

Strategic Immunization and Group Structure

Andrea Galeotti and Brian W. Rogers*

March 28, 2013

Abstract: We consider the spread of a harmful state through a population divided into two groups. Interaction patterns capture the full spectrum of assortativity possibilities. We show that a central planner who aims for eradication optimally either divides equally the resources across groups, or concentrates entirely on one group, depending on whether there is positive or negative assortativity, respectively. We study a game in which agents can, at a cost, immunize. Negative assortative interactions generate highly asymmetric equilibrium outcomes between ex-ante identical groups. When groups have an underlying difference, even a small amount of inter-group contacts generates large asymmetries.

We study the diffusion of a harmful state through a population. Immunity is available, but is costly. The state is meant to capture various kinds of choices or risky behaviors such as, for example, tobacco use, in which case immunity is interpreted as a commitment to avoid the temptation of smoking. The state can also capture the presence of an electronic virus on a computer network; in this case immunity represents the purchase of anti-virus software, or other costly measures taken to avoid the virus. But perhaps the most conventional interpretation is that the state represents human infection of various communicable diseases

*Galeotti, University of Essex, Department of Economics, Wivenhoe Park, Colchester, Essex CO3SQ, United Kingdom, agaleo@essex.ac.uk; Rogers, Northwestern University, Kellogg School of Management, MEDS, and Northwestern Institute on Complex Systems, 2001 Sheridan Road, Evanston, IL, 60208, USA, b-rogers@northwestern.edu. We thank Ben Golub, Sanjeev Goyal, Matt Elliott, and Adrien Vigier for helpful comments. Galeotti is grateful to the European Research Council for support through ERC-starting grant (award no. 283454) and to the UK Economic and Social Research Council for support through the MiSoC research centre (award no. RES-518-285- 001). The authors state that they have no relevant or material financial interests that relate to the research described in this paper.

that spread through social contacts; in this case immunity captures a decision to vaccinate oneself.

We model the diffusion process using a standard Susceptible-Infected-Susceptible (SIS) framework, developed originally for the kind of epidemiological application mentioned above. See Bailey (1975) and Anderson & May (1992) for classical treatments. More recently, the SIS model and its variants have been proposed as useful in understanding a wide variety of processes including such diverse applications as behaviors, information diffusion, learning dynamics, myopic best response, and imitation dynamics.¹

The diffusion process works as follows. Individuals exist in one of two states, S and I , and transition between these states over time. Transitions from I to S occur at a given recovery rate, but transitions from S to I occur at an endogenous rate that depends, in particular, on the expected number of individuals in state I that will be met in a given period. We take a long-run perspective, and base our utility and welfare calculations on the (unique nondegenerate) steady-state of this process, which determines, most importantly, an expected proportion of time that a given individual will be infected with the harmful state.

There are two new features of our framework that enrich the canonical SIS model, and which drive the bulk of our analysis. *First*, we want to study the importance of group structure, and of group-dependent contact patterns, on diffusion outcomes. We therefore model the population structure and contact patterns in a sufficiently detailed manner so as to depart from the restriction that interactions occur through uniform random mixing in a homogenous population. We assume instead that society is composed of two groups, which we label A and B , and we parameterize the degree of positive or negative assortativity in the interactions between the groups.² At one extreme, individuals have contacts only within their respective groups, and so the diffusion of the harmful state within each group is independent of the diffusion in the other group. At the other extreme, all contacts involve A - B pairs, resulting in a bipartite interaction structure. In this case, the diffusion of the harmful state within one group depends entirely on the prevalence in the other group. The range of cases in between capture varying extents of assortativity, and in each of these cases the diffusion process is linked

¹See, for example, Jackson & Rogers (2007), López-Pintado (2008) and Jackson & Yariv (2010), as well as references therein.

²Jackson & Lopez-Pintado (2011) consider a setting with positive assortative interactions but there are no incentives in their model. Montopoli, Bhattacharyya & Bauch (2009) studies numerically a model with vaccination in a two-country setting. Our paper is related more closely to the work of Reluga (2009) which also considers strategic vaccination in a two-group setting. The focus in Reluga (2009) is on the stability of equilibria, whereas our analysis explores different economic questions about eradication policies and comparative statics.

across the two groups. The intermediate case, in particular, involves interaction patterns that are completely blind to group labels.

Most applications of interest involve some degree of positive assortative matching. In social networks, the main driving force behind positively assortative interactions is homophily, whereby individuals exhibit a tendency to interact with others who are similar to themselves.³ Due to the robust nature of homophilic biases, positive assortative interactions have been shown to pervade social networks on many dimensions, including economic status, race, and political preferences.⁴

Notwithstanding the prevalence of positively assortative networks, there are important cases in which negative assortative interactions are prevalent. This includes the application to sexual contact networks, in which the predominance of interactions are male-female. In the realm of communicable diseases, sexually transmitted diseases are among the most important to understand. More generally, negative assortative interactions are typical of settings in which agents match for reasons relating to complementarities of skill or knowledge. A particular application is to a financial market consisting of buyers and sellers who seek each other out for profitable exchange. The diffusion process captures the possibility that rumors of a bad aggregate state will spread through the market, thereby preventing surplus from being realized. Here, immunity would be interpreted as taking the costly action of becoming informed about the true aggregate state, so as not to be influenced by the spread of a pessimistic rumor.

The *second* aspect is that, while diffusion processes are often studied in a purely stochastic framework, we explicitly model the incentives of agents who are exposed to the diffusion process. In the last two decades some progress has been made towards incorporating certain aspects of incentives into epidemiological diffusion models.⁵ This research has identified important effects of modeling agents' incentives, demonstrating crucial differences with the earlier models that are purely mechanical in nature. Our paper contributes to this line of work by modeling the immunization decision of individuals in a strategic setting. In particular, we assume that all agents simultaneously decide whether or not to undertake a costly immuniza-

³In fact, from a network analysis perspective, it may be more natural to use the term homophily in place of positive assortativity. We use the latter because it is easier to unify with the complementary case of negatively assortative interactions, which are also important in our analysis.

⁴There is a large literature on homophily in social networks. Important early work is contained in Marsden (1987) and Marsden (1988). Examples of more recent work include McPherson, Smith-Lovin & Cook (2001), Moody (2001), and Echenique, Fryer & Kaufman (2006).

⁵The literature in economics dates at least to Brito, Sheshinski & Intriligator (1991). Important contributions include Kremer (1996), Geoffard & Philipson (1996), Geoffard & Philipson (1997), Bauch & Earn (2004), Boulier, Datta & Goldfarb (2007), Toxvaerd (2010a) and Toxvaerd (2010b).

tion.⁶ The cost of the immunity must be weighed against the expected cost of exposure to the harmful state during the diffusion process. As such, individuals anticipate the consequences that their collective immunization decisions will have on the outcome of the diffusion process, and decide optimally under these beliefs.

A game theoretic analysis of immunization in which assortativity affects contact patterns is novel and, as we shall see, provides interesting insights into behavior and welfare. In fact, our main results come specifically from the combination of these two factors. We now summarize these findings.

Our first set of results concerns the problem of a central planner attempting to eradicate the harmful state (Section II).⁷ The planner might be a governmental agency deciding how many vaccines to produce for a communicable disease and how to allocate them across the population. Alternatively, the social planner could be a governmental agency that aims to eliminate smoking in schools via an educational program and it has to decide how large the program should be and which students it should target or require to participate. How should the planner's policy depend on the structure of interactions among potential smokers? What information does the government need in order to determine the optimal size of the program (level of immunity) and how to efficiently distribute the resources to the population?

We show that there is a sharp discontinuity in the optimal eradication policy when we move from positive assortative interactions to negative assortative interactions (Proposition 2). In the former case, the optimal allocation involves splitting the resources equally between the two groups and the minimum budget needed to eradicate the harmful state is independent of the exact degree of positive assortativity. In contrast, when there are negative assortative interactions, the optimal allocation is to first focus resources exclusively on one group until it is fully immunized and then immunize part of the other group, if required. Furthermore, under negative assortative interactions, the minimum budget necessary to achieve eradication is decreasing in the intensity of assortativity, reaching its lowest level when the network of interactions is bipartite.

The basic intuition behind this result is the following. Suppose group B receives more immunizations than group A, and the planner has an extra immunization to allocate. Consider

⁶In a complementary paper, Goyal & Vigier (2010) studies the dual roles of vaccination and level of interaction in a game theoretic framework.

⁷Galeotti & Goyal (2009) also study optimal policies in the presence of social effects. While Galeotti & Goyal (2009) model social effects in a reduced form, here we explicitly define the process of diffusion. Another difference is that we focus on the effect that assortativity has in the design of optimal policies, an aspect that is not considered by Galeotti & Goyal (2009).

first the case of positive assortativity. There are two effects to account for. The direct benefit of an additional immunization depends on the infection rate among non-immunized agents, which is greater in group A (Proposition 1). The indirect benefit of an additional immunization is tied to how many of the immunized agent's contacts it protects. Viewed in this way, when interactions are positively assortative, immunizing one group at a higher level than the other is inefficient, since within the high-immunity group, more relationships are already immunized, and thus an immunization in the low-immunity group protects more relationships on the margin. The direct and indirect effects are aligned, and therefore the optimal policy immunizes both groups at the same rate. In the case of negative assortativity, both of these effects are reversed and so the optimal policy focuses on only one group.

An important implication of this result is that in contexts where the harmful state diffuses through social networks that exhibit positive assortativity, it is not a valuable use of resources for a central planner to learn the exact degree of assortativity, as the optimal plan is not sensitive to these details. In contrast, if the harmful state diffuses through social networks that exhibit negative assortativity such as, for example, sexual contact networks, there is value to the social planner in learning the degree of assortativity, as the budget required for eradication depends on that information. In particular, the planner can save on immunity resources the greater is the intensity of assortativity, saving the most when the social network is bipartite.

Our second set of results explore equilibrium behavior (Section III). We posit that an agent can become immune to the harmful state at a cost, which has to be paid at the beginning of the game. The strategic nature of individual immunization decisions in a context in which assortativity affects contact patterns gives rise to the possibility of asymmetric outcomes in which one group immunizes at a higher level than the other, even when the groups are ex-ante identical. We identify the contexts in which such an equilibrium is plausible, along with the associated implications for prevalence and welfare in the two groups.

We show that when there are positive assortative interactions, ex-ante identical groups necessarily obtain symmetric equilibrium outcomes: the immunization rate is the same across groups, with the consequence that the prevalence of the harmful state and welfare are the same across groups as well. Moreover, these outcomes are invariant to the specific level of positive assortativity. In contrast, when there are negative assortative interactions, the equilibrium is highly asymmetric. It is either the case that one group does not immunize at all, or that the other group immunizes completely, or possibly both (Proposition 4). The prevalence of the harmful state is higher in the group with the low immunization rate than in the group with

the high immunization rate, yet the group with the low immunization rate has the highest welfare (Proposition 5). Hence, higher infection prevalence is not indicative of lower welfare. Furthermore, the difference between the welfare of the two groups is higher when the degree of negative assortativity is higher, thereby reaching a maximum in a bipartite social network. We also show that the prevalence of each group may be nonmonotonic with the degree of assortativity (Proposition 6).

Finally, we consider the possibility that there is an underlying economic asymmetry between the groups (Section IV). We capture this asymmetry by imposing different immunization costs in the two groups.⁸ This would be the case if, for example, one group is wealthier than the other, so that its cost of immunization is lower relative to the marginal benefit of being in state S over state I . In this context we show that the introduction of small initial differences across groups is amplified by a social multiplier, and this social multiplier is increasing in cross-group interaction, even when there are positive assortative interactions (Proposition 7 and Corollary 1).

Broadly, our analysis shows how the details of group interactions have a profound effect both on how a planner should respond to reducing the prevalence of the state and on the equilibrium outcomes of the diffusion process. Group interactions tend to generate large asymmetries across groups in the outcomes of interest, namely welfare and prevalence, even when groups are, ex-ante, very similar. It is therefore important to account for the structure of interactions when studying, empirically, the sources of observed asymmetries in behavior and diffusion outcomes across groups. Failing to account for the impact of inter-group interactions may lead to overestimating the underlying differences across groups. Moreover, when interactions are negatively assortative, we demonstrate that the planner will generally have to resolve a tension between the optimal reduction of prevalence and reducing asymmetries in welfare across the groups, whenever fairness is a consideration.

I A model of social diffusion

There is a continuum population. Individuals are classified into two groups each of mass one, labeled A and B . The groups serve purely as arbitrary labels; from a modelling perspective, individuals in the two groups are identical in every respect.⁹ At each point in time, every

⁸Alternative forms of asymmetries, such as different infection rates across groups or different group size are discussed in Section IV and Section V.

⁹We consider asymmetries between groups in Section IV.

individual interacts with k other individuals. A proportion $\beta \in [0, 1]$ of the interactions are with individuals in the same group while the remaining interactions are with individuals in the other group. Thus β measures the extent of positive or negative assortativity of interactions. At one extreme $\beta = 1$, there are no interactions across the groups, and society is segregated into two separate networks. When $\beta > 1/2$, there is positive assortative interaction, in the sense that the majority of contacts involve intra-group interactions. At $\beta = 1/2$, the interaction structure is independent of group labels. For $\beta < 1/2$, most interactions involve individuals from opposite groups and we say that there is negative assortative interaction; at $\beta = 0$, the network becomes bipartite, so that all interactions are across groups.

We model the diffusion of a harmful state through this population in continuous time. For the moment, assume that a proportion π_g of group $g \in \{A, B\}$ has immunity.¹⁰ At each point in time, each individual who is not immunized exists in one of two states, which we call Susceptible (S) and Infected (I). The states of non-immunized individuals change over time. The transitions between states S and I depend on (i) the patterns of interactions between individuals, (ii) the proportion of individuals with immunity, (iii) the rate of infection, and (iv) the recovery rate.

Formally, let ρ_g be the proportion of the non-immunized group g individuals who are in state I . Define θ_g as the probability that a randomly chosen contact of a group g individual is infected. By definition, we have

$$(1) \quad \begin{aligned} \theta_A &= \beta(1 - \pi_A)\rho_A + (1 - \beta)(1 - \pi_B)\rho_B \\ \theta_B &= \beta(1 - \pi_B)\rho_B + (1 - \beta)(1 - \pi_A)\rho_A. \end{aligned}$$

A group g individual in state S becomes infected at a rate $\nu\theta_g k$, where ν is the per-contact rate of infection. When ν is small, which is typically the relevant case, this quantity approximates the chance of becoming infected through independent interactions with each of the k neighbors. Individuals in state I transition to state S at a rate δ . Define $\tilde{\lambda} = \nu/\delta$ as the ratio of the rate of infection to the rate of recovery. Steady-state outcomes are characterized by a stable prevalence in the two groups, where the transition rates between S and I are balanced. We have

$$\frac{\partial \rho_g}{\partial t} = (1 - \rho_g)\nu\theta_g k - \rho_g\delta = 0,$$

¹⁰We assume throughout the paper that immunity is perfect: an agent with immunity cannot be infected by the harmful state. The analysis is easily extended, without additional insights, to the case where immunity is imperfect.

for $g = A, B$. The first term is the in-flow from individuals in state S : the fraction of non-immunized players who are not infected, $1 - \rho_g$, times the rate of infection, $\nu\theta_g k$. The second term is the out-flow: each non-immunized infected agent recovers at a rate δ . Now define $\lambda = k\tilde{\lambda}$ as the intensity of infection. As will be clear below, the effect of k , ν and δ is summarized by the single parameter λ in the analysis. Solving the condition above we obtain:

$$(2) \quad \rho_g = \frac{\lambda\theta_g}{\lambda\theta_g + 1}.$$

We define the prevalence of group g as the proportion of agents in group g that have adopted the harmful state, i.e., $P_g = (1 - \pi_g)\rho_g$, with total prevalence given by $P = P_A + P_B$.

Using standard arguments, one can show the following:

Observation 1 For any given intensity of infection $\lambda > 0$, degree of assortative interaction $\beta < 1$ and a pair of immunization levels (π_A, π_B) , there exists either (i) a unique steady-state in which no agent is infected, i.e., $\rho_g = 0$, $g \in \{A, B\}$, or (ii) two steady-states, one in which no agent is infected and the other with a positive, though potentially different, level of infection in both groups. In case (ii) the unique steady-state with positive level of infection is globally stable.¹¹

Since the groups are inherently identical, when dealing with asymmetric outcomes, we focus, without loss of generality, on the case $\pi_A < \pi_B$.

The extent to which individuals in the two groups interact puts a particular structure on the way that their steady state outcomes relate to one another. These relationships depend crucially on whether matchings are positively or negatively assortative, as we now demonstrate.

Proposition 1 *In the steady state with positive level of infection the following hold:*

1. *Suppose $\beta < 1/2$. $\pi_A < \pi_B$ iff $\rho_A < \rho_B$ iff $\theta_A < \theta_B$ iff $P_A > P_B$.*
2. *Suppose $\beta > 1/2$. $\pi_A < \pi_B$ iff $\rho_A > \rho_B$ iff $\theta_A > \theta_B$ iff $P_A > P_B$.*

Regardless of the degree of assortativity, the group with the lower immunization rate experiences higher prevalence. When interactions are negatively assortative, this same group experiences a *lower* level of infection among the non-immunized agents and a lower proportion of infected neighbors. On the other hand, when interactions are positively assortative, the

¹¹This result can be shown through a straightforward modification of standard results in the SIS framework as can be found, for example in Vega-Redondo (2007).

group with the lower immunization rate experiences a higher level of infection among the non-immunized and a higher proportion of infected neighbors.

In what follows, we first study the problem of a planner who allocates a budget of immunizations with the aim of eradicating the harmful state from the population (Section II). We then turn to the equilibrium analysis (Sections III and IV). There, agents choose whether to become immune at a cost, or to be exposed to the diffusion process. In both cases, we assume that agents are concerned with the steady state outcome, and that the outcome of the diffusion process is the unique positive-valued steady state prevalence, when it exists. Notice that Proposition 1 applies to any pair of immunization rates (π_A, π_B) . We will make use of it in the analysis when considering optimal immunization plans as well as equilibrium outcomes.

II Eradication of Infection

We begin our analysis by considering the problem of a central planner whose objective is to eradicate the harmful state. We ask two questions. First, what is the smallest amount of immunization the social planner requires to eradicate the harmful state? Second, how should the social planner allocate that immunity across the population so as to achieve eradication? We solve for the optimal immunization plan subject to the parameter that describes the transmission rate of the undesirable state, λ , as well as the structure of interactions, as captured by β .¹²

The planner may represent a government deciding how many vaccines for a communicable disease to produce and how to allocate them across distinct groups. Alternatively, we may think that a governmental agency has some resources to use towards reduce smoking in schools via an educational program and it has to decide how large the educational program should be and which students should attend the program. Under reasonable conditions on the production cost of immunizations (e.g., constant marginal cost), the objective to eradicate the undesirable state is equivalent to the objective to maximize social welfare net of immunization costs.¹³

¹²There is a literature on the difficulties of inducing discretionary immunization due to the effect of strategic substitutes. See, for example, Geoffard & Philipson (1997), Barrett (2003), Bauch & Earn (2004). As our goal is specifically to study the impact of interaction patterns across populations, we focus on the case in which the planner can directly implement the desired outcome.

¹³In general, welfare maximization may involve non-trivial prevalence. This would be true, e.g., if the cost function is sufficiently convex. The planner's optimal immunization budget generally depends both on the cost function and on the details of the matching process. Given any immunization budget for the planner, such as the socially optimal one, we are interested in the optimal way to allocate that budget so as to minimize

Interestingly, the conclusion that a social planner would find full eradication to be optimal, if indeed costs are small enough that it prefers some action to taking no action, appears in other analyses, including Rowthorn (2006) and Barrett & Hoel (2007).

What information does the government need to efficiently decide the optimal policy to eradicate the infection? Should immunizations be allocated symmetrically or concentrated on one group? How should the planner's actions depend on the structure of interactions? We find that the answer to these questions depends crucially on whether the interaction patterns are positively or negatively assortative.

Proposition 2 *For every (λ, β) there exists a threshold $\bar{\pi}^*$ such that eradication is possible if and only if the mass of immunizations available to the social planner is $\bar{\pi} \geq \bar{\pi}^*$.*

When there are positive assortative interactions, the threshold $\bar{\pi}^$ is independent of β and eradication is obtained by dividing $\bar{\pi}^*$ equally across groups, i.e. $\pi_A = \pi_B = \bar{\pi}^*/2$.*

When there are negative assortative interactions, the threshold $\bar{\pi}^$ is strictly increasing in β and eradication is obtained by concentrating $\bar{\pi}^*$ as much as possible on one group, i.e., $\pi_B = \min\{1, \bar{\pi}^*\}$ and $\pi_A = \max\{0, \bar{\pi}^* - \pi_B\}$.*

The proof of Proposition 2 develops in two steps. In the first step, we fix the division of immunity across groups and determine, for this division, the smallest mass of immunization needed to eradicate the disease. In the second step, we determine the allocation that minimizes this threshold level of immunization needed to eradicate the disease.

The basic force behind this result relies in part on Proposition 1 and works as follows. First focus attention on the infection rate *among non-immunized agents*, since this is precisely what captures the direct benefit of an additional immunization. When interactions are positively assortative, the infection rate is higher in the group receiving less immunity, while when interactions are negatively assortative, the infection rate is higher in the group receiving *more* immunity. Next observe that, independent of the nature of assortativity, it is always the case that the group receiving the higher level of immunization obtains the lower overall prevalence. The indirect benefit of an additional immunization to a node in the group with lower immunity is summarized by the following.¹⁴ When interactions are positively assortative, the agent interacts mostly within its own group where there is a higher proportion of susceptible agents,

total prevalence. The answer is similar to that provided in Proposition 2: when there are positive assortative interactions, total prevalence is minimized by dividing the immunizations equally across groups; when there are negative assortative interactions, total prevalence is minimized when immunizations are concentrated as much as possible on one group.

¹⁴The converse is true for a node in the group with higher immunity.

thereby generating a large externality, while, conversely, when interactions are negatively assortative, the agent interacts mostly with the other group where the proportion of susceptible agents is lower, generating a smaller externality. Summarizing, the direct and indirect effects are aligned in both cases, such that when interactions are positively assortative it is efficient, on the margin, to immunize an agent in the group with lower immunity, and when interactions are negatively assortative it is efficient to immunize an agent in the group with higher immunity.

Proposition 2 shows that there is a sharp discontinuity in the optimal eradication policy of the planner when we move from positive assortative interactions to negative assortative interactions. Figure 1(a) illustrates how the minimum level of immunization that the social planner needs to eradicate the undesirable state, denoted by $\bar{\pi}^*$, changes with the degree of assortative interaction. Figure 1(b) illustrates how the optimal allocation of $\bar{\pi}^*$ changes with the degree of assortative interaction.

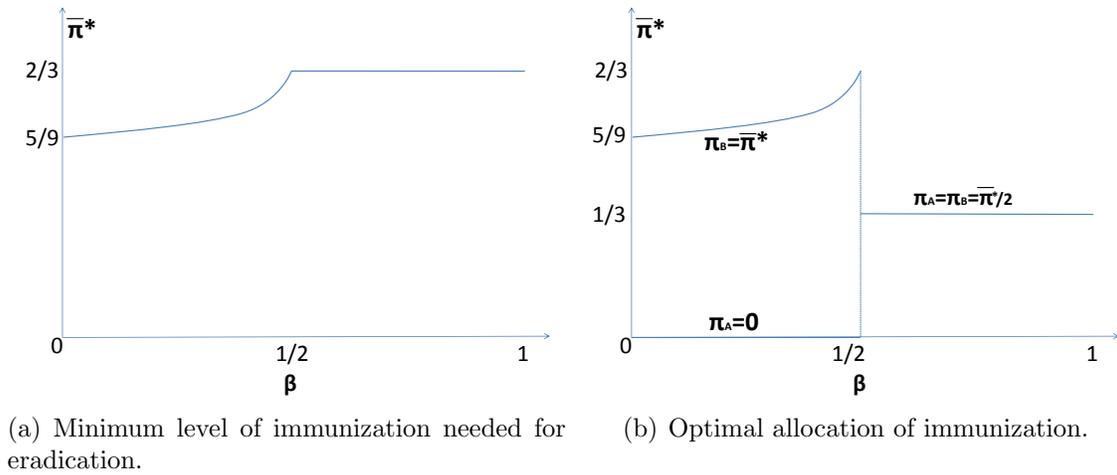


Figure 1: Optimal eradication policy, $\lambda = 1.5$, $\beta \in [0, 1]$.

When there are positive assortative interactions, the threshold $\bar{\pi}^*$ is always higher than in the case of negative assortative interactions and the planner eradicates the undesirable state by splitting equally the resources for immunization among the two groups. The minimum level of immunizations that the social planner needs for eradication is given by $\bar{\pi}^* = 2(1 - 1/\lambda)$ and it has a very natural interpretation: it is equal to the value that total prevalence would take in the absence of any immunity in the population. In this sense, one can think of achieving eradication under positive assortative interaction as being equivalent to providing immunity to every infected individual at a given moment in the steady-state of the diffusion process.

When there are negative assortative interactions the minimum level of immunization needed to eradicate the harmful state is decreasing in the intensity of negative assortativity, reaching therefore the minimum level when the network of interactions becomes bipartite.¹⁵ Moreover, eradication is obtained by allocating immunity exclusively to one group, and then allocating immunity to the other group, if necessary, only once the first group is fully immunized. Fixing the intensity of infection, an increase in inter-group interactions increases the externalities that one group exerts on the other and, by concentrating immunity to one group, the planner leverages these group-level externalities more effectively, reducing the requisite budget for eradication.

Proposition 2 illustrates when and why it is important to incorporate the details of the interaction structure into the design of eradication policies, even when there are no underlying differences between groups. Consider, for example, a social planner who implements an eradication policy ignoring group labels and assumes interactions result from uniform random matching, as in the classical treatments of SIS. In this case, to eradicate the infection, the social planner will allocate randomly to the population a mass of immunizations $\bar{\pi}^* = 2(1 - 1/\lambda)$. In view of Proposition 2, this implies that in contexts where the harmful state diffuses through a positively assortative network, there is no value for the social planner to learn the exact details of the interaction structure, i.e., the value of β . In contrast, if the harmful state diffuses through social networks that exhibit negative assortativity, as in sexual contact networks, there is a positive value for the social planner to learn the exact value of β , since this information allows the social planner to achieve the goal of eradication while using a smaller budget of (presumably costly) immunization. This benefit is greatest when the social network is bipartite. In fact, as β decreases from $1/2$ the benchmark policy of allocating randomly $\bar{\pi}^* = 2(1 - 1/\lambda)$ across groups produces an over-investment of resources which gets increasingly worse relative to the optimal eradication policy.

III Immunity in a decentralized game

We turn now to assessing the outcomes that might be expected in the absence of any action taken by a planner. Before the diffusion process starts, individuals may choose, at a cost, to immunize themselves. We model this decision as a simultaneous move game in which agents

¹⁵The influential work of Anderson & May (1984) provides a result that eradication can be easier to achieve in the presence of population heterogeneity. A related result is obtained in Goyal & Vigier (2010) under particular assumptions.

anticipate the effects of aggregate immunization on the diffusion outcomes. Formally, the action of agent i is $a_i \in \{0, 1\}$, where $a_i = 0$ is interpreted as no immunization. Immunization, when taken, is perfect.

Individuals receive a flow utility of one in state S and zero in state I . Individuals are infinitely patient, so that utility is measured by the limiting proportion of time in state S . Denote the cost of immunization by $c > 0$. An agent who immunizes receives utility $1 - c$, since she is never infected but must pay for the immunity. On the other hand, a group g agent who does not immunize saves the cost, but spends a proportion ρ_g of time in state I , obtaining utility $1 - \rho_g$. Importantly, the steady state values of ρ_g depend on the immunization decisions of individuals in both groups. Thus, immunization decisions are optimal if (and only if) $\pi_g > 0$ implies $\rho_g \geq c$ and $\pi_g < 1$ implies $\rho_g \leq c$. In particular, at an interior equilibrium it must be that

$$(3) \quad \rho_A = \rho_B = c.$$

Having introduced immunization cost, we will be concerned with welfare in the two groups. There are two costs in the model: the explicit cost of immunization and the opportunity cost of being infected. We define the welfare of group g as the expected utility of a randomly selected agent, i.e.,

$$W_g = 1 - \pi_g c - (1 - \pi_g) \rho_g.$$

Our interest is in understanding how the pattern of interactions affects the equilibrium immunization rates of both groups, and consequently their prevalence and welfare.

A Equilibrium characterization

We first observe that if the intensity of infection λ is sufficiently low, or the cost of immunization c is sufficiently high, then in equilibrium no player will immunize. We have:

Proposition 3 *If $\lambda(1 - c) < 1$ there is a unique equilibrium and it is such that no player immunizes.*

The idea behind Proposition 3 is easy to understand. The smaller is the intensity of infection, λ , the lower will be the steady-state prevalence, all else equal, and so the lower will be the incentive to immunize. On the other hand, given any level of prevalence, the incentive to immunize is decreasing in its cost. It is never optimal to immunize when $c > 1$, since

an individual can never gain more than the difference in utility between being susceptible and infected from immunity. Notice, though, that the threshold below which no individual immunizes is independent of β . In light of these observations, we henceforth restrict attention to the case in which immunization decisions are non-trivial.

Assumption 1 $\lambda(1 - c) > 1$.

We maintain Assumption 1 throughout the rest of the analysis. The following result summarizes the equilibrium characterization.

Proposition 4 *For every (λ, β, c) there is exactly one stable equilibrium. When there is positive assortative interaction the stable equilibrium is symmetric, interior, and constant in β . When there is negative assortative interaction the stable equilibrium is asymmetric and such that either $\pi_A = 0$ or $\pi_B = 1$ (possibly both).*

The Appendix contains a complete characterization of equilibria and their stability properties.^{16,17} We offer here a rough intuition that underlies this result. First, recall from the equilibrium condition (3) that, at an interior equilibrium, it must be that the proportion of non-immunized agents that are infected in the two groups is the same, i.e., $\rho_A = \rho_B$. The steady-state condition (2) then implies that for an agent in either group, the expected proportion of neighbors that is infected is the same, i.e., $\theta_A = \theta_B$. For these two conditions to hold, it is necessary that immunization rates are symmetric across groups, i.e., $\pi_A = \pi_B$. Hence, any interior equilibrium is necessarily symmetric.¹⁸ Once that is known, it is easy to explicitly solve for the equilibrium immunization rate, which gives $\pi = 1 - 1/(\lambda(1 - c))$. Notice that this rate is indeed interior under Assumption 1.

By this argument any additional equilibrium must be asymmetric and involve a boundary solution. When there are positive assortative interactions, it is impossible to satisfy the equilibrium conditions at two distinct values of $\pi_A < \pi_B$. In contrast, when there are negative assortative interactions, the fact that most interactions are inter-group implies that one group

¹⁶Under negative assortativity, the symmetric interior equilibrium exists but it is unstable. To understand why this is so, note that as β decreases, the degree of strategic substitutability between the groups' immunization rates increases. That is, the best response of each group to the other group's immunization rate becomes more sensitive. When $\beta = 1/2$, the slope of the reaction functions at the interior equilibrium is unity, and so it is at this point that the interior equilibrium switches from being stable to unstable.

¹⁷In perhaps the most closely related paper, Reluga (2009) obtains results on the existence and stability of equilibria in a similar setting.

¹⁸The knife-edge case of uniform mixing, $\beta = 1/2$, is the only exception, in which case there is a continuum of interior equilibria.

free-rides off the high immunization rate of the other group. This is rational because the high-immunity group interacts mostly with the low-immunity group, which has higher prevalence, justifying the cost of their immunizations. The low-immunity group, meanwhile, takes advantage of the fact that most of their interactions are with the high-immunity group which has low prevalence. The consequence is that the free-riding group, while infected at a higher rate, obtains higher welfare because their high prevalence is necessarily more than offset by the savings in immunization costs. That is, higher infection prevalence is not indicative of lower welfare; indeed, the opposite is true. We have:

Proposition 5 *Consider the case of negative assortative interactions. In the stable equilibrium, the group that immunizes less has both higher prevalence and higher welfare.*

When there are positive assortative interactions, both groups immunize at interior rates that are constant in β , whereas in the case of negative assortative interactions one of the groups is either fully immunized or not immunized at all. Hence, equilibrium group immunization rates are discontinuous around $\beta = 1/2$. Nevertheless, the resulting prevalence in each of the two groups, as well as their welfare, are continuous in β . Furthermore, when there are positive assortative interactions, all interesting outcomes like prevalence and welfare are constant in β . However, this neutrality result in the degree of assortativity on equilibrium outcomes no longer holds when there are negative assortative interactions. In fact, in this case the effect of β on group prevalence, immunization rate and welfare is subtle, and cannot always be signed. We turn now to exploring these comparative statics.

Figure 2 illustrates features of the comparative statics. In the figure, the cost of immunization c is constant throughout, whereas the intensity of infection λ takes values that increase from the left column to the right column. The higher is $\lambda(1 - c)$ the higher is the incentive to immunize, all else equal. So, when we move from the left column to the right column we are considering environments in which, ex-ante, it is more profitable to invest in immunization. The three rows depict immunization rates (top), prevalence (middle) and welfare (bottom), where group A is shown in dashed blue, group B in dotted red, and the total population in solid green.

As we have already discussed the equilibrium under positive assortative interactions, we focus here on negative assortativity. In this range, consider the extreme cases. When $\beta = 0$, it is always the case that group A completely free rides on group B, and that there remain some non-immunized agents in group B, so that $0 = \pi_A < \pi_B < 1$. This is true because when the network is bipartite, the infection will be eradicated even when there remain some

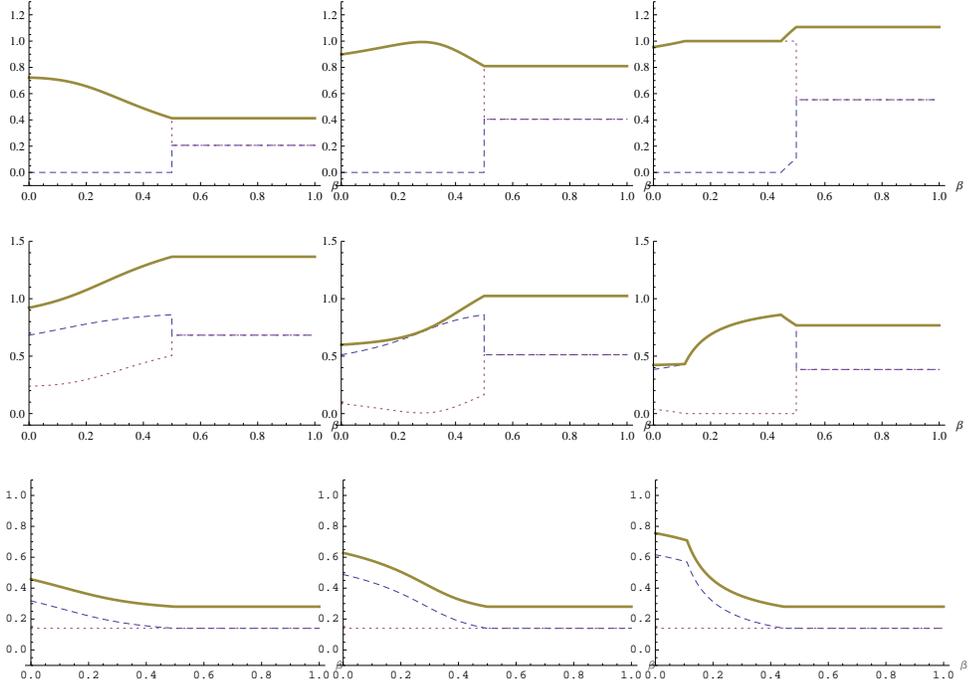


Figure 2: In all cases c is constant (at 0.86). The value of λ increases from the left column to the right column (taking values 9, 12, 16). Each row depicts different quantities with group A in dashed blue, group B in dotted red, and the total in solid green. Row 1: Immunization rates, Row 2: Prevalence, Row 3: Welfare.

non-immunized individuals in group B , due to herd immunity. Consider now values of β near $1/2$. When the incentive to immunize is sufficiently high, $\lambda(1-c) > 2$, the total immunization rate in society is greater than 50%, and so the equilibrium involves $\pi_B = 1$, as in the third column. Otherwise, the equilibrium involves $\pi_A = 0$, as in the first two columns.

These observations characterize which boundary conditions bind the equilibrium immunization rates near the extreme cases of $\beta = 0$ and $\beta = 1/2$ for negative assortative interactions. This is useful because, as we will see, the comparative statics of the equilibrium are very different in these cases.

Consider first the case where the incentive to immunize is high, so that $0 < \pi_A < \pi_B = 1$ near $\beta = 1/2$. Start from uniformly random interactions and consider decreasing β from $1/2$. Since group B is fully immunized, the free-riding of group A becomes more effective and this causes the immunization rate of group A to decrease, until nobody in group A immunizes. As this immunization rate drops, the prevalence in group A is increasing, due to the equilibrium

condition that ρ_A is constant at value c . But even as the prevalence in group A increases, so does their welfare.

Once we are in a region where group A completely free-rides ($\pi_A = 0$), prevalence in group A now changes directions and begins to decrease with a further decline in β , since the immunization rate in group A is now constant at zero. This fact gives rise to the possibility of a nonmonotonicity in prevalence, as illustrated in the middle-right panel of Figure 2. While there is an interval of β in which both boundary conditions bind ($\pi_A = 0$ and $\pi_B = 1$), at some point the prevalence in A becomes low enough that it is no longer optimal for group B to fully immunize. At this point $\pi_B < 1$, which holds until we reach $\beta = 0$.

Notice that, under parameters implying this pattern of interactions, total immunization is always lower than when $\beta > 1/2$ (first row, third column), while the total prevalence is nonmonotonic (second row, third column). Despite the nonmonotonicity in prevalence, it is always true that welfare strictly increases as β decreases, reaching its maximum value at the bipartite network.

Turning now to the remaining case when the stable equilibrium transitions to immunization levels with $0 = \pi_A < \pi_B < 1$ near $\beta = 1/2$, we see that the comparative statics are possibly nonmonotonic. The first and second columns of Figure 2 fall under these parameters. Indeed, in this case there are two opposing forces on the immunization decision of group B as β decreases. First, group B interacts more and more with group A which has higher prevalence, putting upward pressure on π_B . Second, as β decreases the prevalence in group A is decreasing, putting downward pressure on π_B . The net result of these two effects is subtle and depends on the particular parameters in question. In particular, the immunization rate and the prevalence of group B may well be nonmonotonic in β , as in the second column of Figure 2.

From the above discussion, we are able to conclude that when group A fully free-rides, their prevalence increases with β , whereas when group A only partially free-rides, the opposite is true. We derive an even simpler conclusion about welfare. For group B , the high-immunization group, welfare is always constant, because $\pi_B > 0$ implies that welfare is equal to $1 - c$. Finally, although it is much less obvious, we can also show that for group A , welfare is always decreasing in β . This is summarized in the following result.

Proposition 6 *Consider the stable equilibrium when there are negative assortative interactions.*

1. *For the group that immunizes more, welfare is constant in β .*

2. For the group that immunizes less, welfare is decreasing in β .

Furthermore, for the group that immunizes less, prevalence is either (i) increasing in β if $\lambda(1 - c) < 2$, or (ii) first increasing in β for $\beta < 1/\lambda(1 - c)$, and then decreasing in β thereafter, if $\lambda(1 - c) > 2$.

B Efficiency properties of equilibrium

We are now in a position to relate the equilibrium characterization to the properties of the social planner's problem.

As we have anticipated in Section II, we observe that when immunization has marginal cost c , the optimal action for a social planner who wishes to maximize social welfare is either full eradication or to do nothing. When there are positive assortative interactions, Proposition 2 tells us that to eradicate the harmful state the social planner needs to provide a mass of immunizations which is equal to the steady state value that total prevalence would take in the absence of any immunity in the population. Since the benefit of not being infected by the harmful state is normalized to 1, we have that eradication is better than providing no immunity if and only if $c < 1$. Under negative assortative interactions, eradication can be achieved for an even smaller cost, and so when $c < 1$ eradication is again optimal.

Notice that the distribution of immunity across groups in equilibrium is efficient. When there are positive assortative interactions, the two groups immunize at the same rate and so total immunity is, in equilibrium, divided equally across groups. When there are negative assortative interactions, the equilibrium allocation of immunity is as asymmetric as possible, conditional on the level of total immunity, since $\pi_A > 0$ only if $\pi_B = 1$.

Given that immunization carries a positive externality, it is of course not surprising that equilibrium immunization rates are inefficiently low. What these observations imply, however, is that the source of inefficiency rests entirely on the aggregate level of immunization. The welfare loss arising from equilibrium immunizations relative to the first-best outcome is highest for intermediate ranges of λ and c . For extreme values of either λ or c , the equilibrium and the first-best immunization rates coincide.

IV The effects of group-level heterogeneity

Thus far group labels have been relevant only to the extent that they influence the interactions among individuals. Yet there are often underlying differences between individuals of differ-

ent populations that correlate with their matching patterns. In this section we investigate equilibrium outcomes under a specific economic asymmetry between groups. In particular, we study parameters $c_B < c_A$, so that the cost of immunization is group-dependent. This specification captures the case where one group has easier access to immunization, or where one group is wealthier, so that the cost of immunization is lower relative to the difference in utility between being healthy and infected.¹⁹ As we shall show, the introduction of even small underlying differences across groups is amplified by a social multiplier, and the multiplier increases with the intensity of inter-group interactions.

In order to exhibit this point, we focus on positive assortative interactions, i.e., $\beta > 1/2$, and let $\lambda(1 - c_A) > 2$.²⁰ Recall that in this case, under symmetric costs, the unique equilibrium involves equal immunization rates and identical outcomes in the two groups. The following proposition characterizes equilibria under asymmetric costs.

Proposition 7 *Suppose $c_B < c_A$ and $\lambda(1 - c_A) > 2$. There exists $\bar{\beta} > 1/2$ such that: If $\beta \in [1/2, \bar{\beta}]$ in the unique equilibrium all agents in group B immunize, $\pi_B = 1$, while only a proportion of agents in A immunize, $\pi_A \in (0, 1)$; if $\beta \in [\bar{\beta}, 1]$, the unique equilibrium is interior and $\pi_B > \pi_A$.*

We now show how initial differences between the two groups are amplified by the pattern of interactions via a social multiplier effect.²¹ This amplification in fact drives the equilibrium to a boundary solution at $\beta = \bar{\beta}$. In this sense, arbitrarily small differences between the groups result in very asymmetric outcomes as the intensity of positive assortativity decreases. We denote by Δ_π the difference between π_B and π_A when $\beta = 1$ and, hence, the diffusion in the two groups is independent. Similarly, Δ_P is the difference between the prevalence of group A and the prevalence of group B when $\beta = 1$. Using our characterization, one obtains that

$$\Delta_\pi = [\pi_B - \pi_A]_{\beta=1} = \frac{c_A - c_B}{\lambda(1 - c_A)(1 - c_B)} = \Delta_P = [P_A - P_B]_{\beta=1} > 0$$

¹⁹Other forms of ex-ante asymmetries are also plausible. For example, one could consider the case where one group has a more effective natural immunity, or a higher rate of recovery, than the other group, i.e., $\lambda_B > \lambda_A$. This case is not identical, but is qualitatively similar, to the cost heterogeneity case. Details are available upon request. In the conclusion we also elaborate on the case in which the two groups have different sizes.

²⁰The results extend easily to the case of $1 < \lambda(1 - c_A) < 2$.

²¹Under a complementarity assumption, Goyal & Vigier (2010) find a result that emphasizes a different form of heterogeneity.

We now show how immunization rates and prevalence levels are driven apart as β decreases from one.

Corollary 1 *Suppose $c_B < c_A$, $\lambda(1 - c_A) > 1$ and $\beta > 1/2$.*

- *The difference between the equilibrium immunization rate of group B and group A is the product of Δ_π and a social multiplier $\Phi_\pi(\beta, c_A, c_B) > 1$, which is decreasing and convex in β .*
- *The difference between the equilibrium prevalence of group A and group B is the product of Δ_P and a social multiplier $\Phi_P(\beta, c_A, c_B)$, which is decreasing and convex in β .*
- *The difference between the equilibrium welfare of group B and of group A equals $c_A - c_B$ and therefore it is constant in β .*

Together, Proposition 7 and Corollary 1 illustrate how small initial differences across groups are amplified by a social multiplier, even in the case where there are positive assortative interactions. For all $\beta \geq 1/2$ we have that the immunization rate of the low cost group is higher than the immunization rate of the high cost group and that as β decreases from 1 the immunization rate of the low cost group increases, whereas the immunization rate of the high cost group decreases. Similarly, the prevalence of the low cost group is lower than the prevalence of the high cost group and, as β decreases from 1, the prevalence of the low cost group decreases, whereas the prevalence of the high cost group increases.

In light of these results, we conclude that the structure of group interactions that we study tends to produce highly asymmetric outcomes. The only exception to this tendency is under positive assortative interactions when either the two groups are inherently identical, or they are nearly identical, and the intensity of assortativity is very strong.

V Conclusion

We have explored the diffusion of a harmful state through a population in an environment where immunization is possible, but costly, and where contact patterns generally depend on types in the population. This framework allows us to study, in particular, the impact of the degree of positive or negative assortativity on strategic aspects of immunization.

We have first studied the problem of a social planner who designs a policy to eradicate the spread of the harmful state. We have then derived the equilibrium of a game in which agents,

averse to the possibility of being infected by the harmful state, choose whether to immunize or not, anticipating the consequence of their choices on the diffusion process.²²

Our analysis demonstrates that the details of group interactions have a profound effect on how a planner should respond to reducing the prevalence of the harmful state. As a consequence, we have derived the conditions under which it is more valuable for the social planner to know the exact pattern of interactions. Furthermore, we have shown that inter-group interactions generate large asymmetries across groups in the equilibrium outcomes regarding welfare and prevalence, even when groups are, *ex-ante*, very similar. It is therefore important to account for the structure of interactions when studying, empirically, the causes of observed asymmetries in behavior and diffusion outcomes across groups. Failing to account for the effects of group interactions would, in this setting, lead to overestimating the underlying differences across groups, and consequently to suboptimal immunization plans.

Our model is intentionally stark, allowing us to derive analytical results more easily, but also precluding the study of certain issues relevant to the broader motivation of our work. On the one hand, the messages of this paper indicate that incorporating information about the structure of society and the contact patterns among individuals is of primary importance for making inferences about diffusion processes. On the other hand, understanding the precise implications of these details will require further study.

For example, some preventative measures are effective for only a short time, and the decision about whether to take such measures is more appropriately modeled as a dynamic problem, taking into account time preferences and the possibility of a changing environment. There is also a large scope for generalizing the nature of the diffusion process through society. For example, the population could change size over time, in ways that reflect, for example, a death rate that depends on current prevalence. Also, the rule we have specified by which individuals spread infection is but one example, and different diffusion rules have the possibility to generate qualitative differences in results.

One could consider as well more complex type profiles and contact patterns in order to

²²We note here a simple observation that emphasizes the importance of incorporating economic aspects into diffusion models, as has been recognized for some time and is expressed clearly in the work of Bauch & Earn (2004) and Geoffard & Philipson (1996). In our equilibrium model total prevalence tends to *decrease* in the intensity of infection λ . The opposite conclusion is obtained in a non-strategic model where immunization decisions do not react to changes in λ . As the intensity of infection increases, equilibrium immunization rates increase to the level that, conditional on not having immunity, the proportion of time infected remains constant ($\rho = c$). Given that the non-immunized individuals are infected at a constant rate, but more individuals are immunized, prevalence must decrease. As a result, welfare is constant, as the increased cost of immunization are exactly offset by the decrease in prevalence.

generalize the insights from the present paper. We leave most of these matters for future work. We conclude, however, with a brief discussion of an extension to our model that allows us to study the effect of group size. This demonstrates, in particular, that our conclusions are robust to the introduction of this form of asymmetries.

A A consideration of the effects of group size

We set the mass of group A equal to one, as before, while the mass of group B is $q \in (0, 1]$. So group A is the majority group and group B is the minority group. Let us define β as the probability with which a randomly chosen contact of a group B individual belongs to group B . Since group A is larger than group B , the fact that all agents have the same number of interactions implies that group A individuals will be more positively assortative (or less negatively assortative, as the case may be) than group B individuals. Defining γ as the probability with which a randomly chosen contact of a group A agent comes from group A , we have the following consistency condition

$$1 - \gamma = q(1 - \beta).$$

Using this condition, we obtain new expressions for the probabilities that an agent from either group interacts with an infected agent:

$$\begin{aligned}\theta_A &= [1 - (1 - \beta)q](1 - \pi_A)\rho_A + (1 - \beta)q(1 - \pi_B)\rho_B \\ \theta_B &= \beta(1 - \pi_B)\rho_B + (1 - \beta)(1 - \pi_A)\rho_A.\end{aligned}$$

With these expressions, one can characterize steady states of the system as a function of β and q .

We consider first the planner's problem of eradicating the harmful state with the minimal budget. The main intuitions from Section II carry over in a straightforward manner. First, observe that for any value of q , the value of the mixing parameter $\hat{\beta} = \frac{q}{1+q}$ corresponds to homogeneous mixing in which the distribution of all agents' contacts corresponds to the population frequencies. Hence, the value $\hat{\beta} = \frac{q}{1+q}$ is what separates the cases of positively and negatively assortative interactions. Naturally, when $q = 1$ the model reduces to the specification of Section I and we have $\hat{\beta} = \frac{1}{2}$. When there are positive assortative interactions, $\beta > \hat{\beta}$, the planner should distribute immunization to the population evenly. This corresponds to setting immunization levels proportional to population sizes, so that group A receives a

portion $\frac{1}{1+q}$ of the immunization budget and $\pi_A = \pi_B$. On the other hand, under negative assortativity, $\beta < \hat{\beta}$, the optimal allocation involves focusing the resources on one group. Here, the groups are not identical, and it turns out to be optimal to focus first on the minority group, and to devote remaining resources to the majority group only if necessary. The reason is that the majority group has extra interactions amongst themselves, and so immunizations there leverage the negative assortativity less effectively.²³ The minimum budget required to achieve immunization is decreasing in q and β .

The equilibrium characterization is significantly more involved for the case of unequal groups. Here we describe some partial results that suggest the main features are similar to the case of equal size groups, but that additional subtleties arise due to the asymmetry between groups. First, for all values of q and β there exists an interior equilibrium, which has immunization levels that are the same across groups, and constant in β . When there are positive assortative interactions, this is the unique equilibrium, as in the case of equal group sizes. Under negative assortativity, $\beta < \hat{\beta}$, there are asymmetric equilibria. In this case, if β is not too small, there is an equilibrium in which one group fully immunizes and the other group partially immunizes. The range of parameters for which this equilibrium occurs is larger for the case in which the minority group fully immunizes, as compared to the case where the majority group fully immunizes. In this case, the total immunization rate is higher than the case in which it is the majority group that fully immunizes.

²³If one considers cost heterogeneity in this context, then if c_A/c_B is sufficiently small, it can be more efficient to immunize the majority group first.

Appendix A: The Diffusion Model

Proposition 1 The following hold:

1. Suppose $\beta < 1/2$. $\pi_A < \pi_B$ iff $\rho_A < \rho_B$ iff $\theta_A < \theta_B$ iff $P_A > P_B$
2. Suppose $\beta > 1/2$. $\pi_A < \pi_B$ iff $\rho_A > \rho_B$ iff $\theta_A > \theta_B$ iff $P_A > P_B$

Proof of Proposition 1. We first note that, regardless of β , using the steady state conditions, we have that

$$\rho_A - \rho_B = \frac{\lambda}{(1 + \lambda\theta_A)(1 + \lambda\theta_B)}[\theta_A - \theta_B],$$

and so $\rho_A > \rho_B$ iff $\theta_A > \theta_B$. Next, note that

$$\theta_A - \theta_B = \beta P_A + (1 - \beta)P_B - \beta P_B - (1 - \beta)P_A = (2\beta - 1)(P_A - P_B).$$

So, if $\beta > 1/2$, then $\theta_A > \theta_B$ iff $P_A > P_B$, while if $\beta < 1/2$, then $\theta_A < \theta_B$ iff $P_A > P_B$. Hence if $\beta > 1/2$, then $\rho_A > \rho_B$ iff $\theta_A > \theta_B$ iff $P_A > P_B$, while if $\beta < 1/2$, then $\rho_A < \rho_B$ iff $\theta_A < \theta_B$ iff $P_A > P_B$.

Consider now the case of $\beta < 1/2$. Suppose that $\rho_A < \rho_B$, then we know that $P_A = (1 - \pi_A)\rho_A > P_B = (1 - \pi_B)\rho_B$, and so it is necessary that $1 - \pi_A > 1 - \pi_B$, or equivalently, $\pi_A < \pi_B$. Suppose now that $\pi_A < \pi_B$ and, for a contradiction, $\rho_A > \rho_B$. Since $\rho_A > \rho_B$ it has to be the case that $P_A = (1 - \pi_A)\rho_A < P_B = (1 - \pi_B)\rho_B$ which requires that $\pi_A > \pi_B$, a contradiction. Hence, if $\beta < 1/2$, then: $\pi_A < \pi_B$ iff $\rho_A < \rho_B$. This concludes the proof of part 1 of the proposition.

We now move to part 2. So, suppose that $\beta > 1/2$. We first show that $\rho_A < \rho_B$ implies that $\pi_B < \pi_A$. Note that

$$\rho_A = \frac{\lambda[\beta(1 - \pi_A)\rho_A + (1 - \beta)(1 - \pi_B)\rho_B]}{1 + \lambda[\beta(1 - \pi_A)\rho_A + (1 - \beta)(1 - \pi_B)\rho_B]} > \frac{\lambda\rho_A[\beta(1 - \pi_A) + (1 - \beta)(1 - \pi_B)]}{1 + \lambda\rho_A[\beta(1 - \pi_A) + (1 - \beta)(1 - \pi_B)]} = \frac{\rho_A\lambda A}{1 + \rho_A\lambda A}$$

where the first inequality follows because we have replaced ρ_B with ρ_A , the expression is increasing in ρ_B and, by assumption $\rho_A < \rho_B$, the second equality follows by simply defining $A = \beta(1 - \pi_A) + (1 - \beta)(1 - \pi_B)$. From the above we infer that

$$\rho_A > \frac{\rho_A\lambda A}{1 + \rho_A\lambda A} \text{ iff } 1 + \lambda\rho_A A > \lambda A \text{ iff } \rho_A > 1 - \frac{1}{\lambda A}.$$

Next note that

$$\rho_B = \frac{\lambda[\beta(1 - \pi_B)\rho_B + (1 - \beta)(1 - \pi_A)\rho_A]}{1 + \lambda[\beta(1 - \pi_B)\rho_B + (1 - \beta)(1 - \pi_A)\rho_A]} < \frac{\lambda\rho_B[\beta(1 - \pi_B) + (1 - \beta)(1 - \pi_A)]}{1 + \lambda\rho_B[\beta(1 - \pi_B) + (1 - \beta)(1 - \pi_A)]} = \frac{\rho_B\lambda B}{1 + \rho_B\lambda B}$$

where the first inequality follows because we have replace ρ_A with ρ_B , the expression is increasing in ρ_A and, by assumption $\rho_A < \rho_B$, the second equality follows by simply defining $B = \beta(1 - \pi_B) + (1 - \beta)(1 - \pi_A)$. From the above we infer that

$$\rho_B < \frac{\rho_B\lambda B}{1 + \rho_B\lambda B} \text{ iff } 1 + \lambda\rho_B B < \lambda B \text{ iff } \rho_B < 1 - \frac{1}{\lambda B}.$$

Since, by assumption, $\rho_B > \rho_A$, it is necessary that

$$1 - \frac{1}{\lambda B} > 1 - \frac{1}{\lambda A} \text{ iff } B > A \text{ iff } (2\beta - 1)(\pi_B - \pi_A) < 0 \text{ iff } \pi_B < \pi_A.$$

In a similar way one can prove that $\rho_A > \rho_B$ implies that $\pi_A < \pi_B$ (It is enough to take an upper bound for ρ_A and a lower bound for ρ_B , analogously to the above argument). And so we have that $\rho_A > \rho_B$ iff $\pi_A < \pi_B$.

Appendix B: The social planner's problem

Proof Proposition 2 The proof is divided in three steps.

First step. We start by asking what is the minimum mass of immunization needed to eradicate the undesirable state, given an exogenous allocation. Formally, for a total mass of immunizations available to the planner, an allocation is described by a number $\alpha \in \left[\max \left\{ 0, 1 - \frac{1}{\bar{\pi}^*(\alpha)} \right\}, 1/2 \right]$, where $\pi_A = \alpha\pi$ and $\pi_B = (1 - \alpha)\pi$. For a given allocation α , the minimum mass of immunizations needed to eradicate the disease is denoted by $\bar{\pi}^*(\alpha)$ and is the solution to

$$\rho_A(\alpha\bar{\pi}^*(\alpha), (1 - \alpha)\bar{\pi}^*(\alpha)) = \rho_B(\alpha\bar{\pi}^*(\alpha), (1 - \alpha)\bar{\pi}^*(\alpha)) = 0.$$

We derive $\bar{\pi}^*(\alpha)$ for the relevant case in which $\bar{\pi}^*(\alpha) > 0$, i.e., $\lambda > 1$. We assume this throughout.

For a given π_A and π_B , recall that the steady states conditions are

$$(4) \quad \rho_A = \frac{\lambda [\beta(1 - \pi_A)\rho_A + (1 - \beta)(1 - \pi_B)\rho_B]}{1 + \lambda [\beta(1 - \pi_A)\rho_A + (1 - \beta)(1 - \pi_B)\rho_B]},$$

$$(5) \quad \rho_B = \frac{\lambda [\beta(1 - \pi_B)\rho_B + (1 - \beta)(1 - \pi_A)\rho_A]}{1 + \lambda [\beta(1 - \pi_B)\rho_B + (1 - \beta)(1 - \pi_A)\rho_A]}.$$

Using condition (5), we can express ρ_A as follows

$$(6) \quad \rho_A = \rho_B \left[\frac{1 - \lambda\beta(1 - \pi_B)(1 - \rho_B)}{\lambda(1 - \beta)(1 - \pi_A)(1 - \rho_B)} \right].$$

Using condition (6) in condition (4) we obtain

$$(7) \quad \rho_A = \frac{\lambda\rho_B \left[\beta \left[\frac{1 - \lambda\beta(1 - \pi_B)(1 - \rho_B)}{\lambda(1 - \beta)(1 - \rho_B)} \right] + (1 - \beta)(1 - \pi_B) \right]}{1 + \lambda [\beta(1 - \pi_A)\rho_A + (1 - \beta)(1 - \pi_B)\rho_B]}.$$

Combining condition (6) with condition (7), we produce

$$(8) \quad \frac{\lambda \left[\beta \left[\frac{1 - \lambda\beta(1 - \pi_B)(1 - \rho_B)}{\lambda(1 - \beta)(1 - \rho_B)} \right] + (1 - \beta)(1 - \pi_B) \right]}{1 + \lambda [\beta(1 - \pi_A)\rho_A + (1 - \beta)(1 - \pi_B)\rho_B]} = \left[\frac{1 - \lambda\beta(1 - \pi_B)(1 - \rho_B)}{\lambda(1 - \beta)(1 - \pi_A)(1 - \rho_B)} \right].$$

Setting $\rho_A = \rho_B = 0$ in (8), we obtain that condition (8) holds if and only if

$$(9) \quad 1 - \lambda\beta(2 - \bar{\pi}^*(\alpha)) = (1 - \pi_A)(1 - \pi_B)\lambda^2(1 - 2\beta),$$

where $\bar{\pi}^*(\alpha) = \pi_A + \pi_B$. Replacing further $\pi_A = \alpha\bar{\pi}^*(\alpha)$ and $\pi_B = \bar{\pi}^*(\alpha) - \pi_A$, we obtain that 9 holds if and only if

$$1 - \lambda\beta(2 - \bar{\pi}^*(\alpha)) = (1 - \alpha\bar{\pi}^*(\alpha))(1 - \bar{\pi}^*(\alpha) + \alpha\bar{\pi}^*(\alpha))\lambda^2(1 - 2\beta),$$

and solving for $\bar{\pi}^*(\alpha)$ we obtain that the minimum mass of immunization needed to eradicate the disease under allocation α is

$$(10) \quad \bar{\pi}^*(\alpha) = \frac{A - \sqrt{A^2 + BC}}{2\lambda(1 - 2\beta)C},$$

where $A = \beta + \lambda(1 - 2\beta)$, $B = 4(1 - 2\beta)(1 - \lambda\beta - \lambda A)$ and $C = \alpha(1 - \alpha)$. We observe that when $\beta < 1/2$, for $\bar{\pi}^*(\alpha) > 0$ it has to be the case that $A - \sqrt{A^2 + BC} > 0$, whereas when $\beta > 1/2$, for $\bar{\pi}^*(\alpha) > 0$ it has to be the case that $A - \sqrt{A^2 + BC} < 0$.

Second Step. Now that we have derived $\bar{\pi}^*(\alpha)$, we compute the allocation that minimizes the smallest mass of immunizations needed to eradicate the disease. Formally, we look for

$$\alpha^* = \arg \min_{\alpha \in [\max\{0, 1 - \frac{1}{\bar{\pi}^*(\alpha)}\}, 1/2]} \bar{\pi}^*(\alpha).$$

Consider $\beta < 1/2$. Note that the sign of the derivative of $\bar{\pi}^*(\alpha)$ with respect to α is the same as the sign of the derivative of $\bar{\pi}^*(\alpha)$ with respect to C (recall that $\alpha \leq 1/2$). We have

$$\frac{d\bar{\pi}^*(\alpha)}{dC} = \frac{1}{2\lambda C^2 \sqrt{A^2 + BC}} \frac{1}{1 - 2\beta} \left[BC - 2A\sqrt{A^2 + BC} + 2A^2 \right].$$

Since $1 - 2\beta > 0$, it follows that $\frac{d\bar{\pi}^*(\alpha)}{dC} > 0$ if and only if

$$BC - 2A\sqrt{A^2 + BC} + 2A^2 > 0.$$

To see that this is the case note that $BC - 2A\sqrt{A^2 + BC} + 2A^2$ is increasing in C . In fact, simple algebra shows that the derivative of $BC - 2A\sqrt{A^2 + BC} + 2A^2$ with respect to C is positive whenever $B[\sqrt{A^2 + BC} - A] > 0$, which holds because $B < 0$ (as $\lambda > 1$

and $\beta < 1/2$) and $\sqrt{A^2 + BC} - A < 0$ (see observation at the end of step 1). Therefore, $BC - 2A\sqrt{A^2 + BC} + 2A^2 > [BC - 2A\sqrt{A^2 + BC} + 2A^2]_{C=0} = 0$. Hence, if $\beta < 1/2$, the best way to eradicate the disease is to focus first on one group, i.e., $\alpha^* = \max\{0, 1 - \frac{1}{\bar{\pi}^*(\alpha)}\}$.

Consider now the case where $\beta > 1/2$. Recall the expression above for $\frac{d\bar{\pi}^*(\alpha)}{dC}$ and note that, since $1 - 2\beta < 0$, it follows that $\frac{d\bar{\pi}^*(\alpha)}{dC} < 0$ if and only if

$$BC - 2A\sqrt{A^2 + BC} + 2A^2 > 0.$$

If B is positive, we can apply the same argument used above to show that indeed $BC - 2A\sqrt{A^2 + BC} + 2A^2 > 0$. Suppose $B < 0$; since $1 - 2\beta < 0$, the fact that $B < 0$ implies that $1 - \lambda\beta - \lambda A > 0$ or, equivalently, $1 - \lambda\beta > \lambda A$. Simple algebra shows that $BC - 2A\sqrt{A^2 + BC} + 2A^2$ is decreasing in A ; so, expression $BC - 2A\sqrt{A^2 + BC} + 2A^2$ is minimised at the highest value that A can take, which is A such that $1 - \lambda\beta = \lambda A$. But at this value $B = 0$ and therefore we have that $BC - 2A\sqrt{A^2 + BC} + 2A^2 > [BC - 2A\sqrt{A^2 + BC} + 2A^2]_{B=0} = 0$. Hence, if $\beta > 1/2$ the best way to eradicate the disease is to split the mass of immunizations across groups, i.e., $\alpha^* = 1/2$.

Third step. Using expression (10) for $\bar{\pi}^*(\alpha)$ and our conclusion in step 2, we can now derive the minimum mass of immunization needed to eradicate the disease. In particular, if $\beta > 1/2$, the minimum mass of immunization needed to eradicate the disease is obtained by evaluating $\bar{\pi}^*(\alpha)$ at $\alpha = 1/2$, which gives $\bar{\pi}^* = 2(1 - \frac{1}{\lambda})$. Note that this is independent of β .

Consider the case of $\beta < 1/2$. If the total amount of immunization needed is less than 1, $\bar{\pi}^* < 1$, then, the minimum amount of immunization needed to eradicate the disease is obtained by taking the $\lim_{\alpha \rightarrow 0} \bar{\pi}^*(\alpha)$. This gives us

$$\bar{\pi}^* = 1 - \frac{1}{\lambda} \frac{1 - \lambda\beta}{\lambda(1 - 2\beta) + \beta},$$

which is consistent with our initial hypothesis that $\bar{\pi}^* < 1$ if and only if $\beta < 1/\lambda$. It is easy to verify that $\bar{\pi}^*$ is increasing in β for all $\beta < 1/\lambda$. Finally, suppose that $\bar{\pi}^* > 1$, so that in the optimal assignment $\pi_B = 1$ and $\pi_A = \bar{\pi}^* - 1$. Then, by substituting $\alpha = 1 - \frac{1}{\bar{\pi}^*}$ in expression (10) and solving for $\bar{\pi}^*$, we obtain that $\bar{\pi}^* = 2 - \frac{1}{\lambda\beta}$, which is increasing in β .

Appendix C: Equilibrium Analysis

Proof of Proposition 3 and Proposition 4. Propositions 3 and 4 are proved jointly. The proof is divided in two steps. In the first step we characterize all equilibria of our game. The second step identifies the stable equilibria.

First step: equilibrium characterization. For $\pi_A = \pi_B = 0$ we have that $\rho_A = \rho_B = \rho^* = 1 - 1/\lambda$. For this to be an equilibrium, it must be that $\rho^* \leq c$, which is equivalent to $\lambda(1 - c) \leq 1$. Note also that when $\lambda(1 - c) < 1$, the only equilibrium is $\pi_A = \pi_B = 0$, since $\pi_g > 0$ implies $\rho_g < 1 - 1/\lambda < c$, contradicting the equilibrium requirement that $\rho_g \geq c$ when $\pi_g > 0$.

We now focus on the case $\lambda(1 - c) > 1$. We start with a characterization of interior equilibria. An interior equilibrium requires that, using (3), $\rho_A = c = \rho_B$, which implies by (2) that $\theta_A = \theta_B$, which, using (1) and provided $\beta \neq 1/2$, implies that $\pi_A = \pi_B$. Setting $\pi_A = \pi_B = \pi^*$ and solving $\rho_A = c$ produces

$$\pi^* = 1 - \frac{1}{\lambda(1 - c)}.$$

Thus, provided $\beta \neq 1/2$, for $\lambda(1 - c) > 1$ there exists a unique interior equilibrium, in which case it is symmetric. Next, it is never an equilibrium for $\pi_A = \pi_B = 1$, since any player could profitably deviate by not immunizing, given that all other players are immunized.

The only remaining possibility for equilibria are therefore asymmetric. Without loss of generality, we focus on the case $\pi_A < \pi_B$. By the above, any additional equilibrium must be such that either $\pi_A = 0$ or $\pi_B = 1$, as otherwise it would be interior and therefore not asymmetric. There are therefore three possibilities: (1) $\pi_A = 0$ and $\pi_B \in (0, 1)$, (2) $\pi_A = 0$ and $\pi_B = 1$, and (3) $\pi_A \in (0, 1)$ and $\pi_B = 1$. We consider each of these possibilities in turn.

Case 1. Suppose $\pi_A = 0$ and $\pi_B \in (0, 1)$. In equilibrium $\rho_B = c$. Using this and the steady-state conditions (2) we obtain that the following two conditions must hold

$$(11) \quad \rho_A = \frac{\lambda[\beta\rho_A + (1 - \beta)(1 - \pi_B)c]}{1 + \lambda[\beta\rho_A + (1 - \beta)(1 - \pi_B)c]},$$

$$(12) \quad c = \frac{\lambda[\beta(1 - \pi_B)c + (1 - \beta)\rho_A]}{1 + \lambda[\beta(1 - \pi_B)c + (1 - \beta)\rho_A]}.$$

Solving condition (11) for π_B we obtain

$$\hat{\pi}_B(\rho_A) = \frac{\lambda c(1 - \beta) - \lambda \beta \rho_A^2 + \rho_A(\lambda \beta(1 + c) - 1 - \lambda c)}{\lambda c(1 - \beta)(1 - \rho_A)},$$

and, similarly, solving condition (12) for π_B we obtain

$$\tilde{\pi}_B(\rho_A) = 1 - \frac{1}{\lambda \beta(1 - c)} + \frac{1 - \beta}{\beta c} \rho_A.$$

Note that an equilibrium is defined by $\pi_B = \pi^*$ and $\rho_A = \rho^*$ such that $\hat{\pi}_B(\rho^*) = \tilde{\pi}_B(\rho^*) = \pi^* \in (0, 1)$ and $\rho^* \in (0, c]$, where $\rho^* \leq c$ is required for $\pi_A = 0$ to be consistent with equilibrium.

Note that $\hat{\pi}_B(0) = 1$, $\hat{\pi}_B(c) = \frac{\lambda(1-c)-1}{\lambda(1-c)(1-\beta)} > 0$ and $\hat{\pi}_B(\rho)$ is strictly concave in $\rho \in [0, c]$. Thus $\hat{\pi}_B(\rho) > 0$ for all $\rho \in [0, c]$. Also note that $\tilde{\pi}_B(0) = 1 - \frac{1}{\lambda \beta(1-c)} < \hat{\pi}_B(0)$, $\tilde{\pi}_B(c) = \frac{\lambda(1-c)-1}{\lambda(1-c)\beta}$ and $\tilde{\pi}_B(\rho)$ is linear and strictly increasing in $\rho \in [0, c]$. Thus there is at most one such equilibrium (π^*, ρ^*) .

A necessary condition for equilibrium is that $\hat{\pi}_B(c) \leq \tilde{\pi}_B(c)$ which is equivalent to $\beta \leq 1/2$. So suppose $\beta \leq 1/2$. Recall that ρ^* solves $\hat{\pi}_B(\rho^*) = \tilde{\pi}_B(\rho^*)$, which is well-defined and in $(0, c]$. Define $\tilde{\rho}$ such that $\tilde{\pi}_B(\tilde{\rho}) = 1$, which gives

$$\tilde{\rho} = \frac{c}{\lambda(1-c)(1-\beta)}.$$

Define

$$\Psi(\beta, c, \lambda) = (1 - \beta)(1 - c)[\beta \lambda - 1] - \beta c,$$

and notice that $\hat{\pi}_B(\tilde{\rho}) < 1$ if and only if $\Psi(\beta, c, \lambda) < 0$.

The equilibrium exists if, and only if, $\tilde{\pi}_B(\rho^*) < 1$, i.e., $\tilde{\pi}_B(\rho^*) < \tilde{\pi}_B(\tilde{\rho})$. Since $\tilde{\pi}_B(\rho)$ is strictly increasing, this is equivalent to $\rho^* < \tilde{\rho}$. This is equivalent to $\hat{\pi}_B(\tilde{\rho}) < \tilde{\pi}_B(\tilde{\rho}) = 1$. Finally, by the definition of Ψ , this is equivalent to $\Psi(\beta, c, \lambda) < 0$.

The following lemma provides the region of parameters in which this condition is satisfied.

Lemma 1 *Suppose $\beta \in [0, 1/2)$ and $\lambda(1 - c) > 1$; let $\Psi(\beta, c, \lambda) = (1 - \beta)(1 - c)[\beta \lambda - 1] - \beta c$.*

1. *Suppose $c < 1/2$.*

1.1 *If $\lambda(1 - c) < 2$ then $\Psi(\beta, c, \lambda) < 0$.*

1.2 If $\lambda(1 - c) > 2$ then there exists $\underline{\beta} \in (0, 1/2)$ such that $\Psi(\underline{\beta}, c, \lambda) = 0$, and $\Psi(\beta, c, \lambda) < 0$ if and only if $\beta \in [0, \underline{\beta})$.

2. Suppose $c > 1/2$.

2.1 If $\lambda(1 - c) < [1 + 2(c(1 - c))^{1/2}]$ then $\Psi(\beta, c, \lambda) < 0$.

2.2 If $\lambda(1 - c) \in [[1 + 2(c(1 - c))^{1/2}], 2]$ then there exists $0 < \underline{\beta} < \bar{\beta} < 1/2$ such that $\Psi(\underline{\beta}, c, \lambda) = \Psi(\bar{\beta}, c, \lambda) = 0$, and $\Psi(\beta, c, \lambda) \leq 0$ if and only if $\beta < \underline{\beta}$ or $\beta > \bar{\beta}$.

2.3 If $\lambda(1 - c) > 2$ then there exists $\underline{\beta} \in (0, 1/2)$ such that $\Psi(\underline{\beta}, c, \lambda) = 0$, and $\Psi(\beta, c, \lambda) \leq 0$ if and only if $\beta < \underline{\beta}$.

Proof of Lemma 1. The following two observations are used in the proof.

Observation 1. $\Psi(0, c, \lambda) = -(1 - c) < 0$, $\Psi(1/2, c, \lambda) = \frac{1-c}{2} \left[\frac{\lambda}{2} - 1 \right] - \frac{c}{2}$, which is negative iff $\lambda(1 - c) \leq 2$.

Observation 2.

$$\frac{\partial \Psi(\beta, c, \lambda)}{\partial \beta} = (1 - c)[1 + \lambda(1 - 2\beta)] - c,$$

which is positive if $c \leq 1/2$, whereas if $c > 1/2$ it is positive for low β and negative for high β (still below $1/2$).

First suppose that $\lambda(1 - c) > 2$; observations 1-2 imply that there exists a $\underline{\beta} < 1/2$ such that $\Psi(\underline{\beta}, c, \lambda) = 0$ and that $\Psi(\beta, c, \lambda) < 0$ if, and only if, $\beta < \underline{\beta}$; this proves part 1.2. and part 2.3. Second suppose that $c < 1/2$ and $\lambda(1 - c) < 2$; observations 1-2 imply that $\Psi(\beta, c, \lambda) < 0$ for all $\beta < 1/2$, which proves part 1.1.

Third, suppose that $c > 1/2$ and $\lambda(1 - c) < 2$. Recall that observation 1 and observation 2 tell us that $\Psi(0, c, \lambda) < 0$, $\Psi(1/2, c, \lambda) < 0$, and $\Psi(\beta, c, \lambda)$ is first increasing and then decreasing in β . So, there are two possibilities. One possibility is that the maximum value of $\Psi(\beta, c, \lambda)$ is negative, in which case $\Psi(\beta, c, \lambda) < 0$. The other possibility is that the maximum value of $\Psi(\beta, c, \lambda)$ is positive, in which case there exist $\underline{\beta} < \bar{\beta} < 1/2$ such that $\Psi(\underline{\beta}, c, \lambda) = \Psi(\bar{\beta}, c, \lambda) = 0$ and, $\Psi(\beta, c, \lambda) < 0$ for all $\beta < \underline{\beta}$ and for all $\beta > \bar{\beta}$. To distinguish between these two cases, define $\beta^* = \arg \max_{\beta \in [0, 1/2]} \Psi(\beta, c, \lambda)$. This is given by

$$\beta^* = \frac{1}{2} - \frac{2c - 1}{2\lambda(1 - c)},$$

and note that

$$\Psi(\beta^*, c, \lambda) = \frac{1}{2\lambda(1-c)} \left[\frac{1-c}{2}(\lambda(1-c) + 2c - 1) \left(\lambda - \frac{1}{1-c} \right) - c(\lambda(1-c) - 2c + 1) \right].$$

It is now easy to verify that $\Psi(\beta^*, c, \lambda) \leq 0$ iff $\lambda(1-c) < 1 + 2(c(1-c))^{1/2}$. This concludes part 2.1 and 2.2. The proof of lemma 1 is now completed.

Case 2. Suppose $\pi_A = 0$ and $\pi_B = 1$. For equilibrium $\rho_A \leq c \leq \rho_B$. If $\beta > 1/2$ then $\theta_B < \theta_A$ which implies that $\rho_B < \rho_A$, a contradiction. Hence, $\beta \leq 1/2$. Using the steady state condition (2) we obtain

$$\rho_A = \frac{\lambda\beta - 1}{\lambda\beta},$$

and we derive ρ_B as

$$\rho_B = \frac{(1-\beta)(\lambda\beta - 1)}{\beta + (1-\beta)(\lambda\beta - 1)}.$$

Since $\rho_A \in (0, c]$, then $\beta > \frac{1}{\lambda}$ and $\beta \leq \hat{\beta} = \frac{1}{\lambda(1-c)}$. Since $\rho_B \geq c$ then $\Psi(\beta, c, \lambda) \geq 0$. Using Lemma 1, we conclude that $\beta > \frac{1}{\lambda}$, $\beta \leq \hat{\beta} = \frac{1}{\lambda(1-c)}$ and $\Psi(\beta, c, \lambda) \geq 0$ if and only if $\lambda(1-c) > 2$ and $\beta \in [\underline{\beta}, \hat{\beta}]$, where we recall that $\Psi(\underline{\beta}, c, \lambda) = 0$.

Case 3. Suppose $\pi_A \in (0, 1)$ and $\pi_B = 1$. In equilibrium $\rho_A = c \leq \rho_B$ and a necessary condition for $\rho_A \leq \rho_B$ is that $\beta \leq 1/2$. So, let $\beta \leq 1/2$. Solving for $\rho_A = c$ we obtain

$$\pi_A = 1 - \frac{1}{\lambda\beta(1-c)},$$

and $\pi_A \in (0, 1)$ if and only if $\beta > \hat{\beta} = \frac{1}{\lambda(1-c)}$. It is easy to check that $\rho_B \geq c$ for all $\beta \leq 1/2$. Note that $\hat{\beta} \leq 1/2$ if and only if $\lambda(1-c) > 2$. So, this equilibrium exists whenever $\lambda(1-c) > 2$ and $\beta > \in (\hat{\beta}, 1/2]$. This concludes the characterization of equilibria.

Second step: stability. The equilibrium characterisation shows that: 1. there is only one interior equilibrium and the interior equilibrium exists for all β , 2. for $\beta > 1/2$ the equilibrium is unique, 3. for $\beta < 1/2$ there is only one (upon relabelling) asymmetric equilibrium. Hence, to prove the result on global stability, it is sufficient to show that the derivative of the best response of group B with respect to group A 's immunization rate, evaluated at the interior

equilibrium, is larger than -1 if, and only if, $\beta > 1/2$. Formally,

$$\frac{d\pi_B^*(\pi_A)}{d\pi_A}\Big|_{\pi_A=\pi_B=\pi^*} > -1 \text{ if, and only if, } \beta > 1/2.$$

We now prove this statement. We first observe that for every $(\pi_B^*(\pi_A), \pi_A)$ with π_A sufficiently close to the interior equilibrium value π^* , we have that $\rho_B(\pi_B^*(\pi_A), \pi_A) = c$, i.e., the proportion of infected agents of group B who are not immunized does not change with π_A , taking into account that immunization in B is optimal given π_A . This directly implies that $\theta_B(\pi_B^*(\pi_A), \pi_A)$ is also constant when π_A is sufficiently close to π^* , i.e.,

$$\frac{d\theta_B(\pi_B^*(\pi_A), \pi_A)}{d\pi_A}\Big|_{\pi^*} = 0.$$

Note that

$$\begin{aligned} \frac{d\theta_B(\pi_B^*(\pi_A), \pi_A)}{d\pi_A} &= -\beta c \frac{d\pi_B^*(\pi_A)}{d\pi_A} + (1 - \beta) \left[(1 - \pi_A) \frac{d\rho_A(\pi_A, \pi_B^*(\pi_A))}{d\pi_A} - \rho_A \right], \\ \frac{d\rho_A(\pi_A, \pi_B^*(\pi_A))}{d\pi_A} &= \frac{\partial \rho_A(\pi_A, \pi_B^*(\pi_A))}{\partial \pi_A} + \frac{\partial \rho_A(\pi_A, \pi_B^*(\pi_A))}{\partial \pi_B} \frac{d\pi_B^*}{d\pi_A}, \end{aligned}$$

and that at π^* we have that $\rho_A = c$. Hence, $\frac{d\theta_B(\pi_B^*(\pi_A), \pi_A)}{d\pi_A}\Big|_{\pi^*} = 0$ if and only if

$$\frac{d\pi_B^*(\pi_A)}{d\pi_A}\Big|_{\pi^*} = \frac{(1 - \beta) \left[(1 - \pi^*) \frac{\partial \rho_A(\pi_A, \pi_B^*(\pi_A))}{\partial \pi_A}\Big|_{\pi^*} - c \right]}{\beta c - (1 - \beta)(1 - \pi^*) \frac{\partial \rho_A(\pi_A, \pi_B^*(\pi_A))}{\partial \pi_B}\Big|_{\pi^*}}.$$

Next, $\frac{d\pi_B^*(\pi_A)}{d\pi_A}\Big|_{\pi^*} > -1$ if and only if

$$(1 - \beta) \left[(1 - \pi^*) \frac{\partial \rho_A(\pi_A, \pi_B^*(\pi_A))}{\partial \pi_A}\Big|_{\pi^*} - c \right] > (1 - \beta)(1 - \pi^*) \frac{\partial \rho_A(\pi_A, \pi_B^*(\pi_A))}{\partial \pi_B}\Big|_{\pi^*} - \beta c,$$

or, equivalently,

$$(1 - \beta)(1 - \pi^*) \left[\frac{\partial \rho_A(\pi_A, \pi_B^*(\pi_A))}{\partial \pi_A}\Big|_{\pi^*} - \frac{\partial \rho_A(\pi_A, \pi_B^*(\pi_A))}{\partial \pi_B}\Big|_{\pi^*} \right] > c(1 - 2\beta).$$

Using the fact that $\rho_A = \lambda\theta_A/[\lambda\theta_A + 1]$ and that $\theta_A = \beta(1 - \pi_A)\rho_A + (1 - \beta)(1 - \pi_B)\rho_B$, we

obtain

$$\frac{\partial \rho_A(\pi_A, \pi_B^*(\pi_A))}{\partial \pi_A} \Big|_{\pi^*} - \frac{\partial \rho_A(\pi_A, \pi_B^*(\pi_A))}{\partial \pi_B} \Big|_{\pi^*} = \frac{c(1-2\beta)}{(1+\lambda\theta_A(\pi^*))^2},$$

and we can rewrite the condition above as

$$(1-\beta)(1-\pi^*) \frac{c(1-2\beta)}{(1+\lambda\theta_A(\pi^*))^2} > c(1-2\beta),$$

which holds if and only if $\beta > 1/2$. This completes the proof for stability.

Proof of Proposition 5. In view of Proposition 4 we consider two cases. The first case is when in equilibrium $\pi_B = 1$. Since $\pi_B = 1$, then $P_B = 0$ and $W_B = 1 - c$. Prevalence in group A must be positive, for otherwise $\rho_B = 0 < c$, which contradicts equilibrium. So group A 's prevalence is higher than group B 's prevalence. Furthermore, if $\pi_A > 0$, then $\rho_A = c$ and so $W_A = W_B$. If $\pi_A = 0$, then equilibrium implies that $\rho_A < c$, and so $W_A > W_B$. The second case is when in equilibrium $\pi_B \in (0, 1)$. Since $\pi_B \in (0, 1)$, equilibrium implies that $\rho_B = c$ and therefore $P_B = (1 - \pi_B)c$ and $W_B = 1 - c$. Furthermore, if $\pi_B \in (0, 1)$ then $\pi_A = 0$ and so $\rho_A < c$. Hence, $W_A = 1 - \rho_A > 1 - c = W_B$. Finally, since $\rho_A < c = \rho_B$ it follows that $\theta_B > \theta_A$. Given that $\beta < 1/2$ and that $\theta_B - \theta_A = (P_B - P_A)(2\beta - 1)$, we have that $\theta_B > \theta_A$ if and only if $P_A > P_B$.

Proof of Proposition 6. In view of Proposition 4 we observe that since $\pi_B > 0$, $W_B = 1 - c$. We now consider three cases. The first is when $\pi_A = 0$ and $\pi_B = 1$; in this case $P_A = \rho_A = 1 - \frac{1}{\lambda\beta}$ proving that P_A increases in β and welfare $W_A = 1 - P_A$ is decreasing. The second case is when $\pi_B = 1$ and $\pi_A \in (0, 1)$. The welfare is $W_A = 1 - c$; the prevalence is $P_A = (1 - \pi_A)\rho_A = (1 - \pi_A)c$. Since $\pi_A = 1 - \frac{1}{\lambda\beta(1-c)}$, an increase in β , increases π_A , and so it decreases P_A . The last case is when $\pi_A = 0$ and $\pi_B \in (0, 1)$. The welfare is $W_A = 1 - \rho_A$ and $P_A = \rho_A$. Solving equilibrium conditions (11) and (12) we obtain an explicit solution for ρ_A . Taking the derivative with respect to β , one can show that the sign of $\frac{d\rho_A}{d\beta}$ is the opposite of the sign of

$$F(\beta) = c + \gamma(1 - 2c) + \beta(1 - 2\gamma)(1 - 2c) - \sqrt{-4\gamma c(1 - \beta)(1 - 2\beta) + [\beta + (1 - 2\beta)(c + \gamma)]^2},$$

where $\gamma = \lambda(1 - c)$. By inspection $F(\beta) < 0$ for all $\beta < 1/2$, which shows that ρ_A is increasing in β .

To conclude the proof we note that if $\lambda(1 - c) < 2$, the equilibrium characterization implies that in equilibrium $\pi_A = 0$, and so the proof follows. If $\lambda(1 - c) > 2$ in equilibrium $\pi_A = 0$ if

and only if $\beta < 1/\lambda(1 - c)$.

Proof of Proposition 7. Solving the equilibrium conditions of an interior equilibrium together with the steady state conditions we obtain:

$$(13) \quad \pi_A = 1 - \frac{c_A(1 - c_B)\beta - (1 - \beta)c_B(1 - c_A)}{\lambda(2\beta - 1)c_A(1 - c_A)(1 - c_B)}$$

$$(14) \quad \pi_B = 1 - \frac{c_B(1 - c_A)\beta - (1 - \beta)c_A(1 - c_B)}{\lambda(2\beta - 1)c_B(1 - c_A)(1 - c_B)}.$$

By investigation of these two expressions it is easy to check that for all $\beta \in (1/2, 1]$ the following holds: (1.) $\pi_B > \pi_A$, (2.) π_B is decreasing in β , while π_A is increasing in β . Next, define $\bar{\beta}$ to be such that $\pi_B(\bar{\beta}) = 1$. Solving we obtain that

$$\bar{\beta} = \frac{c_A(1 - c_B)}{c_A(1 - c_B) + c_B(1 - c_A)} > 1/2,$$

and that $\pi_A(\bar{\beta}) > 0$. These considerations imply that for all $\beta \in [\bar{\beta}, 1]$ there exists an interior equilibrium, π_A and π_B are defined by expression (13) and (14), and $\pi_B > \pi_A$. We now show that for all $\beta \in [1/2, \bar{\beta}]$ there exists an equilibrium where $\pi_B = 1$ and $\pi_A \in (0, 1)$. Suppose $\pi_B = 1$ and $\pi_A \in (0, 1)$. For equilibrium, $\rho_A = c_A$; solving the equilibrium condition for π_A we obtain that

$$(15) \quad \pi_A = 1 - \frac{1}{\lambda\beta(1 - c_A)},$$

and this implies that

$$\rho_B = \frac{(1 - \beta)c_A}{\beta(1 - c_A) + (1 - \beta)c_A} \geq c_B$$

if, and only if, $\beta < \bar{\beta}$. Again, note that at $\beta = 1/2$, π_A is still positive. It is easy to establish that for $\beta > 1/2$ there is only one equilibrium.

Proof of Corollary 1. Using Proposition 7 and taking the difference between the equilibrium immunization rate of group B and group A, it is easy to verify that this equals $\Delta_\pi \Phi_\pi(\beta, c_A, c_B)$, where

$$\Phi_\pi(\beta, c_A, c_B) = \begin{cases} \frac{1 - c_B}{\beta(c_A - c_B)}, & \text{if } \beta \in (1/2, \bar{\beta}) \\ \frac{c_A c_B(2\beta - 1) + (1 - \beta)(c_B + c_A)}{c_A c_B(2\beta - 1)}, & \text{if } \beta > \bar{\beta} \end{cases}$$

and that $\Phi_\pi(\beta, c_A, c_B) > 1$ and it is decreasing and convex in β . Similarly, taking the difference

between the equilibrium prevalence of group A and group B it is easy to verify that it equals the product of Δ_P and a social multiplier $\Phi_P(\beta, c_A, c_B)$, where

$$\Phi_P(\beta, c_A, c_B) = \begin{cases} \frac{c_A(1-c_B)}{\beta(c_A-c_B)}, & \text{if } \beta \in (1/2, \bar{\beta}) \\ \frac{1}{2\beta-1}, & \text{if } \beta > \bar{\beta} \end{cases}$$

and that $\Phi_P(\beta, c_A, c_B) > 1$ and it is decreasing and convex in β . Finally, direct computation shows that the difference between the equilibrium welfare of group A and of group B equals $c_A - c_B$ and therefore it is constant in β .

References

- Anderson, Roy M., and Robert M. May.** 1984. “Spatial, temporal, and genetic heterogeneity in host populations and the design of immunization programmes.” *IMA Journal Of Mathematics Applied In Medicine And Biology*, 1(3): 233–266.
- Anderson, Roy M., and Robert M. May.** 1992. *Infectious diseases of humans: dynamics and control*. Vol. 26, Wiley Online Library.
- Bailey, Norman T. J.** 1975. “The mathematical theory of infectious diseases and its applications. 2nd edition.” *The mathematical theory of infectious diseases and its applications 2nd edition*, 413.
- Barrett, Scott.** 2003. “Global Disease Eradication.” *Journal of the European Economic Association*, 1(2-3): 591–600.
- Barrett, Scott, and Michael Hoel.** 2007. “Optimal disease eradication.” *Environment and Development Economics*, 12(5): 627–652.
- Bauch, Chris T., and David J. D. Earn.** 2004. “Vaccination and the theory of games.” *Proceedings of the National Academy of Sciences of the United States of America*, 101(36): 13391–4.
- Boulier, Bryan L., Tejwant S. Datta, and Robert S. Goldfarb.** 2007. “Vaccination externalities.” *The BE Journal of Economic Analysis & Policy*, 7(1): 23.
- Brito, Dagobert L., Eytan Sheshinski, and Michael D. Intriligator.** 1991. “Externalities and compulsory vaccinations.” *Journal of Public Economics*, 45(1): 69–90.
- Echenique, Frederico, Roland G. Fryer, and Alex Kaufman.** 2006. “Is school segregation good or bad?” *The American economic review*, 96(2): 265–269.
- Galeotti, Andrea, and Sanjeev Goyal.** 2009. “Influencing the influencers: a theory of strategic diffusion.” *The RAND Journal of Economics*, 40(3): 509–532.
- Geoffard, Pierre-Yves, and Tomas Philipson.** 1996. “Rational epidemics and their public control.” *International Economic Review*.
- Geoffard, Pierre-Yves, and Tomas Philipson.** 1997. “Disease eradication: private versus public vaccination.” *The American Economic Review*, 87(1): 222–230.

- Goyal, Sanjeev, and Adrien Vigier.** 2010. “Endogenous Interaction and Vaccination.”
- Jackson, Matthew O., and Brian W. Rogers.** 2007. “Relating network structure to diffusion properties through stochastic dominance.” *The BE Journal of Theoretical Economics*, 7(1): 6.
- Jackson, Matthew O., and Dunia Lopez-Pintado.** 2011. “Diffusion and Contagion in Networks with Heterogeneous Agents and Homophily.”
- Jackson, Matthew O., and Leeat Yariv.** 2010. “Diffusion, strategic interaction, and social structure.” *Handbook of Social Economics, Elsevier*, 1–50.
- Kremer, Michael.** 1996. “Integrating behavioral choice into epidemiological models of AIDS.” *The Quarterly Journal of Economics*, 111(2): 549.
- López-Pintado, Dunia.** 2008. “Diffusion in complex social networks.” *Games and Economic Behavior*, 62(2): 573–590.
- Marsden, Peter V.** 1987. “Core discussion networks of Americans.” *American sociological review*, 122–131.
- Marsden, Peter V.** 1988. “Homogeneity in confiding relations.” *Social networks*, 10(1): 57–76.
- McPherson, Miller, Lynn Smith-Lovin, and James M. Cook.** 2001. “Birds of a feather: Homophily in social networks.” *Annual review of sociology*, 415–444.
- Montopoli, Leta, Samit Bhattacharyya, and Chris T. Bauch.** 2009. “The Free Rider Problem In Vaccination Policy And Implications For Global Eradication Of Infectious Diseases: A Two-Country Game Dynamic Model.” *Canadian Applied Mathematics Quarterly*, 17(2): 317–338.
- Moody, James.** 2001. “Race, School Integration, and Friendship Segregation in America1.” *American Journal of Sociology*, 107(3): 679–716.
- Reluga, Timothy C.** 2009. “An SIS epidemiology game with two subpopulations.” *Journal of Biological Dynamics*.

- Rowthorn, Robert.** 2006. “The Optimal Treatment of Disease Under a Budget Constraint.”
In *Explorations in Environmental and Natural Resource Economics: Essays in Honor of Gardner M. Brown, Jr.* Ed. R. Halvorsen and DF Layton. Cheltenham, UK: Edward Elgar Pub. 20–35.
- Toxvaerd, Flavio.** 2010*a*. “Infection, Acquired Immunity and Externalities in Treatment.”
- Toxvaerd, Flavio.** 2010*b*. “Recurrent Infection and Externalities in Prevention.”
- Vega-Redondo, Fernando.** 2007. *Complex social networks*. Cambridge University Press.