium protection rates for homogenous populations $V(\theta_J) \in (0,1)$, $J \in \{H, L\}$. At $\beta = 0$, the unique equilibrium is thus interior in both communities. Let Q_J denote the incidence within community J, i.e.

$$Q_J = \frac{\int_{i \in J} (1 - v_i) p_i k_i di}{\int_{i \in J} k_i di}. \tag{19}$$

Making β and Q_{-J} into parameters and focusing on community J we can – by analogy with the analysis of homogenous populations – define within community J the pure interaction game given protection rate V_J. This is a game played among unprotected individuals in community J choosing how much to interact when the remaining fraction V_J of community J protects, a fraction β of interaction occurs across communities, and the incidence of the disease in the other community is Q_{-J}. Let $\Pi(V_J; \beta, Q_{-J})$ denote the equilibrium payoff achieved by the

players in this game; notice that $\Pi(V_J; 0, Q_{-J}) = \Pi(V_J)$, where $\Pi(V_J)$ denotes the equilibrium payoff achieved by unprotected individuals in the PIGGV of Proposition 1. Repeating the same steps as in the proof of Proposition 1 establishes that $\Pi(V_J; \beta, Q_{-J})$ is strictly increasing in V_J. Taylor expanding yields

$$\Pi(V_J; \beta, Q_{-J}) = \Pi(V_J; 0, Q_{-J}) + O(\beta) = \Pi(V_J) + O(\beta). \tag{20}$$

The payoff to protected individual in community J is $r_J(\overline{k}) - \gamma \in (\Pi(0), \Pi(1))$, $J \in \{H, L\}$, since the unique equilibrium in a homogenous population of type θ is interior for both types $\theta = \theta_H$ and $\theta = \theta_L$. This implies, by (20), that for β small $r_J(\overline{k}) - \gamma \in (\Pi(0, \beta, Q_{-J}), \Pi(1, \beta, Q_{-J}))$, $J \in \{H, L\}$, and arbitrary Q_{-J}. Hence, for β small any equilibrium of our game must be interior for both communitys (i.e. $V_J \in (0, 1)$, $J \in \{H, L\}$).

Lastly, we show that, for any β, at most one equilibrium exists that is interior for both communities. In an interior equilibrium for both communities, an individual in community J must be indifferent between vaccinating and not vaccinating. The payoff from protection is constant, and so is the payoff to an unprotected individual. The payoffs to an unprotected individual are strictly falling in disease incidence, so an individual in community J must be facing a constant incidence $q(\theta_J)$ (recall here that $q(\theta_J)$ denotes the equilibrium incidence in a homogenous population with type θ_J). An interior equilibrium thus entails $q(\theta_J) = \beta Q_{-J} + (1 - \beta) Q_J$, $J \in \{H, L\}$. Solving for Q_J, $J \in \{H, L\}$, yields

$$Q_H = \frac{1}{2}[q(\theta_L) + q(\theta_H) + \frac{q(\theta_L) - q(\theta_H)}{2\beta - 1}] \tag{21}$$

$$Q_L = \frac{1}{2}[q(\theta_L) + q(\theta_H) - \frac{q(\theta_L) - q(\theta_H)}{2\beta - 1}]. \tag{22}$$

Q_H and Q_L are thus uniquely determined. Since individuals in community J face $q(\theta_J)$, interaction of unprotected individuals in community J is uniquely determined too, as is the probability of infection for these individuals. Equation (19) now determines V_J. This suffices to show that at most one equilibrium exists that is interior for both communities.

We next establish parts 1-3 of the Proposition.

From Proposition 3, we know that incidence is falling in type of population; so $q(\theta_L) - q(\theta_H) > 0$. Using (21)-(22), in equilibrium, Q_H is therefore strictly decreasing in β while Q_L is strictly increasing. Since individuals in community J face $q(\theta_J)$ throughout, interaction of unprotected individuals in community J remains fixed (in equilibrium) as we vary β. Simi-

larly, as we vary β, the probability that an unprotected individual becomes infected remains constant. An inspection of the expression for Q_J in (19) reveals that changes in incidence must occur entirely through adjustments in the rate of protection, V_J. V_H therefore strictly increases in β while V_L strictly decreases. This establishes part 1 of the Proposition.

Next consider Part 2: Observe that infection depends on the rate of protection and the level of interaction of the unprotected. As we vary β, the is rising in community H and falling in community L, while the interaction remains constant in both communities.

Finally, consider part 3: It follows then that an increase in β induces a cross-sectional increase in interaction of community H and a cross-sectional decrease in interaction in community L.

■

Derivations for the case $\beta = 1$: $q(\theta_J)$, recall, denotes the equilibrium incidence of the disease in a homogenous population with type θ_J. We first show that $q(\theta_J) \in (0, \epsilon)$ implies $V(\theta_J) \in (0, 1)$, i.e. that the unique equilibrium in a homogenous population with type θ_J is interior. $V(\theta_J) = 1$ is impossible, as this would imply $q(\theta_J) = 0$. $V(\theta_J) = 0$, on the other hand, would imply $q(\theta_J) = p(\theta_J)$, the equilibrium probability of infection for unprotected individuals. Thus, by (4), $q(\theta_J) > \epsilon$ (recall, ϵ is the probability of contracting the disease exogenously), contradicting our initial assumption. This finishes to establish that $V(\theta_J) \in (0, 1)$, $J \in \{H, L\}$.

Since the unique equilibrium in a homogenous population is interior for both types, individuals of type θ_J are exactly indifferent between vaccinating or not when facing the incidence $q(\theta_J)$. By the same token, they strictly prefer vaccinating if facing $q > q(\theta_J)$, and strictly prefer remaining unprotected if facing $q < q(\theta_J)$.

As in Proposition 5, let q_J denote the incidence facing individuals in community J and Q_J the incidence within community J. Thus here $q_J = Q_{-J}$, since an individual in community J interacts exclusively with individuals in communities $-J$. Suppose an equilibrium exists in which some individual in community J protects. We will first show that if all individuals in community $-J$ remain unprotected, then it must be that all individuals in community J in fact protect. We will then show that if some individual in community $-J$ protects, then the equilibrium must be interior.

If all individuals in community $-J$ remain unprotected, each of them must be infected with probability ϵ at least , inducing $q_J = Q_{-J} \geq \epsilon$. But $q(\theta_J) < \epsilon$, so if all individuals in community $-J$ remain unprotected then *all* individuals in community J must protect. This in turn makes it optimal for individuals in community $-J$ not to protect.

Next, suppose that some individual in $-J$ protects. We must then have $q_{-J} = Q_J \geq$

$q(\theta_{-J})$. Since $q(\theta_{-J}) > 0$, some individual in community J must therefore be unprotected. Thus, in community J, some individuals protects while others do not. By symmetry, the same must then be true for community $-J$. This shows that if some individual in community $-J$ protects, then the equilibrium must be interior in both communities. Furthermore, this equilibrium is fully characterized by the condition $Q_J = q(\theta_{-J})$, $J \in \{H, L\}$. The arguments above establish that there exist exactly three equilibria, satisfying the description given in the statement of the Proposition.

Proof of welfare ranking: Let W_1, W_2 and W_3 be the average social welfare in the three equilibria identified above. We establish that $W_1 > W_2 > W_3$.

As in Proposition 3, let $p(k_i, q_i)$ denote the infection probability of i interacting according to k_i and facing incidence q_i. We then have

$$W_1 = \frac{1}{2}\Big(\theta_H g(\overline{k}) - \gamma\Big) + \frac{1}{2}\Big(1 - p(k_u(\theta_L), 0)\Big)\theta_L g(k_u(\theta_L)) \tag{23}$$

$$W_2 = \frac{1}{2}\Big(1 - p(k_u(\theta_H), 0)\Big)\theta_H g(k_u(\theta_H)) + \frac{1}{2}\Big(\theta_L g(\overline{k}) - \gamma\Big) \tag{24}$$

where $k_u(\theta_J)$ denotes the optimal interaction of unprotected individual in community J facing zero incidence. Thus $W_1 - W_2 = 1/2[T(\theta_H) - T(\theta_L)]$, where

$$T(\theta) = \Big(\theta g(\overline{k}) - \gamma\Big) - \Big(1 - p(k_u(\theta), 0)\Big)\theta g(k_u(\theta)). \tag{25}$$

Differentiating yields

$$\frac{dT}{d\theta} = g(\overline{k}) - \Big(1 - p(k_u(\theta), 0)\Big)g(k_u(\theta)) \tag{26}$$

where we have used the fact that $k_u(\theta)$ is optimal for an unprotected individual in community J facing zero incidence. But then, clearly, $\frac{dT}{d\theta} > 0$ and so $W_1 > W_2$.

That both W_1 and W_2 are strictly greater than W_3 follows from the observation that $W_3 = 1/2(\theta_H g(\overline{k}) - \gamma) + 1/2(\theta_L g(\overline{k}) - \gamma)$, while the maximum payoff attained by unprotected individuals is strictly greater when facing zero incidence than when facing $q(\theta_J) > 0$ in an interior equilibrium.

■

7 References

1. Alpcan, T. and T. Basar(2011), *Network Security: A Decision and Game Theoretic Approach.* Cambridge University Press. Cambridge, England.

2. Anderson, R. (2008), *Security Engineering.* Second Edition. Wiley.

3. Anderson, R., Bohme, R., Clayton, R., and T. Moore (2007), *Security economics and the internal market*, European Network and Information Security Agency.

4. Anderson R. M., and R.M. May (1991), *Infectious Diseases of Humans: Dynamics and Control.* OUP, Oxford.

5. Aspnes, J., K. Chang, A. Yampolskiy (2006), Inoculation strategies for victims of viruses and the sum-of-squares partition problem, *Journal of Computer and System Sciences*,

6. Brito, D.L., E. Sheshinski and M.D. Intriligator (1991), Externalities and compulsory vaccinations, *Journal of Public Economics*, 45, 69-90.

7. Chen, F., and F. Toxvaerd (2014), The economics of vaccination, *Mimeo*, Cambridge University.

8. Center for Strategic and International Studies (2013), *The Economic Impact of Cybercrime and Cyber Espionage.*

9. Detica and Cabinet Office (2011), *The Cost of Cyber Crime.*

10. Francis, P. J. (1997), Dynamic epidemiology and the market for vaccination, *Journal of Public Economics*, 63, 3, 383-406.

11. Galeotti, A., and B. Rogers (2013), Strategic Immunization and Group Structure, *American Economic Journal-Micro*, 5, 2, 1-32.

12. Geoffard, P-Y., and T. Philipson (1996), Rational Epidemics and Their Public Control, *International Economic Review*, 37, 3, 603-624.

13. Geoffard, P-Y., and T. Philipson (1997), Disease eradication: private versus public vaccination, *American Economic Review*, 87, 1, 222-230.

14. Gersovitz, M. (2011), The economics of infection control, *Annual Review Resource Economics*, 3, 27796

15. Gersovitz, M., and J. S. Hammer (2004), The economical control of infectious diseases, *Economic Journal*, 114, 1-27.

16. Goldman, S. M. and J. Lightwood (2002), Cost optimization in the SIS model of infectious disease with treatment, *Topics in Economic Analysis and Policy*, 2, 1, 1-22.

17. Hays, J. N. (2009), *The Burdens of Disease: Epidemics and Human Response in Western History.* Rutgers University Press.

18. Kremer, M. (1996), Integrating Behavioral Choice into Epidemiological Models of AIDS, *Quarterly Journal of Economics*, 111, 2, 549-73.

19. Moore, T., Clayton, R., and R. Anderson (2011), Economics and Internet Security: a Survey of Recent Analytical, Empirical and Behavioral Research *Harvard University Computer Science Department, Discussion Paper TR-03-11.*

20. Philipson, T. and R. Posner (1993), *Private Choices and Public Health: The AIDS Epidemic in an Economic Perspective.* Harvard University Press.

21. World Health Organization (2013), *HIV/AIDS.* Fact Sheet Number 2013.

22. World Health Organization (2014a), *Influenza (Seasonal).* Fact Sheet Number 211.

23. World Health Organization (2014b), *Tuberculosis.* Fact Sheet Number 104.