

Interaction, protection and epidemics[☆]Sanjeev Goyal^a, Adrien Vigier^b^a Faculty of Economics & Christ's College, University of Cambridge, United Kingdom^b Økonomisk Institutt, University of Oslo, Norway

ARTICLE INFO

Article history:

Received 18 September 2013

Received in revised form 23 February 2015

Accepted 24 February 2015

Available online 6 March 2015

JEL classification:

I12

D85

Keywords:

Social interaction

Epidemics

Vaccination

ABSTRACT

Individuals respond to the risk of contagious infections 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 infection rates.

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 infection rates 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.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

Interactions between individuals generate value, but facilitate the spread of infections. This tension is salient in diseases such as influenza, HIV and tuberculosis, but also applies to the Internet and other digital networks.¹ In all these examples, infection spreads primarily through interpersonal contacts: so prevalence can be reduced by restricting interaction and/or 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 spread of infections.

[☆] 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. Adrien Vigier thanks the Gates Cambridge Fellowship for its financial support.

E-mail addresses: sg472@cam.ac.uk (S. Goyal), adrien.vigier@uio.econ.no (A. Vigier).

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,b). 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, cyber-security has emerged as a major concern. We discuss the relevance of our analysis for cyber-security later in the introduction.

In our model, a population faces the risk of becoming infected. Every individual chooses how much to interact with others in the population, and whether or not to protect himself. Interactions generate benefits but increase the risk of infection from infected others. Protection is available, at a fixed cost. The protection rate, the profile of interaction, and the contagiousness of the infection together determine the extent of the disease in the population.

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 determines the probability of becoming infected from interacting with infected individuals, and 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 only. This threshold reflects the fixed costs associated with protection: below the threshold incurring the costs is not worthwhile. Above the threshold, returns from protection outweigh costs. Greater contagiousness now induces greater

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

protection, and a first-order stochastic dominance shift in the profile of 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 falling (rising) infection 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 et al. (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 et al. (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 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.

Consider next the ‘pure interaction’ benchmark, where protection is unavailable. The more a susceptible interacts, the greater the chances that he becomes infected and, in turn, transmits the disease to others around him. ‘Pure interaction’ models are thus characterized by the property that increasing returns from interaction raises infection (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 ‘pure protection’ benchmark. 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 (2008). Aspnes et al. (2006) (and the literature that follows them) study protection choices by nodes faced with a viral infection that spreads through a 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 basic model is presented in Section 2, and analyzed in Section 3. Section 4 studies heterogeneity. Section 5 concludes. All proofs are presented in the online-appendix.

2. Model

The basic model has the following features. A continuum of individuals live for two periods. Each period, the agents individually decide how much social interaction to have (in a bar say, or in on-line activity). Social interaction is beneficial, but has drawbacks: interaction in the first period raises chances of a contagious infection, which reduces pay-offs in the second period. To mitigate the chances of infection, each agent faces two options. He may reduce interaction in the first period, or invest in protection. The protection rate, the profile of interaction, and the contagiousness together determine the fraction of the population who become infected. We next lay out the details and notation of this model.

2.1. Social interaction

A typical agent is labeled i . We let $k_{it} \geq 0$ denote the socialization ‘effort’ of agent i in period $t = 1, 2$: each period, i selects k_{it} individuals uniformly at random from the pool of available contacts.⁵ We borrow this interpretation from Kremer (1996).⁶ Given an arbitrary subset A of individuals, the probability that a new contact of i is with an individual in A is

$$\mathbb{P}_i(\text{select contact in } A) = \frac{\int_{j \in A} k_{jt} dj}{\int_{j \in [0,1]} k_{jt} dj}. \quad (1)$$

2.2. Infection

There are two ways to become infected. An individual may become infected exogenously, with probability $\epsilon > 0$. Or an individual may

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 and Cabinet Office, 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 et al. (2011).

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. 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.

⁵ The assumption on uniform selection is relaxed in Section 4, when considering heterogeneous populations. There, we allow for ‘search’ to be directed.

⁶ As in Kremer (1996), we abstract from strategic complementarity in social interactions. Making strategic complementarities more explicit and significant would introduce the possibility of multiple equilibria. While this would enrich the analysis, our thought is that the key trade-offs we identify would remain important in this richer framework.

contract the disease by contagion. We let α denote the contagiousness of the disease, i.e. the probability of contracting the disease from interacting with an infected individual.

It is useful to define the incidence q_{it} facing agent i in period t , indicating the probability that a new meeting of i is with an infected individual. The total probability that (unprotected) individual i becomes infected during period t is then

$$p_{it} = \epsilon + (1 - \epsilon)(1 - (1 - \alpha q_{it})^{k_{it}}). \quad (2)$$

Throughout, we assume that α is small so that p_{it} is well approximated by

$$p_{it} = \epsilon + (1 - \epsilon)\alpha q_{it} k_{it}. \quad (3)$$

2.3. Protection

Each individual faces the option to invest in protection, for a cost $\gamma > 0$.⁷ The cost may be financial (as in the purchase of a condom, a vaccine or a computer security software) or reflect direct disutility (as in the case of condoms, or possible side-effects in the case of a vaccine). The binary variable v_{it} records the protection status of individual i in period t : $v_{it} = 1$ if individual i protects, and $v_{it} = 0$ otherwise. Protection perfectly immunizes an individual, so $p_{it} = 0$ if $v_{it} = 1$.⁸ We will say that an individual is unprotected or susceptible if $v_{it} = 0$. The protection rate V_t denotes the fraction of the population investing in protection. We let I_t denote the infection rate, i.e. the fraction of individuals who become infected during period t :

$$I_t := \int_{i \in [0,1]} (1 - v_{it}) p_{it} di. \quad (4)$$

2.4. Returns from interaction

We let $r_i(k_{it}) \geq 0$ denote the returns from interaction for individual i interacting according to k_{it} during period t . These are defined by

$$r_i(k_{it}) := \theta_i g(k_{it}), \quad (5)$$

where θ_i denotes the type of individual i , and g is a twice continuously differentiable strictly concave function maximized at $\bar{k} > 0$. Higher θ_i may reflect greater returns from interaction due to differences in human, physical or financial capital. A population is (ex ante) homogenous if all individuals have the same type, i.e. $\theta_i = \theta$ for all i , and heterogeneous otherwise. Section 3 studies homogenous populations; Section 4 takes up the case of heterogeneous populations.

2.5. Timing

The timing of the game is as follows. Each agent decides, when a period begins, whether or not to invest in protection and how much to interact during that period. During period t , unprotected agents become infected according to Eq. (3). An agent who becomes infected during

period t 'dies' when the period ends. The game ends after period 2, where all agents die.

2.6. Strategies and payoffs

A strategy for individual i determines protection and interaction in each period before death. We assume that agents maximize expected lifetime returns from interaction, net of protection costs.⁹

2.7. 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.

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 existence of a unique equilibrium, and compare its properties with those obtained in 'pure interaction' models on the one hand and 'pure protection' models on the other. We explore the effects of varying the contagiousness of the disease, and the returns from interaction.

We begin with some preliminary observations on the framework defined earlier allowing us to simplify the formal analysis of our model. Since all agents die when the final period ends, they need not be concerned about infection during that period. This yields $v_{i2} = 0$ and $k_{i2} = \bar{k}$, for all i , in any equilibrium. An unprotected individual i 's expected payoff may thus be simplified to

$$\Pi_i(k_{i1}, v_{i1} = 0) = r_i(k_{i1}) + (1 - p_{i1})r_i(\bar{k}), \quad (6)$$

with p_{i1} given by Eq. (3): interacting according to k_{i1} in period 1 confers returns $r_i(k_{i1})$ during that period, but also reduces chances of a payoff in the period which follows.

Observe similarly that a protected individual i 's expected payoff must in any equilibrium be given by

$$\Pi_i(v_{i1} = 1) = 2r_i(\bar{k}) - \gamma. \quad (7)$$

Eqs. (6) and (7) allow us to reduce the initial setup to an equivalent but simpler one-shot game, and to focus exclusively on first-period variables. Henceforth time subscripts will be dropped, under the convention that a variable without a time subscript always refers to the first period. We will say that an equilibrium is interior if the protection rate V lies in the open interval $(0, 1)$.

Our first result establishes existence and uniqueness of an equilibrium in a homogenous population.

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

All proofs are presented in the on-line appendix. We briefly describe the ways in which interaction and protection shape the payoffs of a susceptible. The negative externality exerted by unprotected individuals is governed by the incidence of the disease q —the probability that an individual with whom an individual interacts happens to be infected.

⁷ We have opted for a simple formulation of protection: the costs are invariant with respect to levels 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.

⁸ We assume that protection is perfect, i.e. an individual who invests in protection never contracts the disease. We note that (i) this assumption is justified if the probability of becoming infected is sufficiently reduced by protection and (ii) our results are unchanged if there is a small probability of getting the disease from the vaccine itself.

⁹ A simple way to account for risk aversion would involve subtracting $cp_{i1}(1 - p_{i1})$ to the RHS of Eq. (6) giving a susceptible's expected payoff, c representing the cost of uncertainty for a susceptible. 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 this point.

The incidence of the disease can be written as the product of two terms: (i) the probability z of meeting another susceptible and (ii) the probability p that a susceptible is in fact infected.

Since in a homogenous population the incidence faced by all individuals is the same, and payoffs are strictly concave in interaction, then all susceptibles must in any equilibrium choose the same interaction. Let here k_u denote the interaction of unprotected individuals in a given equilibrium. Since $\bar{k}$ maximizes protected individuals' payoffs then, using (1) where A now refers to the subset of susceptibles, the probability z of meeting an unprotected individual is

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

Substituting $q = zp$ in Eq. (3) and solving for p yields

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

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

$$k_u = \arg\max_k \left[r(k) + (1-\epsilon)(1-\alpha z p k) r(\bar{k}) \right] \quad (10)$$

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

The proof of the Proposition has 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 fixed, existence and uniqueness of equilibrium in the overall game of protection and interaction follows.

The dual response to infection (restricted interaction and protection) plays a central role in our model. To illustrate its role, we now compare our equilibrium (a) with equilibrium in the 'pure protection' model with interaction fixed exogenously at the no-disease optimum $\bar{k}$ (benchmark 1) and (b) with equilibrium in the 'pure interaction' model with no protection (benchmark 2). 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 study of epidemiology is that a part of the benefits from protection comes from the reduced risk of infecting others, which 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 'pure protection' benchmark. 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.

Proposition 2. Consider a homogenous population. Let (V^*, I^*) and (V^1, I^1) denote the equilibrium protection and infection rates in our model and benchmark model 1, respectively. Then:

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 the returns from protection upwards, 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.

Consider next the implications of increasing returns from interaction. In what follows $V(\theta)$ denotes the equilibrium protection rate in a homogeneous population with type θ .

Proposition 3. Suppose $V(\theta) \in (0, 1)$. An increase in the returns θ induces a first-order stochastic dominance shift in the equilibrium profile of interaction and raises equilibrium protection. Moreover, depending on the function g , increasing θ may reduce equilibrium infection rates.

The intuition behind the first part of Proposition 3 is straightforward. Higher returns from interaction enhance the appeal of protection relative to the option of reducing interaction. In equilibrium, more agents protect. This lowers the incidence q of the disease, allowing in turn susceptibles to interact more.

The second part of the Proposition is perhaps most interesting. Increased protection tends to lower infection; yet higher interaction from the susceptible part of the population pushes infection up. This raises the question: how is infection globally affected by increased returns from interaction? The key issue is the elasticity of interaction: if this elasticity is not too high, then the effect of higher protection will prevail. We show in the proof, using an example with logarithmic returns, that infection rates may fall with a rise in θ .

The latter observations are 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 interaction, the greater the negative equilibrium externality each agent imposes on others. In the 'pure interaction' benchmark, therefore, an increase in θ unambiguously raises the infection rate (Kremer, 1996). As shown in the discussion above, in our model the effects of increasing returns θ are very different. The more individuals value interaction, the less inclined they are to respond to an epidemic by reducing interaction: this pushes protection up and leads to falling infection rates for a broad class of applications.

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. Social interaction among the unprotected stays constant, but because more individuals protect (and protected individuals interact more) the profile of interaction undergoes a first-order stochastic dominance shift. Hence, somewhat surprisingly, social interaction is non-monotonic in contagiousness: it first decreases but eventually rises.

Proposition 4. Consider a homogenous population. Equilibrium protection V increases with contagiousness α . Moreover, there exist $\underline{\alpha}$ and $\bar{\alpha}$, $\underline{\alpha} \leq \bar{\alpha}$, such that¹⁰:

1. If $\alpha \leq \underline{\alpha}$ then a reduction in α induces a first-order stochastic dominance shift in the equilibrium profile of interaction. Infection rate changes depend on the elasticity of interaction with respect to α .

¹⁰ Lemma 3 in the online appendix shows that both threshold values are decreasing in ϵ , the exogenous probability of infection.

2. If $\underline{\alpha} < \alpha < \bar{\alpha}$ then an increase in α induces a first-order stochastic dominance shift in the equilibrium profile of interaction, and a decrease of the infection rate I .
3. If $\alpha \geq \bar{\alpha}$ then all agents protect, and $I = 0$.

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 the equilibrium interaction of unprotected individuals, k_u , decreases in α and completes the arguments for part 1.

Consider next an interior equilibrium: an individual must be indifferent between protecting or not. Since payoffs to those with protection (viz. (7)) are unaffected by α , the payoffs earned by susceptibles (viz. (6)) must be unchanged as well. This implies that the equilibrium externality $\alpha q(\alpha)$ is constant in α . As α is rising, the incidence $q(\cdot)$ must be falling. Now consider protection. Suppose that protection were 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 the equilibrium k_u is constant in α , and the fraction of protected individuals rises, an increase in contagiousness α thus induces a first-order stochastic dominance shift in the equilibrium profile of interaction. This completes the argument for part 2.

We conclude with a brief 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 of Proposition 4 follows similar logic. The contrast is, however, sharp if we look at infection rates in an interior equilibrium (i.e. $\alpha \in (\underline{\alpha}, \bar{\alpha})$): 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. But $I = (1 - V)p$, so the infection rate I falls, since equilibrium V is rising in α .

Fig. 1 illustrates these points. Contagiousness α is on the x-axis. We measure on the y-axis the equilibrium average interaction in the population, the equilibrium vaccination rate V , and the equilibrium interaction of unprotected agents k_u . The parameter $\bar{k}$ is the optimal interaction without contagion.




Fig. 1. Responses to Contagiousness.

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.¹¹ The probability that a new meeting of an individual in community J is with an individual in subset $A \subset J'$ is thus

$$\mathbb{P}(i \in J \text{ selects contact in } A \subset J') = \left((1 - \beta)I_{\{J=J'\}} + \beta I_{\{J \neq J'\}} \right) \frac{\int_{i \in A} k_i di}{\int_{i \in J'} k_i di},$$

where as usual $I_{\{J\}}$ 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. $V(\theta_j) \in (0, 1)$ and $J \in \{H, L\}$.¹²

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 . So 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 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 such that equilibrium is interior in both isolated communities, so that $V(\theta_j) \in (0, 1)$, $J \in \{H, L\}$. Then an equilibrium exists for all levels of integration β , and is unique over a range $[0, \bar{\beta}]$, where $\bar{\beta} \in (0, 1/2)$. Moreover, over this range, increasing β induces (i) higher (lower) protection in community H (in community L), (ii) lower (higher) infection in community H (in community L) and (iii) a (inverse) first-order stochastic dominance

¹¹ 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 $J_1 \cup J_2$ 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.

¹² 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.

shift in the equilibrium profile of interaction of community H (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. Equilibrium payoffs of unprotected individuals therefore remain unchanged as β varies, given that payoffs of protected individuals are constant. Since payoffs of unprotected individuals 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 first-order stochastic dominance shift in the equilibrium profile of interaction of community H . Similar reasoning establishes the claim for community L .

There are interesting implications of Proposition 5. To illustrate, suppose that prior to the stage at which individuals make a decision a social planner is able to allocate protection 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 protection 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. Greater integration typically implies multiple equilibria. To bring out this point in the simplest way, we next consider an example where $\beta = 1$ and focus on a range of parameters where isolated communities exhibit interior equilibria. The case $\beta = 1$ corresponds, e.g., to a situation where all sex is heterosexual. We show in the example that three equilibria exist. 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 H types protect and L types do not. Extremal equilibria are attractive since they maximize meetings between protected and unprotected individuals. Of the two extremal equilibria, that in which high types protect dominates due to the fact that the loss from reduced interaction is worse for high types than low.

Proposition 6. Consider a heterogeneous population, $\beta = 1$, and $q(\theta_j) \in (0, \varepsilon)$, $j \in \{H, L\}$. There exist exactly three equilibria. In equilibrium 1 all individuals in groups H protect and all individuals in group L remain unprotected. In equilibrium 2 all individuals in group L protect and all individuals in group H remain unprotected. Equilibrium 3 is interior in both groups. If W_k denotes average social welfare in these equilibria then $W_1 > W_2 > W_3$.

5. Conclusion

Social and economic interactions generate rewards but also facilitate the spread of infections. 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 an infection may be contracted either exogenously 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. 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.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <http://dx.doi.org/10.1016/j.jpubeco.2015.02.010>.

References

- Alpcan, T., Basar, T., 2011. *Network Security: a Decision and Game Theoretic Approach*. Cambridge University Press, Cambridge, England.
- Anderson, R., 2008. *Security Engineering*. Second ed. Wiley.
- Anderson, R.M., May, R.M., 1991. *Infectious Diseases of Humans: Dynamics and Control*. OUP, Oxford.
- Anderson, R., Bohme, R., Clayton, R., Moore, T., 2007. *Security Economics and the Internal Market*. European Network and Information Security Agency.
- Aspnes, J., Chang, K., Yampolskiy, A., 2006. Inoculation strategies for victims of viruses and the sum-of-squares partition problem. *J. Comput. Syst. Sci.* 72 (6), 1077–1093.
- Brito, D.L., Sheshinski, E., Intriligator, M.D., 1991. Externalities and compulsory vaccinations. *J. Public Econ.* 45, 69–90.
- Chen, F., and F. Toxvaerd (2014), *The economics of vaccination*, Mimeo, Cambridge University.
- Detica and Cabinet Office, 2011. *The Cost of Cyber Crime*.
- Francis, P.J., 1997. Dynamic epidemiology and the market for vaccination. *J. Public Econ.* 63 (3), 383–406.
- Galeotti, A., Rogers, B., 2013. Strategic immunization and group structure. *Am. Econ. J. Microecon.* 5 (2), 1–32.
- Geoffard, P.-Y., Philippon, T., 1996. Rational epidemics and their public control. *Int. Econ. Rev.* 37 (3), 603–624.
- Geoffard, P.-Y., Philippon, T., 1997. Disease eradication: private versus public vaccination. *Am. Econ. Rev.* 87 (1), 222–230.
- Gersovitz, M., 2011. The economics of infection control. *Ann. Rev. Resour. Econ.* 3, 277–296.
- Gersovitz, M., Hammer, J.S., 2004. The economical control of infectious diseases. *Econ. J.* 114, 1–27.
- Goldman, S.M., Lightwood, J., 2002. Cost optimization in the SIS model of infectious disease with treatment. *Top. Econ. Anal. Policy* 2 (1), 1–22.
- Kremer, M., 1996. Integrating Behavioral Choice into Epidemiological Models of AIDS. *Q. J. Econ.* 111 (2), 549–573.
- Moore, T., Clayton, R., Anderson, R., 2011. *Economics and internet security: a survey of recent analytical, empirical and behavioral research*. Harvard University Computer Science Department, Discussion Paper TR-03-11.
- Philippon, T., Posner, R., 1993. *Private Choices and Public Health: The AIDS Epidemic in an Economic Perspective*. Harvard University Press.
- Wellings, K., Field, J., Johnson, A.M., Wadsworth, J., Bradshaw, S., 1994. *Sexual Behaviour in Britain: The National Survey of Sexual Attitudes and Lifestyles*. Penguin Books.
- World Health Organization, 2013. *HIV/AIDS. Fact Sheet Number 2013*.
- World Health Organization, 2014a. *Influenza (seasonal). Fact Sheet Number 211*.
- World Health Organization, 2014b. *Tuberculosis. Fact Sheet Number 104*.