Herd Immunity and a Vaccination Game*

Wooyoung Lim† Franco P. Zhang‡

November 12, 2018

Abstract

Would the affected communities in West Africa voluntarily obtain herd immunity if a cure for Ebola was available? This paper theoretically and experimentally investigates people's vaccination choices in the context of a nonlinear public good game. We first consider the standard SIR model (Kermack and McKendrick, 1927) to capture the dynamic process of contagion and to show that the long-run infection probability decreases concavely as vaccination coverage rises, with the probability eventually vanishing when herd immunity is obtained. The nonlinear nature of the externality dramatically changes the social planner's problem such that it is socially optimal to eliminate an epidemic by obtaining herd immunity. We then consider an individual's strategic vaccination choice problem. A "vaccination game" is defined in which costly commitments (vaccination) are required of a fraction of the population to reach the critical level needed for herd immunity, without which defectors are punished by the natural contagion of epidemics. Our game-theoretic approach to herd immunity reveals that endogenous epidemic punishment is an effective mechanism, resulting in voluntary vaccination to obtain herd immunity, for which the orthodox principle of positive externalities fails to account. As a result, the socially optimal outcome can be approximated through voluntary vaccination. Our experimental implementation of a vaccination game in a controlled laboratory setting provides strong support for the major theoretical predictions.

Keywords: Vaccination Game, Economic Epidemiology, Nonlinear Public Goods, Laboratory Experiment

JEL classification: C72, C91, D62, H00, I18

*We are grateful to Marco Battaglini, Dongkyu Chang, Stephen Coate, John Duffy, Lata Gangadharan, Charles Holt, Duk Gyoo Kim, Philip Neary, In-Uck Park, Siyang Xiong, Jane Zhang, and Xiaojian Zhao for their valuable comments and suggestions. For helpful comments and discussions, we also thank seminar and conference participants at Cornell University, HKUST, Korea University, NYU Shanghai, Sungkyunkwan University, Tsinghua University, Yonsei University, the 2017 Asian Meeting of the Econometric Society, the 2017 Asia-Pacific Economic Science Association Conference, and the 92nd Western Economic Association Internal Annual Conference. This study is supported by a grant from the Research Grants Council of Hong Kong (Grant No. GRF-16502015).

†Department of Economics, Hong Kong University of Science and Technology, wooyoung@ust.hk

‡Department of Economics, Cornell University, pz82@cornell.edu

1 Introduction

Globally, infectious diseases are responsible for one-quarter of all deaths each year; even when infections do not kill, they reduce the quality of life of the hundreds of millions of people affected (Laxminarayan et al., 2006). Recent outbreaks, such as SARS in 2003, Ebola in 2014 and MERS in 2015, have forced citizens around the world to confront and reconsider the high social and economic costs of epidemic control.[1] The problem of epidemic control is of fundamental economic importance because disease and infection are irrefutably associated with uncertainty and externalities, which have long been central concepts in economics. According to a WHO report (WHO Ebola Response Team, 2015), behavioral changes at the community level are key to the prevention and elimination of epidemics, especially in low- and middle-income countries. Understanding people's choices for epidemic control thus has significant policy implications.

Any model of epidemic control must consider 1) how an infectious disease is transmitted among individuals and 2) how individuals respond to the threat of being infected. To capture these two components, we construct our epidemic model based on two key building blocks. The first block is the dynamic epidemiological model of Bauch and Earn (2004), which captures the contagion and elimination of infectious diseases. The second block is a vaccination game, a nonlinear public good game in which costly commitments (vaccination) are required of a fraction of the population to reach the critical level required for herd immunity, without which defectors are punished by the natural contagion of epidemics.[2,3]

More precisely, our work builds critically on the nonlinear nature of vaccination externality, the foundation of which is a class of well-established epidemiological models. The susceptible-infected-recovered (SIR) framework, introduced by Kermack and McKendrick (1927), describes the transmission of infectious disease through individuals in a society with a fixed population that consists of three compartments: susceptible, infected, and recovered. Bauch and Earn (2004) extended the model to incorporate vaccination with fixed timing and revealed that the long-run probability of infection exhibits an interesting property: it decreases concavely as vaccination

[1] The word "epidemic" originates from the Greek "epidemos," which literally means among (epi) the people (demos), and it refers to the rapid spread of infectious disease to a large number of individuals in a given population within a short period, usually two weeks or less (Cambridge Dictionary, 2015).

[2] More precisely, we derive the steady-state probability of infection from the dynamic epidemic transmission model of Bauch and Earn (2004) and consider forward-looking individuals who simultaneously and independently make (strategic) decisions on whether to receive a vaccine on day 1 by considering the threat of being infected as the steady-state probability of infection. Thus, our model can be regarded as a steady-state model of dynamic contagion with static strategic decision making.

[3] Galeotti and Rogers (2013) also considered a dynamic epidemic model of the susceptible-infected-susceptible (SIS) framework to study the diffusion process of a harmful state and combine it with a strategic decision-making problem of individuals exposed to the diffusion process. See the literature review section for further discussion.

coverage increases and eventually vanishes when a critical mass of the population obtains vaccinations.[4] The minimum fraction of a population that must be vaccinated to prevent an outbreak is called "herd immunity."

We explore how the nonlinear nature of this externality affects the social planner's problem and individuals' vaccination choice problem. We first show that, due to the nonlinearity of the externality, the marginal social cost of vaccination becomes discontinuous at the vaccination coverage required to obtain herd immunity. As a result, eliminating epidemics by obtaining herd immunity is a necessary and sufficient condition for the social optimum.[5] This result is new to the economic literature on epidemic transmission and control, the previous conclusion of which was that social optimum should be achieved when an epidemic is not eliminated (e.g., Brito, Sheshinski and Intriligator, 1991; Francis, 1997; Geoffard and Philipson, 1997; Gersovitz and Hammer, 2004; Kureishi, 2009).

We then turn to investigating the outcome of strategic vaccination when the probability of infection is endogenously determined by vaccination coverage.[6] Each player faces a trade-off between the nonlinear probabilistic benefit of vaccination and the deterministic cost of vaccination, although vaccination behavior benefits society as a whole. The nonlinear public goods game, in which forward-looking individuals make vaccination choices to optimize their health assets to guard against epidemics, with the steady-state probability of infection derived from the dynamic epidemic contagion model of Bauch and Earn (2004), is in our formulation the "vaccination game." We show that there is a strategic advantage from being unpredictable in immunization, which could lead individuals to randomize accepting vaccination. In this equilibrium, an increase in the relative benefit of vaccination results in unambiguously higher vaccination coverage and, consequently, a higher likelihood of herd immunity.

That the equilibrium likelihood of herd immunity monotonically increases with the relative benefit of vaccination is vital in predicting whether voluntary vaccinations can be an effective mechanism to achieve the social optimum. From a social planner's perspective, it indicates that voluntary vaccination can approximate elimination of any disease if the relative benefit of vacci-

[4]This modeling choice is especially appropriate when individuals are assumed not to be informed about the prevalence of the disease.

[5]Since Bauch and Earn (2004), the nonlinear nature of this externality has been well understood in the epidemiological literature. However, to our knowledge, we are the first to investigate how the nonlinear nature of this externality changes the social planner's problem.

[6]Heal and Kunreuther (2005) and Galeotti and Rogers (2013) were pioneers in adopting a game-theoretic approach to explore the interdependency of vaccination decisions. Following in their footsteps, in this paper, we extend the game-theoretic approach to incorporate the nonlinear nature of epidemic externalities. The model presented in this paper is, to the best of our knowledge, the first attempt in economics to apply a nonlinear public good game to the study of vaccination choice.

nation is sufficiently large, i.e., when the individual benefit from vaccination is sufficiently larger than the individual's utility cost of vaccination. The social planner's well-defined policy goal should be to increase the relative benefit. One can imagine several natural policy tools available to the social planner to achieve the goal, such as a subsidy to decrease the cash cost of vaccination, a nationwide public vaccination campaign to lower the psychological cost of vaccination, and financial aid to encourage R&D activities to minimize the side effects of vaccines. However, a mandatory vaccination program is *not* necessary to achieve the social optimum.

The results of the controlled laboratory experiments provide strong support for our major theoretical predictions from the mixed-strategy equilibrium. We consider a laboratory vaccination game in which there are eight individuals in a society, each of whom independently and simultaneously decides whether to be vaccinated. For each individual, the benefit of being vaccinated monotonically and nonlinearly decreases with the number of other people being vaccinated. To the best of our knowledge, we are the first to experimentally explore how the nonlinear nature of the positive prevention externality affects people's free-riding behavior in the context of a vaccination choice problem. We explore various treatments that differ in the relative benefit of vaccination and the characteristics of the disease (the reproduction ratio). The mean vaccination coverage and the likelihood of herd immunity reported in our experimental data indicate that endogenous epidemic punishment is an effective mechanism for voluntary vaccination, indicating that herd immunity can be achieved through voluntary, private vaccination.

The rest of this paper proceeds as follows. The remainder of this section reviews the related literature on the economics of vaccination and nonlinear public good games. Section 2 presents the steady-state characterization of the dynamic epidemiological model of Bauch and Earn (2004) and formally derives herd immunity and epidemic punishment. Section 3 explores the social planner's problem and reveals that eliminating an epidemic by obtaining herd immunity is socially optimal. In Section 4, we formally define the vaccination game and characterize the unique, symmetric, mixed-strategy Nash equilibrium. Section 5 describes the experimental design and procedures of the vaccination game. Section 6 presents the experimental results. In Section 7, we provide experimental evidence for the role of concavity in the long-run probability of infection in facilitating voluntary vaccination. Section 8 concludes the study by inviting further discussion of the relevance of the vaccination game.

1.1 Literature Review

Many leading contributions to the field of economic epidemiology have focused on the market failure of vaccination. Brito et al. (1991) showed that laissez-faire competitive equilibrium is not socially optimal since individuals do not consider the benefits and costs to others who will be infected or protected as a consequence of their infection or vaccination. Subsequent formal treatments reached similar conclusions by incorporating intertemporal disease transmission (Geoffard and Philipson (1997), Gersovitz and Hammer (2004), Boulier et al. (2007)).[7] A notable exception is Francis (1997), who examined a case of dynamic vaccination in which the competitive equilibrium coincided with the socially optimal solution.

In the above mainstream welfare literature on vaccination, decentralized decision making regarding vaccination considered the probability of infection as given. As noted by Heal and Kunreuther (2005), the notion of the strategic interplay among individuals deserves particular attention because vaccination decisions are interlocking, at least at a community level.[8] Our paper contributes to the emerging literature on strategic vaccination (Heal and Kunreuther (2005), Galeotti and Rogers (2013), Chen and Toxvaerd (2014)). Heal and Kunreuther (2005) adopted a static epidemic formulation, while we explicitly model the dynamic transmission of epidemics. Chen and Toxvaerd (2014) considered a two-player dynamic game of vaccination, while our formalization allows us to study the threshold implications of herd immunity for a general population and thus captures the crucial feature of collective action in vaccination problems. In addition, the prevalence elastic strategy in Chen and Toxvaerd (2014) makes it difficult to attribute the results to either the prevention externality or the infection externality.[9] Identifying the source of market failure is essential to its remediation measures. To isolate the positive externalities of vaccination, we contend that a vaccination game with inflexible timing embedded in a dynamic epidemiological context provides an ideal tool.[10]

[7]For example, Geoffard and Philipson (1997) argued that public health interventions, including price subsidies and stand-alone vaccination programs, are generally unable to remedy vaccination market failures, as the increase in demand is always offset by an increase in free riding. Note that they considered eradication, while the focus of this paper is elimination, which means the interruption of epidemics of an infectious disease, achieved by maintaining vaccination coverage to ensure that the proportion of immune individuals remains greater than the critical threshold for herd immunity. Eradication means the worldwide reduction of infective organisms in the wild to zero. To achieve eradication, elimination must be achieved in all world regions.

[8]Investigating vaccination choices in a game-theoretic framework is also critical for examinations of the welfare aspects of preventive behavior, as the individual risk of infection is directly influenced by the choices made by others in the community.

[9]A prevalence elastic strategy is an action plan contingent on how many people are infectious over the course of an epidemic. It is usually employed in models featuring flexible timing of vaccination, in which feedback about the epidemic is available. As noted by Gersovitz and Hammer (2004), the prevention externality, wherein the vaccination choice of an individual can reduce the likelihood that others are immunized, should be separated from the infection externality, as one's own infection can increase the likelihood that others are infected.

[10]Our social planner problem can also be seen as complementary to that of Gersovitz and Hammer (2004), who

This paper is closest in form to Galeotti and Rogers (2013), although we differ in focus. Galeotti and Rogers (2013) integrated strategic decision making into the susceptible-infected-susceptible (SIS) framework, in which forward-looking individuals consider the long-run infection probability.[11] Their focus was on how group structure, particularly homophily and assortative matching, affects immunization behavior, while we focus on how a nonlinear externality affects vaccination decisions and the conditions under which voluntary vaccination can be socially optimal. There are also major differences in assumptions between our model and that of Galeotti and Rogers (2013). First, unlike Galeotti and Rogers (2013), we consider the SIR framework under uniform random matching, in which a recovered individual obtains individual immunity and thus does not suffer any further infection. Second, Galeotti and Rogers (2013) considered a society with a continuum of individuals, while we consider a society with finitely many individuals.

In the economics literature, nonlinear public good games have been studied primarily to test the robustness of the typical overcontribution result observed in the standard linear public good games. Several papers, such as Keser (1996), Sefton and Steinberg (1996) and Isaac and Walker (1998), have shown that the overcontribution results are persistent in nonlinear settings with interior Nash equilibria. Laury, Walker and Williams (1999) explored how robust the overcontribution results are to changes in the information describing the payoff structure. They showed that, although experimental subjects continue to allocate more resources than the Nash prediction, providing participants with detailed descriptions of the declining marginal benefit to the public good leads to a significant decrease in the provision of public goods.

Nonlinear public good games have been also investigated in the context of common pool resources (e.g., Ostrom et al., 1994; Casari and Plott, 2003) and of the volunteer's dilemma (e.g., Diekmann, 1985; Goeree and Holt, 2000; Goeree et al., 2017).[12] In a recent contribution, Cason and Gangadharan (2015) studied behavior in linear and nonlinear social dilemma games with costly punishment opportunities and demonstrated that the impact of punishment is weaker and takes longer to be effective in a nonlinear environment. In the volunteer's dilemma initiated by Diekmann (1985), only a single volunteer's costly commitment to provide a public benefit is required, and a symmetric equilibrium usually involves mixed strategies. In this regard, the vaccination game is an extension of the volunteer's dilemma when the critical level requires a fraction of the

studied the effects of an infection externality without considering the threshold implications of herd immunity.

[11] The proper modeling choice between the SIS framework and the SIR framework depends crucially on the nature of the disease. For instance, it might be more natural to consider the SIS model for pertussis, diphtheria and AIDS/HIV among the diseases mentioned in Appendix A.

[12] Note, however, that threshold public goods, also known as lumpy public goods or provision point public goods, are *only* extreme examples of nonlinear public goods. In particular, our nonlinear public good game is *not* a threshold public good game. Threshold public goods have been widely studied in the economics literature, and we do not review this literature in this paper.

population. In fact, Goeree and Holt (2000) classified the volunteer's dilemma as a special case of a threshold public good, for which the critical level is one. Moreover, a standard provision point-discrete public good game (see Palfrey and Rosenthal, 1984) usually exhibits multiple equilibria in pure strategies, and mixing is usually not a concern. To the best of our knowledge, the vaccination game is the first to consider mixed-strategy Nash equilibria in nonlinear threshold public good games.[13]

2 Preliminaries: Epidemic Model of Bauch and Earn (2004)

A plausible vaccination choice at a fixed interval should account for the fact that the probability of catching a disease is not constant but will converge to a certain level in the long run. In this section, we present a model of dynamic epidemic transmission based on the analysis of Bauch and Earn (2004), who described the interaction of different health compartments of the population. As we shall see, the vaccination choice of each agent gives rise to the vaccination coverage of a community, which eventually determines the probability that an unimmunized agent is infected.

2.1 Population Dynamics

Bauch and Earn (2004)'s formulation was based on the well-established SIR model in mathematical epidemiology (refer to, e.g., Kermack and McKendrick, 1927; Geoffard and Philipson, 1997). In the classical model, susceptible individuals in a population interact with one another and are exposed to an infectious disease. Once infected, an individual shifts from the susceptible to the infected state. Upon infection, an individual has a chance of making a spontaneous recovery and transitioning to the recovered state, where he/she remains thereafter.

Define $S(t)$, $I(t)$, and $R(t)$ as the number of susceptible, infectious and recovered individuals, respectively, at time t, which exhaustively and mutually exclusively categorize a population of size $N \gg 0$. Instantaneously, there is probability β of infectious contact between a susceptible individual and an infectious individual.[14] We introduce vaccination in such a way that immunized agents remain in the recovered state, while unvaccinated agents automatically enter the susceptible state. The dynamics of the SIR model thus evolve as the following nonlinear ODE

[13] In their experimental investigation of the volunteer's dilemma, Goeree et al. (2017) focused on the symmetric mixed-strategy equilibrium. Dixit and Olson (2000) considered a mixed-strategy equilibrium in a two-stage lumpy public good game, but their focus was on whether voluntary participation undermines the Coase theorem.

[14] Note that infectious contact can only occur between a susceptible individual and an infectious individual.

system:

$$\frac{dS(t)}{dt} = \mu N(1-P) - \frac{\beta S(t) I(t)}{N} - \mu S(t) \tag{1}$$

$$\frac{dI(t)}{dt} = \frac{\beta S(t) I(t)}{N} - \gamma I(t) - \mu I(t) \tag{2}$$

$$\frac{dR(t)}{dt} = \mu N P + \gamma I(t) - \mu R(t), \tag{3}$$

where, for any given t, $S(t) + I(t) + R(t) = N$, γ is the probability of recovery and μ indicates the birth rate and the death rate ($1/\mu$ is the expected life expectancy).[15] P denotes the vaccination coverage in the population, which is a public choice component added to the original model to fit settings in which a new medication is introduced or in which new members are introduced into the community (e.g., infants, immigrants).[16] It is noteworthy that, in our setting, the vaccination term has been defined independently of the time trend of prevalence.

More precisely, (1) specifies the rate of change in the number of susceptible individuals due to the entry of nonvaccinated members and exit due to new infections and non-disease-related mortality; (2) specifies the rate of change in prevalence due to the entry of infectious individuals, while exits are caused by post immunity and mortality; and (3) specifies the rate of change in the number of recovered individuals due to the entry of vaccinated members and those who recover from infection, with exits purely due to mortality.

We rewrite the nonlinear ODE system in dimensionless form:[17]

$$\begin{aligned} \frac{d\hat{S}}{dt} &= \mu(1-P) - \beta \hat{S}\hat{I} - \mu \hat{S} \\ \frac{d\hat{I}}{dt} &= \beta \hat{S}\hat{I} - \gamma \hat{I} - \mu \hat{I} \\ 1 &= \hat{S} + \hat{I} + \hat{R}, \end{aligned} \tag{4}$$

where $\hat{S}$, $\hat{I}$, and $\hat{R}$ respresent the respective proportion of the population.

Define P_{crit} as the vaccination coverage needed for herd immunity (after which the entire population will be safe from an epidemic). To prevent a disease from becoming an epidemic, we

[15] Following Geoffard and Philipson (1997), we assume that infection does not induce death since we are not interested in the effects of population change. See Keeling and Rohani (2008) for SIR models that incorporate infection-induced mortality. When disease-induced mortality is added to the SIR model with density-dependent transmission, the equilibrium and stability property simply reflect a change in parameters (Chapter 2, pages 24-25).

[16] We focus on the case of perfect vaccination and no virus evolution, although it is straightforward to extend our model to incorporate imperfect vaccination and virus evolution. Our main qualitative result does not depend on the perfect vaccination assumption.

[17] Note that Equation (3) is redundant after normalization. We further assume that the initial condition $\hat{R}(0) = 0$, $\hat{I}(0)$ is small but nontrivial, and thus $\hat{S}(0) \approx 1$.

should at least have no increase in the proportion of infected individuals, i.e.,

$$\frac{d\hat{I}}{dt} = 0,$$

which implies that

$$1 - P_{crit} = \frac{\gamma + \mu}{\beta}.$$

It follows that

$$P_{crit} = 1 - \frac{1}{R_0}, \tag{5}$$

where $R_0 = \frac{\beta}{\gamma+\mu}$ is the basic reproduction ratio of the disease. Note that $R_0 > 1$ for any epidemic (see Anderson and May (1992) for a detailed discussion).[18] As listed in Appendix A, the value of R_0 varies with the disease. In the subsequent analysis, we use R_0 as a general characteristic that differentiates epidemics.

Definition 1 (Herd Immunity)**.** *A community achieves **herd immunity** to a disease if a proportion of the population weakly more than P_{crit} is vaccinated in the steady state.*

2.2 Long-run Behavior of the Nonlinear Differential System

We denote the steady state of the system as a profile (S^*, I^*, R^*). By setting $\frac{d\hat{S}}{dt} = 0$ and $\frac{d\hat{I}}{dt} = 0$, we obtain (S^*, I^*, R^*):

$$(1 - P, 0, P) \text{ if } P \geqslant P_{crit}, \tag{6}$$

$$(\frac{1}{R_0}, \frac{\mu}{\gamma + \mu}(1 - P) - \frac{\mu}{\beta}, 1 - \frac{1}{R_0} - \frac{\mu}{\gamma + \mu}(1 - P) + \frac{\mu}{\beta}) \text{ if } P < P_{crit}. \tag{7}$$

It is readily seen that the system has a "good" steady state and a "bad" steady state. Steady state (6) is seemingly more socially desirable than steady state (7), as no infection occurs in convergence to the former. In the subsequent welfare analysis, we verify that this relationship is indeed the case. Vaccination coverage P plays an indispensable role here, as the steady state toward which the system will converge depends crucially on P.

Figure 1(a) plots the direction field of $\hat{S}$ and $\hat{R}$ and the contour line of $\hat{I}$ as a function of $\hat{S}$ and $\hat{R}$ when $P = 90\%$, while Figure 1(b) reproduces the graph with the much lower vaccination coverage of $P = 10\%$ as the only parameter change. We set $\beta = 1$, $\gamma = 0.1$, and $\mu = 0.01$, which results in $R_0 = 9.1$ and $P_{crit} = 89\%$. The difference between the two figures illustrates the distinct

[18] Barrett (2003) and Kureishi (2009) derive the same result in their static models.

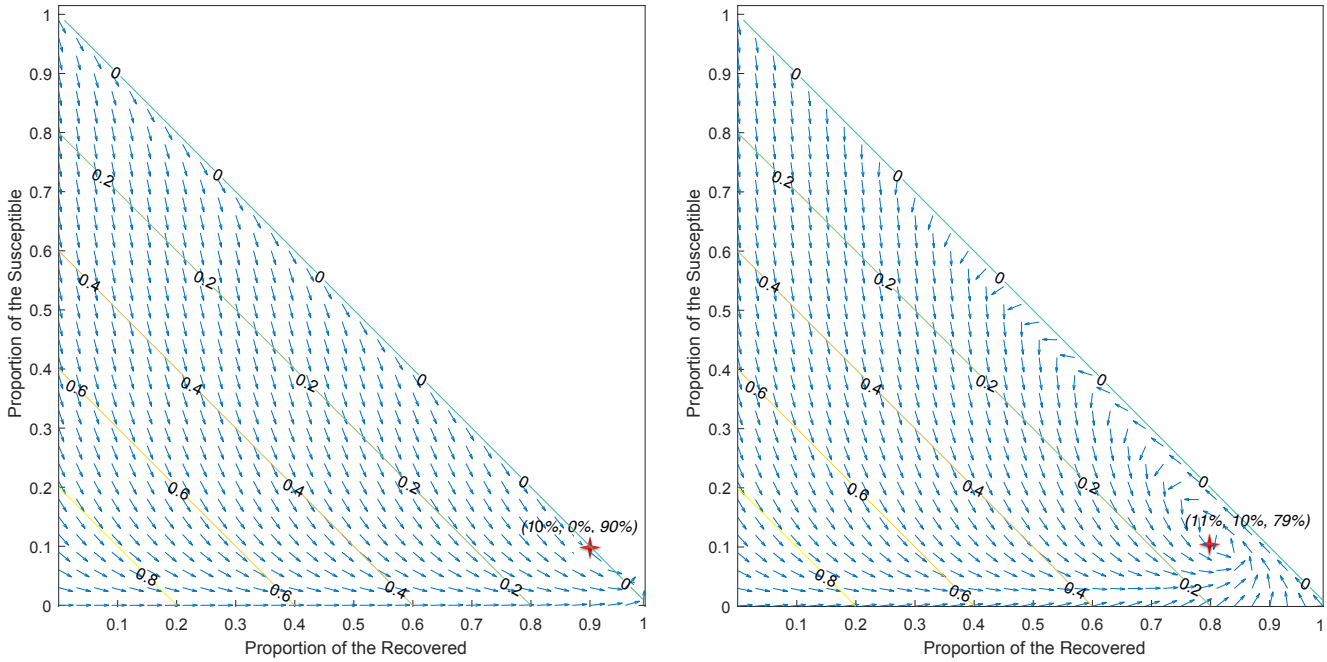

(a) Herd Immunity

(b) Without Herd Immunity

Figure 1: Direction Field and Contour Line of the SIR disease

steady states with and without herd immunity. In Figure 1(a), starting from any arbitrary initial point in the upper-left corner, the epidemic system will eventually move toward the steady state (6) marked by the cross (10%, 0, 90%), which lies on the diagonal line $\hat{I} = 0$ (i.e., zero prevalence). In comparison, given an arbitrary plausible initial point, the epidemic system in Figure 1(b) will always move to a stable node corresponding to the steady state (7) marked by the other cross (11%, 10%, 79%), making elimination impossible.

The implications of the steady states are far reaching: one of them is to determine the long-run infection probability. Define π_P as the probability that an unvaccinated individual will eventually be infected, with the vaccination coverage of the population being P. We simply call π_P the long-run probability of infection. The characterization of π_P leads us to the following proposition.

Proposition 1. *[Bauch and Earn (2004)] Given any $P \in [0,1]$, there exists a unique π_P that is strictly decreasing and concave in P until P reaches the elimination threshold P_{crit}. Furthermore, $\pi_P = 1 - \frac{1}{R_0(1-P)}$ for any $P < P_{crit}$, and $\pi_P = 0$ for any $P \geqslant P_{crit}$.*

Proof. See Appendix B.

Figure 2, directly adopted from Figure 1 in Bauch and Earn (2004), plots the long-run probability of infection π_P as a function of the vaccination coverage P for different types of diseases. For all types of diseases, the threat of the disease decreases as more people choose to immunize, eventually vanishing when vaccination coverage reaches the critical level P_{crit}. Moreover, for vac-

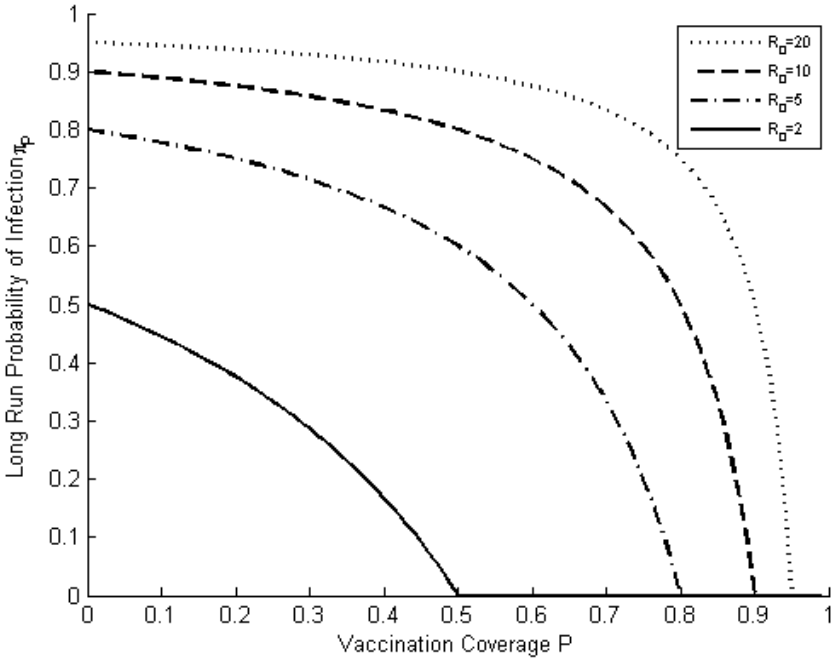

Figure 2: long-run Probability of Infection

cination coverage less than the critical level, a more contagious disease indicates a greater threat of infection. This figure thus visualizes the nonlinear, threshold properties of the endogenously determined vaccination externality identified by the dynamic epidemic process.

Note that the concavity of the long-run probability of infection π_P originates from the standard SIR model implicitly assuming uniform random matching between individuals in S and I or, equivalently, a complete contact network.[19] Thus, the disease transmission ratio $\beta SI/N$ satisfies the property of increasing returns to scale, and the individual externality of vaccination increases with respect to the population vaccination coverage. This type of phenomenon is called a "positive network externality" (refer to, e.g., Katz and Shapiro, 1985; Farrell and Saloner, 1985). Obviously, the positive network externality comes into play in any regular graph such that the concavity of π_P is a natural and unavoidable property of epidemic transmission.

[19]Epidemic models with richer interaction structures were studied by Anderson and May (1992) and Sattenspiel and Simon (1988). Network structure was considered more explicitly by Kretzschmar and Morris (1996) and Pastor-Satorras and Vespignani (2001a,b, 2002). For a comprehensive survey, see Jackson (2010), Draief and Massoulié (2010) and Just et al. (2015).

3 Social Planner's Problem

The notion of herd immunity and the nonlinear nature of externality make it necessary to redefine the social planner's problem. In this section, we show that the elimination of epidemics is a necessary and sufficient condition for the socially optimal vaccination coverage when the notion of herd immunity and the concavity of the long-run infection probability are considered.

For each player i, there are three possible health statuses: susceptible (S), infected (I), and recovered (R). For any health status $\theta \in \{S, I, R\}$, $u(\theta)$ denotes the instantaneous utility from the state θ. Assume that

$$u(S) = u(R) > u(I), \tag{8}$$

such that a player values the uninfected status qualitatively more than the infected status. Define the utility cost of infection as $L = u(R) - u(I)$.

Let $U_{nv}(P)$ denote a schedule representing the expected utility of an unvaccinated individual given the vaccination coverage P in the steady state. Then,

$$U_{nv}(P) = \begin{cases} \frac{1}{\mu}u(R) & \text{if } P \geqslant P_{crit} \\ \frac{1}{\mu}u(S) - \pi_P \cdot d_{R_0} \cdot L & \text{if } P < P_{crit}, \end{cases}$$

where $d_{R_0} > 0$ is the duration of the infection that is uniquely determined by R_0, a characteristic of the disease. $1/\mu$ is the life expectancy.

Let $U_v(P)$ denote a schedule representing the actual utility of a vaccinated individual *inclusive of vaccination cost* in the steady state. Then, we have

$$U_v(P) = \frac{1}{\mu}u(R) - C(P),$$

where $C(P)$ is the P-percentile individual's utility cost of vaccination that can reflect a combination of cash costs, psychological costs and possible side effects. Note that this cost is the only individual heterogeneity allowed in the model considered in this section.[20]

In an epidemic context, social welfare is defined as the aggregate health assets of the society

[20]To render our analysis comparable to the findings in the literature, we focus on a quasisymmetric environment allowing for the heterogenous vaccination costs in this section. However, we consider the case of homogeneous vaccination costs to provide a full characterization of equilibria in Section 4. The main result of this section – the optimality of herd immunity – does not depend on whether the vaccination cost is heterogenous or homogenous.

(i.e., utilitarian). Formally, the social welfare function $W(P)$ can be written as

$$W(P) = \int_0^P U_v(\hat{P})d\hat{P} + U_{nv}(P)(1-P).$$

The social optimum can be found by solving the social planner's problem of maximizing the social welfare function $W(P)$. From the first-order condition, we have

$$\underbrace{U_v(P)}_{MSB(P)} = \underbrace{U_{nv}(P)}_{MPC(P)} - \underbrace{\frac{dU_{nv}(P)}{dP}(1-P)}_{Externality} \equiv MSC(P) \tag{9}$$

The left-hand side of Equation (9) represents the marginal social benefit (MSB) of vaccination. The first term on the right-hand side of Equation (9) represents the marginal *private* cost (opportunity cost) of vaccination, while the second term represents the prevention externality of vaccination. The combination of these two terms, denoted MSC, is thus the marginal *social* cost of vaccination. The social optimum is achieved exactly when the marginal social benefit of vaccination is the same as its marginal social cost.

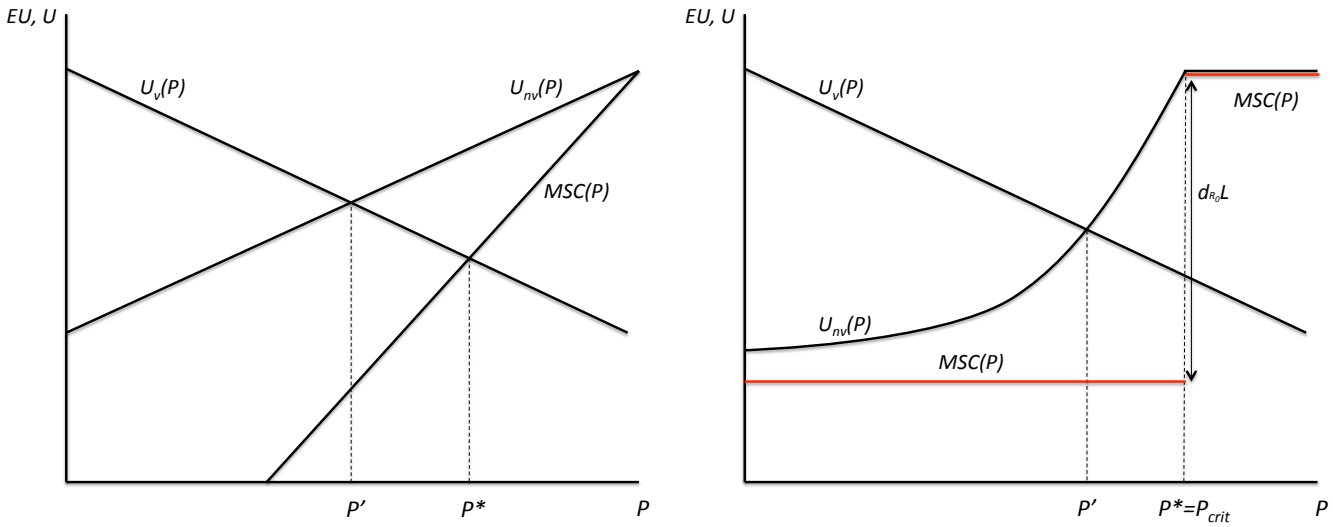

(a) Brito et al. (1991) (b) Our case

Figure 3: Social Planner's Problem

Figure 3(a) illustrates how Brito et al. (1991) established their main result that market equilibrium, denoted as P', cannot achieve the socially optimal outcome, denoted as P^*, the solution obtained by the social planner's problem.[21] To achieve this result, a few assumptions are required. First, $U_{nv}(P)$ is increasing and linear. Second, a lower-cost person is always vaccinated before a higher-cost person so that $U_v(P)$ is nonincreasing in P.

[21] The figure is adopted from Figure 2 in Francis (1997).

Figure 3(b) demonstrates how the social planner's problem is changed when the notion of herd immunity and the concavity of the long-run probability of infection are considered. Our result in Proposition 1 implies that $U_{nv}(P)$ is increasing and convex for any $P < P_{crit}$ and constant for any $P \geqslant P_{crit}$. Consequently, the shape of the $MSC(P)$ curve drastically changes, as depicted by the red curve in Figure 3(b). Specifically, for any $P < P_{crit}$, the right-hand side of Equation (9) becomes

$$\begin{aligned} MSC(P) &= \frac{1}{\mu}u(S) - \pi_P d_{R_0} L - \frac{1}{R_0(1-P)^2} d_{R_0} L(1-P) \\ &= \frac{1}{\mu}u(S) - \pi_P d_{R_0} L - (1-\pi_P) d_{R_0} L \\ &= \frac{1}{\mu}u(S) - d_{R_0} L, \end{aligned}$$

and $MSC(P) = U_{nv}(P)$ for any $P \geqslant P_{crit}$. Observe that $MSC(P)$ does not depend on P when $P < P_{crit}$. This result is driven primarily by the concavity of the long-run infection probability. The concavity of the long-run infection probability implies that, when P increases, an individual not only encounters a higher opportunity cost of vaccination, but it also creates a stronger prevention externality.[22] The higher marginal private cost of vaccination is canceled out by the stronger externality, resulting in a constant marginal social cost of vaccination.

Hence, for any $P < P_{crit}$ we have

$$MSB(P) - MSC(P) = \frac{1}{\mu}u(R) - C(P) - \frac{1}{\mu}u(S) + d_{R_0} L = d_{R_0} L - C(P),$$

and for any $P \geqslant P_{crit}$ we have

$$MSB(P) - MSC(P) = \frac{1}{\mu}u(R) - C(P) - \frac{1}{\mu}u(S) = -C(P) < 0.$$

Note that, for any $P < P_{crit}$, $U_{nv}(P) > MSC(P)$ because, at $P = 0$,

$$U_{nv}(P) = \frac{1}{\mu}u(S) - (1 - \frac{1}{R_0}) d_{R_0} L > \frac{1}{\mu}u(S) - d_{R_0} L = MSC(P).$$

Therefore, unless the society has a significant proportion of individuals ($> 1 - P_{crit}$) whose vaccination costs are strictly greater than $d_{R_0} L$, the marginal social benefit of vaccination strictly dom-

[22] When $U_{nv}(P)$ is linear, as illustrated in Figure 3(a), the degree of the externality is linearly decreasing in P because the externality depends solely on the proportion of the unvaccinated $(1-P)$. However, when $U_{nv}(P)$ is convex, Equation (9) shows that the degree of the externality is jointly determined by the proportion of the unvaccinated and the slope of $U_{nv}(P)$. The effect of the slope dominates that of the proportion of the unvaccinated. As a result, the degree of the externality is monotonically increasing in P.

inates the marginal social cost of it for any $P < P_{crit}$. This result is summarized in the following proposition.

Proposition 2. *Assume that* $C(P_{crit}) \leqslant d_{R_0}L$. *Then,*

(1) $P' < P^*$*; i.e., the market equilibrium cannot achieve the socially optimal outcome; and*

(2) $P^* = P_{crit}$*; i.e., the socially optimal outcome is achieved if and only if herd immunity is achieved.*

Three remarks are in order. First, we consider a society that consists of a continuum of individuals (so that P is a continuous variable) in this section for ease of comparison between our result and the result presented in Brito et al. (1991). However, the optimality of the herd immunity presented in Proposition 2 is robust to the environment with a finite number of individuals. Second, this result reveals that achieving herd immunity is socially optimal unless the society has a significant proportion of individuals whose vaccination costs are unreasonably high. Thus, in the remainder of our paper, we focus on the case in which achieving herd immunity is socially optimal, and we address whether voluntary vaccination can achieve the elimination of epidemics. Third, the market equilibrium P' is a traditional solution obtained by assuming that individuals consider the vaccination coverage P as given when making their vaccination choices. In the next section, we consider a strategic vaccination choice problem among individuals in the society in which vaccination coverage P is endogenously determined. We show that this strategic interdependency can completely change the prediction.

4 Vaccination Game and Equilibrium Analysis

In this section, we present a game-theoretic model of vaccination choices by incorporating the main result of the epidemic dynamics presented in Section 2. The long-run infection probability π_P generates an endogenous punishment for free-riders, i.e., unvaccinated individuals.

4.1 Vaccination Game

Consider a society with a finite set of players $\mathcal{N} := \{1, 2, \ldots, n\}$ whose preferences satisfy the Von Neumann-Morgenstern axioms such that they maximize expected utility.[23] Each player $i \in$

[23] To be consistent with the epidemic model, let $n = \mu N$ denote an integer-valued number of players. Risk preferences can alter the benefits of vaccination; altruism can reduce the costs of vaccination. However, in either case, the decision rule for the binary choice would not be affected, and the same result applies.

$\mathcal{N}$ simultaneously and independently makes a vaccination choice b_i in $\mathcal{B}_i := \{vc, nv\}$, where vc and nv refer to vaccination and nonvaccination choices, respectively. For any player i, let $b_{-i} = (b_1, ...b_{i-1}, b_{i+1}, ..., b_n)$ denote the pure-strategy profile of other players. For a given pure-strategy profile $\boldsymbol{b} = (b_1, ..., b_n) = (b_i, b_{-i})$, let $V(\boldsymbol{b}) = \{i \in \mathcal{N} | b_i = vc\}$ denote the set of vaccinated players. The proportional vaccine coverage of the population P is endogenously determined as follows:

$$P(\boldsymbol{b}) = |V(\boldsymbol{b})|/n,$$

where $|V(\boldsymbol{b})| = \nu$ denotes the cardinality of the set $V(\boldsymbol{b})$. Let ν_{crit} denote the smallest integer value greater than $P_{crit}n$.

For a given P, let π_P denote the long-run infection probability. As described in Proposition 1, π_P is strictly decreasing and concave in P for any $P < P_{crit}$, and $\pi_P = 0$ for any $P \geqslant P_{crit}$. $\mathcal{E} = \{\beta, \gamma, \mu\}$ describes the unique characteristics of the nature of the epidemic. We consider forward-looking players who consider the steady-state, long-run probability of infection π_P when making their vaccination decisions b_i.[24]

For each player i, there are three possible health statuses *ex post*: susceptible (S), infected (I), and recovered (R). For any health status $\theta \in \{S, I, R\}$, $u(\theta)$ denotes the instantaneous utility from the state θ. Assume that

$$u(S) = u(R) > u(I), \tag{10}$$

such that a player values the uninfected status qualitatively more than the infected status.[25] Define the utility cost of infection as $L = u(R) - u(I)$.[26]

For any player i, the expected utility of not vaccinating given a fixed strategy profile of other players b_{-i} (and thus for a fixed vaccine coverage P) is as follows:

$$\mathcal{U}_i(nv, b_{-i}) = \begin{cases} \frac{1}{\mu}u(R) & \text{if } P(nv, b_{-i}) \geqslant P_{crit} \\ \frac{1}{\mu}u(S) - \pi_{P(nv,b_{-i})} \cdot d_{R_0} \cdot L & \text{if } P(nv, b_{-i}) < P_{crit}, \end{cases}$$

where $d_{R_0} > 0$ is the duration of the infection uniquely determined by R_0, a characteristic of the disease. $d'_{R_0} \geqslant 0$ because a larger R_0 indicates a lower possibility of recovery.

[24] We aim at modeling a situation in which epidemic transmission occurs in a short period of time, e.g., two weeks or less, allowing us to focus on the steady-state payoffs and outcomes with no time discounting.

[25] We acknowledge that u(R) can be larger than u(S), and the same conclusion applies in this case but is considerably more notationally burdensome.

[26] As one might have already noted, the model and notations we have in this section are consistent with those in Section 3. We restate them to render this section self-contained.

For any player i, the expected utility of vaccination for any b_{-i} is

$$\mathcal{U}_i(vc, b_{-i}) = \mathcal{U}_i(vc) = \frac{1}{\mu}u(R) - C,$$

where C is the utility cost of vaccination that can reflect a combination of cash costs, psychological costs and possible side effects. All of these costs are common knowledge.[27]

We call the simultaneous-move game defined by the tuple $\{\mathcal{N}, \{\mathcal{B}_i\}_{i\in\mathcal{N}}, \{\mathcal{U}_i\}_{i\in\mathcal{N}}\}$ the vaccination game $\mathcal{G}$. In what follows, we focus on the case in which $C \leqslant \pi_0 d_{R_0} L$ because our welfare analysis shows that herd immunity is the social optimum (i.e. $P^* = P_{crit}$) if and only if that condition is satisfied. If $C > \pi_0 d_{R_0} L$, then the only equilibrium outcome is zero vaccination coverage.

4.2 Mixed-strategy Equilibria

Can epidemics be eliminated by voluntary vaccinations? It is clear that there is a spectrum of pure-strategy equilibria in which precisely $n \cdot P_{crit}$ players choose vaccination, and the remainder choose to free ride. Consequently, the answer should be yes if one focuses on the pure-strategy equilibria. Nonetheless, the asymmetry of these equilibria is undesirable because they all arbitrarily require identical players to choose different strategies in a precisely coordinated manner.[28] In this section, we thus answer this question by focusing on unique symmetric mixed-strategy equilibrium. The full characterization of Nash equilibrium outcomes is presented in Appendix B.

Before presenting a formal derivation, we discuss why an individual would use a mixed strategy for a binary immunization choice in the simplest three-player setting when one of them has already been infected, and the vaccination of only one susceptible individual suffices to obtain herd immunity in the society. If the other susceptible individual has not been vaccinated for certain, an individual should protect himself/herself from infection unless the vaccination cost is unreasonably high; however, if the other individual is certain to vaccinate, an individual would choose to free ride. Therefore, there is a strategic advantage of being unpredictable.

Let $\sigma_i \in [0,1]$ denote the probability that player i chooses vaccination. $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_n)$ denotes a mixed-strategy profile. The expected payoff for player i from randomization with σ_i

[27] Following Heal and Kunreuther (2005) and Galeotti and Rogers (2013), we focus on the case of homogenous vaccination cost, which enables us to provide a full characterization of equilibria in the next subsection. However, our main result – the existence of a symmetric mixed-strategy equilibrium in which herd immunity is approximated as the relative benefit of vaccination increases – does not depend on the homogenous vaccination cost assumption. Also note that $\mathcal{U}_i(nv, b_{-i})$ and $\mathcal{U}_i(vc, b_{-i})$ correspond to $U_{nv}(P)$ and $U_{vc}(P)$, respectively, defined in Section 3.

[28] See Dixit and Olson (2000) for a detailed discussion of this coordination issue. For the same reason, Goeree et al. (2017) focused on the symmetric mixed-strategy equilibrium in their experimental investigation of the volunteer's dilemma.

can be expressed as follows:

$$EU_i(\sigma_i, \sigma_{-i}) = \frac{u(R)}{\mu} - \sigma_i C - (1 - \sigma_i) d_{R_0} L \mathbb{E}[\pi_{P(\sigma)}],$$

where $\mathbb{E}[\pi_{P(\sigma)}]$ denotes the expected infection probability given the mixed-strategy profile $\boldsymbol{\sigma}$.

Definition 2. *A strategy profile $\boldsymbol{\sigma}^* = (\sigma_1^*, \ldots, \sigma_n^*) \in [0,1]^n$ is a totally mixed-strategy Nash equilibrium for the game $\mathcal{G}$ if we have for any $i \in \mathcal{N}$*

$$\sigma_i \in (0,1), \tag{11}$$

and for all $\sigma_i^* \in [0,1]$,

$$EU_i(\sigma_i^*, \sigma_{-i}^*) \geqslant EU_i(\sigma_i, \sigma_{-i}^*). \tag{12}$$

The following proposition implicitly characterizes the unique totally mixed-strategy Nash equilibrium of the vaccination game $\mathcal{G}$, the uniqueness of which is proved in Appendix B.

Proposition 3. *Let $\frac{L}{C} = r \in (\frac{R_0}{(R_0-1)d_{R_0}}, +\infty)$ be the relative benefit of vaccination, and $n > R_0$. There exists a unique, totally mixed-strategy equilibrium, where $\sigma_i^* = \sigma^*$ and is implicitly defined by*

$$\frac{1}{d_{R_0} r} = \sum_{k=0}^{\nu_{crit}} \left(1 - \frac{1}{R_0} - \frac{k}{n}\right) \binom{n}{k} \sigma^{*k} (1 - \sigma^*)^{n-1-k}. \tag{13}$$

Proof. See Appendix B.

This proposition reveals that our vaccination game has no asymmetric, totally mixed-strategy Nash equilibrium. However, as noted earlier, there is a whole collection of asymmetric pure-strategy equilibria, in which precisely $n \cdot P_{crit}$ players choose vaccination, and the remainder free ride. This outcome conflicts with the noncooperative nature of our vaccination game because identical players are required to behave differently but in a precisely coordinated manner. Therefore, we focus on the unique symmetric mixed-strategy equilibrium, also considered by Palfrey and Rosenthal (1984), Dixit and Olson (2000), and Goeree et al. (2017). We determine later whether any of the asymmetric pure-strategy equilibria could explain the observed laboratory behavior well.

4.3 Equilibrium Vaccination Coverage

In this section, we focus on the mixed-strategy equilibrium and ask the following two questions. First, when is it possible that players contribute sufficiently for the society to achieve herd immunity? Second, are the players more likely to reach immunity if they are faced with a more threatening epidemic?

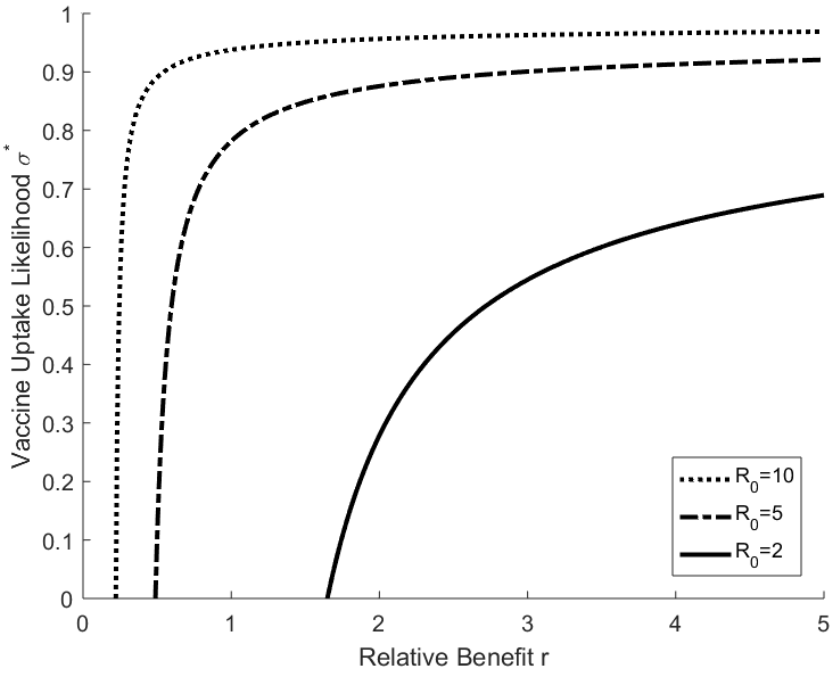

Figure 4: Mixed Strategy and Herd Immunity

Assuming a linear form of duration $d_{R_0} = 0.5R_0$ and $n = 200$, Figure 4 plots the equilibrium likelihood of vaccine uptake σ^* as a function of r for different types of diseases. It illustrates that the individual vaccination uptake likelihood $\sigma^*(r, R_0)$ is strictly increasing in r and converges to 1 as r approaches infinity. That is, if the cost of vaccination is far less than the loss of contracting the disease, people are almost certain to opt for vaccination. In the extreme case of a fatality, the substantially high relative benefit r is not only effective in enforcing voluntary vaccinations but also deters free riders who leave their fates to others.

Let $\mathcal{P}^*(r, R_0)$ denote the probability distribution over the vaccination coverage induced in the mixed-strategy equilibrium. The following proposition shows that an increase in the relative benefit r leads to an equilibrium distribution yielding an unambiguously higher vaccination coverage.

Proposition 4. *For any $R_0 > 1$, $\mathcal{P}^*(r, R_0)$ first-order stochastically dominates (FOSD) $\mathcal{P}^*(r', R_0)$ if and only if $r > r'$.*

Proof. See Appendix B.

The stochastic dominance presented in Proposition 4 implies that $Pr^*(P \geqslant P_{crit})$, the equilibrium probability for the society to achieve the vaccination coverage needed to obtain herd immunity, is monotonically increasing in r and converges to 1 as r approaches infinity. By nature of mixed-strategy equilibria, it is impossible to obtain $Pr^*(P \geqslant P_{crit}) = 1$ and obliterate an epidemic. However, the society can still approximate the complete immunity via voluntary vaccination. The following corollary summarizes this discussion.

Corollary 1. $\Pr^*(P \geqslant P_{crit}) \to 1$ *as* $r \to \infty$.

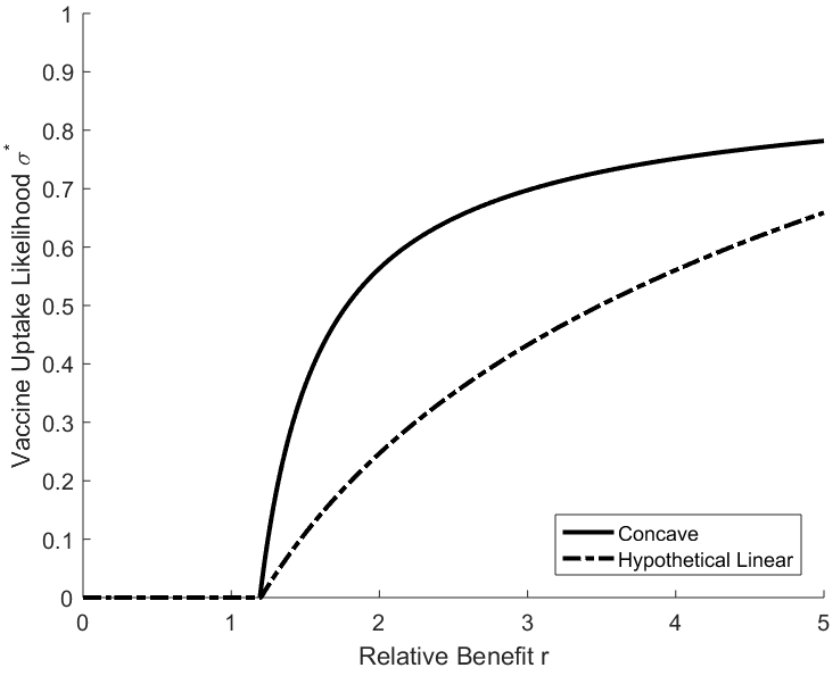

Figure 5: Elimination Thresholds in Concave vs. Linear Environments

We now investigate the role of the concavity of the long-run infection probability in facilitating the elimination of an epidemic. The concavity of π_P implies that the vaccination uptake likelihood σ^* from the mixed-strategy Nash equilibrium is strictly increasing and *concave* in the relative benefit r, as illustrated in Figure 5. Figure 5 also presents the vaccination uptake likelihood based on π_P^L, a counterfactual, linear approximation of π_P. In this hypothetical "linearized" environment, the vaccination uptake likelihood becomes less concave in r. It is thus more difficult to achieve herd immunity via voluntary vaccination.[29] Let $\mathcal{P}^L(r, R_0)$ be the equilibrium vaccination coverage in the linearized environment with $\pi_P^L = 1 - \frac{1}{R_0} - P$ for any $P < P_{crit}$, and $\pi_P^L = 0$ for any $P \geqslant P_{crit}$. Then, we have the following result.

[29] We design experiments to investigate this issue and present supporting evidence in Section 7.

Proposition 5. *For any* $r > \frac{R_0}{(R_0-1)d_{R_0}}$ *and* $R_0 > 1$, $\mathcal{P}^*(r, R_0)$ *FOSD* $\mathcal{P}^L(r, R_0)$.

The results presented in Propositions 4 and 5 have an important policy implication. Proposition 4 implies that the social planner should have a well-defined policy goal of increasing the relative benefit r. Proposition 5 shows that targeting a higher relative benefit r is particularly effective due to the concave nature of the externality. There are a few ways to achieve a high relative benefit r. First, a subsidy should be an effective policy tool to lower the vaccination cost C. Second, a nationwide vaccination campaign can also decrease the psychological cost of vaccination. Third, financial aid can be provided to encourage R&D activities to develop new vaccines with fewer and weaker side effects. However, the propositions imply that no mandatory vaccination program is needed to achieve the socially optimal outcome, especially when the network externality is strong in the society, and the nature of the externality is concave.

We now address our second question. An exogenous increase in the reproduction ratio R_0 has two competing effects on how easily the society can achieve herd immunity via voluntary vaccination. On the one hand, it increases the long-run probability of infection π_P, as illustrated in Figure 2, indicating that individuals in the mixed-strategy Nash equilibrium are more likely to vaccinate. On the other hand, a higher R_0 also increases the critical level needed for herd immunity P_{crit}. Figure 4 shows that the first effect dominates the second effect so that it is easier to achieve herd immunity when R_0 is higher. This result is summarized in the following proposition.

Proposition 6. *For any* $r \in (\frac{R_0}{(R_0-1)d_{R_0}}, +\infty)$, $\mathcal{P}^*(r, R_0)$ *FOSD* $\mathcal{P}^*(r, R_0')$ *if and only if* $R_0 > R_0'$.

Proof. See Appendix B.

This result implies that a more contagious disease is unambiguously easier to manage. A higher R_0 encourages people to be vaccinated voluntarily. Hence, epidemics like Ebola, with substantially low R_0, are particularly difficult to control based on voluntary vaccination.

5 Experimental Design

We now present our experimental design to test the major theoretical predictions from the unique symmetric/totally mixed-strategy Nash equilibrium shown in the previous section. There are a few reasons that it is necessary to offer supporting evidence for our theoretical findings. First, there is a trivial equilibrium selection issue in our vaccination game such that, without experimental evidence, it is unclear whether the general conclusion that we draw from our theoretical analysis is reliable. Second, except for the work by Ibuka et al. (2014), no experimental study has

directly investigated how the benefit and cost of vaccination influence free-riding behavior in individuals' vaccination decisions. To the best of our knowledge, we are the first to experimentally explore how the nonlinear nature of the positive prevention externality affects people's free-riding behavior in the context of a vaccination choice problem.

5.1 Treatments and Hypotheses

Our experimental implementation considers the following reduced form of the vaccination game. There are eight individuals in a society (i.e., $n = 8$), and they are ex ante identical. Each individual is initially endowed with $u(R)/\mu = 80$. It is assumed that the duration of infection takes the functional form $d_{R_0} = 0.8R_0$ and the utility cost of vaccination $C = 5$. That is, with herd immunity, all free riders receive 80, while people who receive the vaccine receive 75. In the absence of herd immunity, the probability of infection π_P is endogenously determined by P, according to Proposition 1. For example, when the utility cost of infection $L = 25$, the basic reproduction ratio is $R_0 = 4$, and the vaccination coverage is $P = 1/2$, so $\pi_P = 1/2$; i.e., the payoff for a nonvaccinator is equally likely to be 80 and 0 ($= 80 - 0.8 \times 4 \times 25$).

The major treatment variables correspond to the relative benefit of vaccination (r: relative benefit) and the basic reproduction ratio of the epidemic (R_0: reproduction ratio). We choose parameter specifications with $r = 1$ or 5 and $R_0 = 2$ or 4 to create a qualitative difference in the vaccination coverage predicted by the symmetric mixed-strategy Nash equilibrium across treatments. Table 1 presents the four treatments, each of which involves a unique combination of the relative benefit and basic reproduction ratio as well as the equilibrium predictions.[30]

	Parameter Specification		**Theoretical Predictions**			
Treatment	Relative Benefit r	Reproduction Ratio R_0	Herd Immunity nP_{crit}	Mixed Strategy σ^*	Expected Number of Vaccinated People $n\mathbb{E}[P]$	Likelihood of Herd Immunity $\Pr(P \geqslant P_{crit})$
1	1	2	4	0	0	0%
2	5	2	4	0.52	4.16	67.95%
3	1	4	6	0.68	5.44	50.13%
4	5	4	6	0.89	7.12	95.13%

Table 1: Experimental Treatments and Theoretical Predictions

[30] We also have one more treatment (Treatment 3L) that incorporates the hypothetical linearized environment discussed in the previous section into the set of parameters chosen for Treatment 3. We present the experimental design and results from this additional treatment in Section 7.

A unique feature of Treatment 1 is that the reproduction ratio R_0 is large relative to the relative benefit r, indicating that the expected gain from vaccination is dominated by the expected cost. As a result, no vaccination is predicted. With the three other treatments, however, the individual probability of vaccination is positive in the mixed-strategy equilibrium, and the resulting vaccination coverage for the society reaches the critical level needed for herd immunity in more than 50% of cases. Thus, we have our first hypothesis as follows.

Hypothesis 1. *(a) In Treatment 1, the individual probability of vaccination is not significantly different from 0, and as a result, the likelihood of herd immunity is not significantly different from 0.*

(b) In Treatments 2, 3 and 4, the individual probability of vaccination is significantly higher than 0.

(c) In Treatments 2, 3 and 4, the likelihood of herd immunity is significantly higher than 0.

We next consider whether the endogenous epidemic punishment is effective. Given the reproduction ratio R_0, the higher relative benefit r in Treatment 2 (Treatment 4) than in Treatment 1 (Treatment 3) makes it more likely that an individual vaccinates, and thus, the society achieves herd immunity in Treatment 2 (Treatment 4) than in Treatment 1 (Treatment 3).

Hypothesis 2. *Given a fixed reproduction ratio R_0,*

(a) the individual likelihood of vaccination is higher in Treatment 2 than in Treatment 1 and higher in Treatment 4 than in Treatment 3.

(b) the likelihood of herd immunity is higher in Treatment 2 than in Treatment 1 and higher in Treatment 4 than in Treatment 3.

We investigate the effect of the reproduction ratio on individuals' vaccination choices. Given the relative benefit r, the higher reproduction ratio R_0 in Treatment 3 (Treatment 4) than in Treatment 1 (Treatment 2) makes it more likely for an individual to take vaccination in Treatment 3 (Treatment 4) than in Treatment 1 (Treatment 2).

Hypothesis 3. *Given a fixed relative benefit r,*

(a) the individual likelihood of vaccination is higher in Treatment 3 than in Treatment 1 and higher in Treatment 4 than in Treatment 2.

(b) the likelihood of herd immunity is higher in Treatment 3 than in Treatment 1 and higher in Treatment 4 than in Treatment 2.

Notably, the predictions from the asymmetric pure-strategy Nash equilibria are invariant to the changes in the treatment variables. With every treatment, there exists a collection of asymmetric equilibria in which exactly $n \cdot P_{crit}$ players are vaccinated, and the remainder are not. As a result, the likelihood of herd immunity is expected to be 100% in all four treatments. Our experimental data are informative and enable us to reject this prediction from the asymmetric equilibria.

5.2 Procedures

The experiments were conducted in English at the Hong Kong University of Science and Technology Experimental Lab using z-Tree (Fischbacher, 2007). A *between-subjects* design and *random-matching* protocol were used. Four sessions were conducted for each of the four treatments, and each session included twenty-four subjects. Using sessions as independent observation units, we have four observations for each treatment. A total of 384 subjects with no prior experience with our experiment were recruited from the undergraduate and graduate populations of the university and participated in 16 sessions.[31]

Upon arrival at the laboratory, subjects were instructed to sit at separate computer terminals. Each participant was given a copy of the experiment instructions, which were read aloud and supplemented with slide illustrations. In each session, subjects first participated in one practice round and then in 20 official rounds.

We illustrate the instructions for Treatment 1. The full instructions can be found in Appendix D. In each round, a subject was randomly matched with seven other participants to form a group of eight. In each group, the eight members were asked to make decisions that would affect their earnings in the round. The participants were randomly rematched after each round to form new groups.

We asked the subjects to imagine that these eight individuals in a group live in a village. Initially, every individual in the village begins with the same green status. There is a red circle that carries the source of redness, from which subjects want to protect themselves. Each individual independently and simultaneously decides whether to buy the shield. The price of the shield is fixed at 5 experimental currency units (ECU). With the shield, a subject is immune to redness and stays green; without the shield, a subject will either turn red or remain green, depending on how many other individuals in the village have the shield. Table 2 presents the probability of turning red.

[31] We also conducted 4 sessions for an additional treatment that are discussed in Section 7, each of which had 24 subjects.

If you choose "No Shield"	
# of others having the shield	Prob. of turning red
0	50.0%
1	42.9%
2	33.3%
3	20.0%
4	0.00%
5	0.00%
6	0.00%
7	0.00%

Table 2: Probability that you turn red

The earnings in each round are determined by the ex post status of an individual and by whether he/she buys the shield. If the individual does not buy the shield but his/her status remains green, then he/she earns 80 ECU. If his/her status turns red, resulting in the loss of 8 ECU, his/her earnings are $80 - 8 = 72$ ECU. If an individual buys the shield, he/she must pay 5 ECU, and his/her status stays green so that he/she earns $80 - 5 = 75$ ECU.

We randomly selected one round to determine the subjects' payments. A subject was paid the amount of ECU that he/she earned in the selected round at an exchange rate of 10 ECU = 1 HKD. A session lasted for approximately forty-five minutes, and the subjects earned, on average, HK$109 (≈US$14), including a HK$30 show-up fee.[32]

6 Experimental Results

Table 3 presents a summary of our experimental results aggregated over all of the sessions for each treatment. According to Column (1) of Table 3, the individual frequencies of choosing vaccination were significantly greater than zero with all four treatments ($p < 0.066$, signed rank test). Column (2) of the table reveals that, consistent with the theoretical predictions, the mean vaccination coverage was lower than herd immunity in Treatments 1 and 3 and higher than herd immunity in Treatment 2. Moreover, the mean vaccination coverage was surprisingly close to herd immunity in Treatment 4.

Figure 6 reports the mean vaccination coverage aggregated over all of the rounds for each session of each treatment and compares it with the threshold for herd immunity (represented by solid lines) and the theoretical prediction (represented by dashed lines) from the mixed-strategy Nash equilibrium (Nash prediction hereafter). The mean vaccination coverages in all four ses-

[32]Under Hong Kong's currency board system, the HK dollar is pegged to the US dollar at a rate of 1 USD = 7.8 HKD.

Treatment	(1) Vaccination Frequency	(2) Number of Vaccinated People	(3) Likelihood of Herd Immunity
1	0.243	1.944	8.75%
2	0.551	4.408	78.75%
3	0.604	4.832	32.50%
4	0.742	5.936	66.67%

Table 3: Summary of Experimental Results

sions in Treatment 1, as shown by Figure 6(a), were significantly less than the herd immunity level ($p = 0.066$, signed rank test) but significantly greater than the boundary Nash prediction ($p = 0.066$, signed rank test). However, this regularity of a higher actual contribution level than the Nash prediction has been well documented in the standard literature on public good games (e.g., Andreoni, 1995). Similarly, Figure 7 shows that the likelihood of herd immunity aggregated over all of the rounds for each session of each treatment is significantly greater than 0 in all four treatments ($p < 0.07$ for all four cases, signed rank tests). Thus, we confirm Hypotheses 1(b) and 1(c) but reject 1(a). However, inconsistent with the predictions from the pure-strategy equilibria, the likelihood of herd immunity is significantly lower than 100% ($p < 0.058$, all four cases, signed rank test). Thus, we reject the efficient provision of herd immunity predicted by the pure-strategy equilibria.

Result 1. *In all of the treatments, the individual frequency of choosing vaccination and the likelihood of herd immunity are significantly greater than 0. Moreover, in all of the treatments, the likelihood of herd immunity is significantly less than 100%.*

The endogenous epidemic punishment is an effective mechanism, inviting voluntary vaccination to obtain the public good of herd immunity. Figure 7 shows that a higher relative benefit led to a significant increase in the likelihood of herd immunity. The Mann-Whitney test confirms that the likelihood of herd immunity was significantly higher in Treatment 2 than in Treatment 1 ($p = 0.018$) and higher in Treatment 4 than in Treatment 3 ($p = 0.02$). We also observed that a higher reproduction ratio led to a significant increase in the likelihood of herd immunity. Again, the nonparametric Mann-Whitney test confirmed that the likelihood of herd immunity was significantly higher in Treatment 3 than in Treatment 1 ($p = 0.02$) and higher in Treatment 4 than

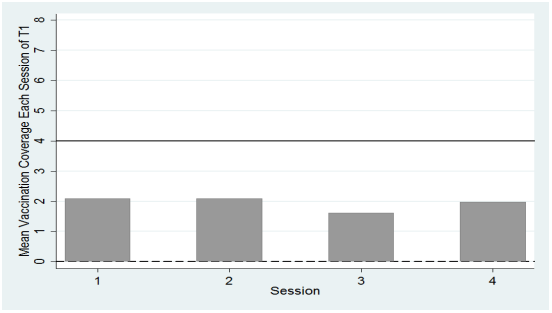

(a) Treatment 1

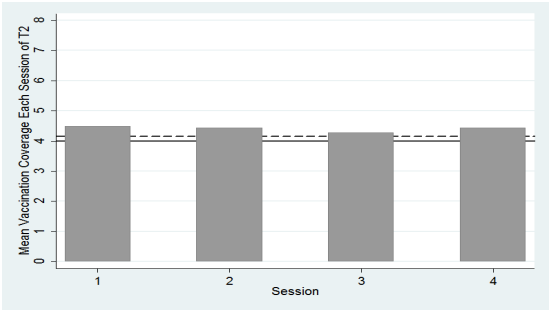

(b) Treatment 2

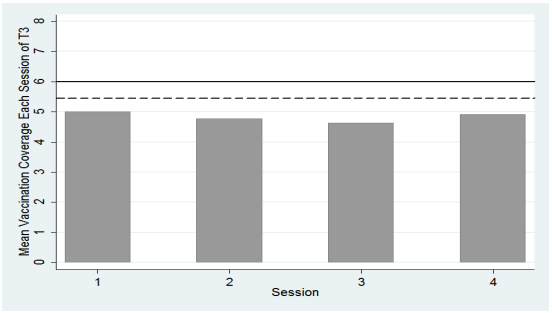

(c) Treatment 3

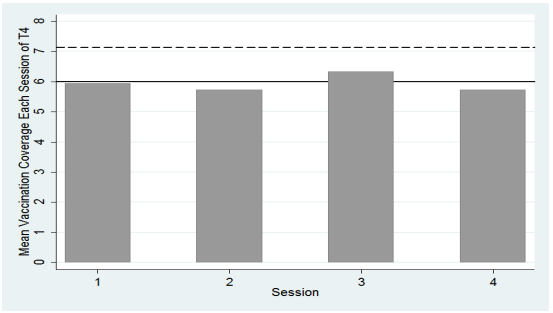

(d) Treatment 4

Note: The solid line represents herd immunity, and the dashed line represents the Nash prediction.

Figure 6: Mean Vaccination Coverage

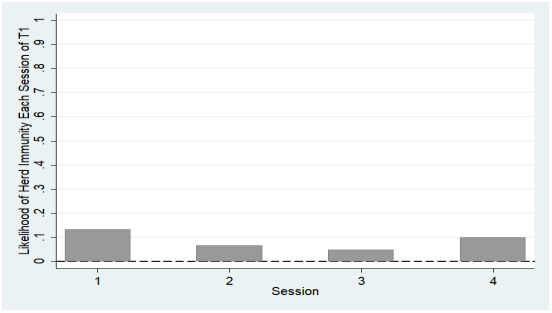

(a) Treatment 1

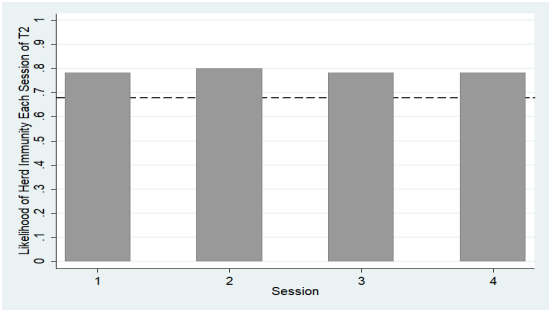

(b) Treatment 2

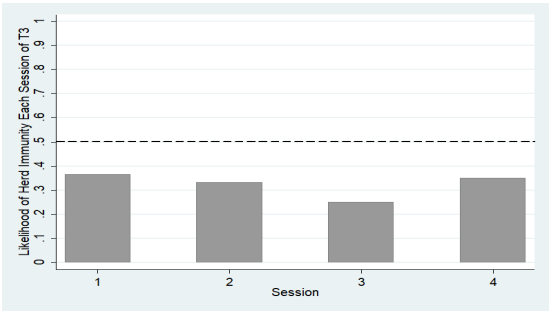

(c) Treatment 3

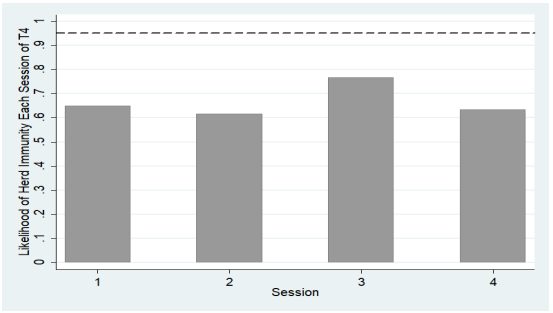

(d) Treatment 4

Note: The dashed line represents Nash prediction.

Figure 7: Likelihood of Herd Immunity

in Treatment 2 ($p = 0.018$), thus confirming our Hypotheses 2 and 3.[33] We have the following results.

Result 2. *Given a fixed reproduction ratio R_0, the likelihood of herd immunity was higher in Treatment 2 than in Treatment 1 and higher in Treatment 4 than in Treatment 3. Given a fixed relative benefit r, the likelihood of herd immunity was higher in Treatment 3 than in Treatment 1 and higher in Treatment 4 than in Treatment 2.*

7 Role of Concavity: Experimental Evidence

In this section, we provide experimental evidence for the role of concavity in enhancing the individual vaccination uptake likelihood and thus the likelihood of herd immunity. To investigate this issue, we considered a new, counterfactual treatment, Treatment 3L, a variant of Treatment 3 in which the probability of infection is linearized. Recall that, in Treatment 3, the long-run probability of infection was nonlinear, as $\pi_P = 1 - \frac{1}{4(1-P)}$. In Treatment 3L, we linearized it to be $\pi_P^L = \frac{3}{4} - P$.[34] There is no other difference between Treatment 3 and Treatment 3L.[35]

The theoretical prediction confirms our intuition and reveals that the linearized probability of infection decreases the equilibrium vaccination uptake likelihood σ^* and the likelihood of herd immunity $Pr(P \geqslant P_{crit})$. Table 4 presents the parameter choices for Treatment 3L and the theoretical predictions, compared to those of Treatment 3.

	Parameter Specification			**Theoretical Predictions**			
Treatment	Relative Benefit r	Reproduction Ratio R_0	Long-run Infection Prob. π_P	Herd Immunity nP_{crit}	Mixed Strategy σ^*	Expected Number of Vaccinated People $n\mathbb{E}[P]$	Likelihood of Herd Immunity $\Pr(P \geqslant P_{crit})$
3	1	4	$1 - \frac{1}{4(1-P)}$	6	0.68	5.44	50.13%
3L	1	4	$\frac{3}{4} - P$	6	0.50	4.00	14.45%

Table 4: Treatment 3L and Its Theoretical Predictions

[33] Figure 10(a) reported in Appendix D plots the group-level time-trend data for the likelihood of herd immunity. The high volatility observed in Figure 10(a) stems from herd immunity generating a binary outcome, and the social cost was tightly associated with the realization of this outcome. Consistent with the Nash prediction, groups rarely contributed sufficiently to achieve herd immunity in Treatment 1, in sharp contrast to the high frequencies of herd immunity achieved in Treatments 2 and 4. The likelihood of herd immunity was substantially higher than 0 in Treatment 3 (32.5%), also consistent with the theoretical prediction.

[34] Note that π_P^L is chosen such that $\pi_P = \pi_P^L$ at $P = 0$ and at $P = P_{crit}$.

[35] Our selection of Treatment 3 for linearizing the environment is guided by the magnitude of change in the individual vaccination uptake likelihood from the linearized π_P being most salient in Treatment 3.

We conducted four sessions of Treatment 3L, and each session had 24 subjects. The same experimental procedure was used. A session lasted for approximately 45 minutes, and subjects earned, on average, HK$105.6 (≈US$13.5), including a HK$30 show-up fee.

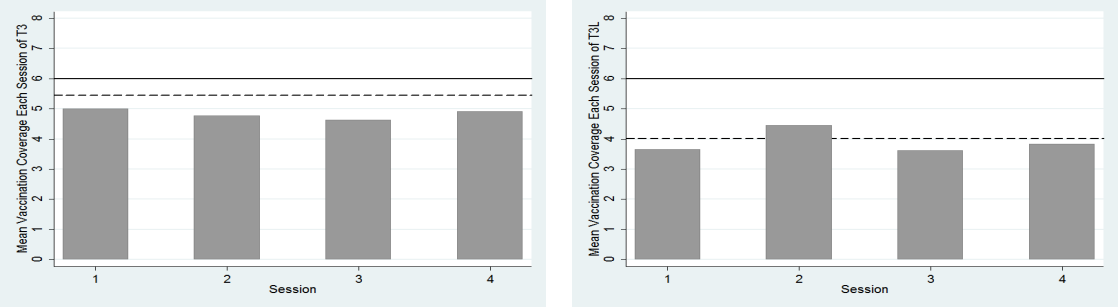

(a) Treatment 3 (b) Treatment 3L

Note: The solid line represents herd immunity, and the dashed line represents the Nash prediction.

Figure 8: Mean Vaccination Coverage

Figure 8 reports the mean vaccination coverage from Treatment 3L. Confirming the theoretical prediction that the concavity of the long-run probability of infection facilitates individual vaccination uptake likelihood, the Mann-Whitney test revealed that the mean vaccination coverage was significantly higher in Treatment 3 than in Treatment 3L (p= 0.02).

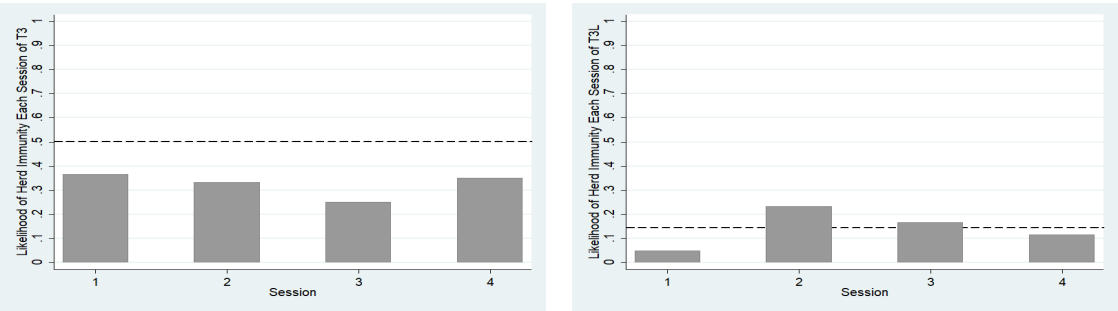

(a) Treatment 3 (b) Treatment 3L

Note: The dashed line represents the Nash prediction.

Figure 9: Likelihood of Herd Immunity

Figure 9 shows that Treatment 3L is dominated by Treatment 3 with respect to the likelihood of herd immunity. Confirming the theoretical prediction that the concavity of the long-run probability of infection increases the likelihood of herd immunity, the Mann-Whitney test revealed that the likelihood of herd immunity was significantly higher in Treatment 3 than in Treatment 3L (p= 0.02). Similarly, Figure 11(a), presented in Appendix D, shows the same dominance relationship between Treatments 3 and 3L. These results are summarized in the following proposition.

Result 3. *The individual frequency of choosing vaccination was significantly higher in Treatment 3 than*

in Treatment 3L. Similarly, the mean likelihood of herd immunity is significantly higher in Treatment 3 than in Treatment 3L.

8 Concluding Remarks

This paper is the first to incorporate a game-theoretic, nonlinear public good game framework into dynamic epidemic modeling. We obtain a novel and striking result in contrast to the previous literature: for each type of disease, there is a range of relative benefit of vaccination such that epidemics can be eliminated through voluntary vaccination. This observation implies that policies regarding epidemic control can be better informed by considering individuals' strategic vaccination choices.

Note that, although our main result is partly consistent with the welfare theorem presented by Francis (1997), it does not hinge on a peculiar set of assumptions such as no spontaneous recovery and infinitely lived, homogenous individuals. The main driver of our result is the concavity of the long-run probability of infection, derived from the dynamic process of epidemics, which dramatically changes the social welfare analysis and equilibrium prediction. Moreover, as in Chen and Toxvaerd (2014), we consider an inflexible vaccination timing problem and show that herd immunity can still be achieved by voluntary vaccination. It is also straightforward to see that our model can be easily extended to the imperfect vaccination case and that the same results can be established.

We conclude our paper by discussing a few natural extensions. First, the marriage between the game-theoretic framework and SIR modeling can be applied to other economic problems. For example, the SIR model has been shown to be useful in explaining the diffusion of technological innovation (Stoneman, 1983) and the spread of information through a population of asset-market traders (Shiller et al., 1984). Investigating these issues by considering players' strategic interplay, in addition to the SIR model, might be able to provide some new insights.

Second, it would be interesting to explicitly consider the optimal vaccination timing problem. Suppose that people receive regular feedback about the proportion of vaccinated individuals in society such that, from time to time, people update their beliefs before making their vaccination choices. This assumption does not appear unrealistic, as was the case with Ebola in underdeveloped countries. People do receive feedback in some instances, especially regarding the number of infected people. In this case, the strategic optimal timing of vaccination is important. It is thus of intellectual interest to consider a supplementary game in which every player has one opportunity

to vaccinate and in which he/she chooses when to obtain the vaccine as the epidemic spreads. We leave these extensions for future research.

Appendix A. Values of R_0 of well-known infectious diseases

Disease	Transmission	R_0
Measles	Airborne	12-18
Pertussis	Airborne droplets	12-17
Diphtheria	Saliva	6-7
Smallpox	Airborne droplets	5-7
Polio	Fecal-oral route	5-7
Rubella	Airborne droplets	5-7
Mumps	Airborne droplets	4-7
HIV/AIDS	Sexual contact	2-5
SARS	Airborne droplets	2-5
Ebola	Bodily fluids	1.5-2.5

Anderson and May (1992)

Appendix B. Proofs

Proof of Proposition 1.

We first show that the steady state automatically satisfies the stability condition of the nonlinear differential system, resulting in the convergence of π_P. We transform the nonlinear differential system into a locally linear system around an open ball of the fixed points (7) and (6).

Define

$$f(\hat{S},\hat{I}) := \mu(1-P) - \beta\hat{S}\hat{I} - \mu\hat{S}$$
$$g(\hat{S},\hat{I}) := \beta\hat{S}\hat{I} - \gamma\hat{I} - \mu\hat{I}.$$

The Jacobian matrix evaluated at the steady state is thus given by

$$J^* = \begin{bmatrix} f_{\hat{S}} & f_{\hat{I}} \\ g_{\hat{S}} & g_{\hat{I}} \end{bmatrix}_{(S^*,I^*)} = \begin{bmatrix} -\beta I^* - \mu & -\beta S^* \\ \beta I^* & \beta S^* - \gamma - \mu \end{bmatrix}.$$

If follows that

$$\begin{aligned} \det|J^* - rI| &= r^2 - traceJ^* r + |J^*| \\ |J^*| &= \beta(\gamma+\mu)I^* - \mu\beta S^* + \mu(\gamma+\mu) \\ traceJ^* &= -(\gamma + 2\mu + \beta I^* - \beta S^*). \end{aligned} \tag{14}$$

The characteristic roots of the equation

$$det|J^* - rI| = 0,$$

thus have the following relations

$$\begin{aligned} r_1 + r_2 &= traceJ^* \\ r_1 r_2 &= |J^*|. \end{aligned} \tag{15}$$

We proceed to proof by cases:

- Herd Immunity Case: $P \geqslant P_{crit}$.
  Substitute $R_0 = \frac{\beta}{\gamma+\mu}$ and rearrange the inequality, we have

  $$\gamma + \mu - \beta(1 - P) > 0. \tag{16}$$

  Note that the steady state 6 gives us $S^* = 1 - P$, $I^* = 0$, which implies

  $$|J^*| > 0 \text{ and } traceJ^* < 0. \tag{17}$$

  It follows that

  $$r_1 < 0, \quad r_2 < 0.$$

- Under-provision Case: $P < P_{crit}$.
  Substitute $R_0 = \frac{\beta}{\gamma+\mu}$ and rearrange the inequality, we have

  $$\beta(1 - P) - (\gamma + \mu) > 0. \tag{18}$$

  Note that the steady state 7 gives us $S^* = \frac{1}{R_0}$, $I^* = \frac{\mu}{\gamma+\mu}(1 - P) - \frac{\mu}{\beta}$, which implies

  $$|J^*| > 0 \text{ and } traceJ^* < 0. \tag{19}$$

  It follows that

  $$r_1 < 0, \quad r_2 < 0.$$

From the stability theorem of system of differential equations, a fixed point of a first-order nonhomogeneous linear system with two variables is locally asymptotically stable if and only if the eigenvalues have negative real parts (i.e., $r_1 < 0, r_2 < 0$). Therefore, the steady state of the nonlinear differential system is locally asymptotically stable for any $P \in [0, 1]$. When r_1 and r_2 are real, the steady state is a stable node. When r_1 and r_2 are complex, the steady state is a stable focus, establishing the dynamic stability of the epidemic.

At the steady state, we have

$$\begin{aligned}\pi_P &= \lim_{t\to\infty} \frac{\beta\hat{S}\hat{I}}{\beta\hat{S}\hat{I}+\mu\hat{S}} = \frac{\beta S^* I^*}{\beta S^* I^* + \mu S^*} = 1 - \frac{1}{R_0(1-P)} \text{ when } P < P_{crit} \\ \pi_P &= \lim_{t\to\infty} \frac{\beta\hat{S}\hat{I}}{\beta\hat{S}\hat{I}+\mu\hat{S}} = \frac{\beta S^* I^*}{\beta S^* I^* + \mu S^*} = 0 \text{ when } P \geqslant P_{crit}.\end{aligned} \tag{20}$$

The uniqueness of π_P can be easily shown by contradiction.

When $P < P_{crit}$, calculating a partial derivative of π_P w.r.t. P gives us

$$\frac{\partial \pi_P}{\partial P} = -\frac{1}{R_0(1-P)^2} < 0. \tag{21}$$

It follows that the maximum of π_P is obtained when P takes the value of its lower bound 0. This step completes the proof. ■

Characterization of All Nash Equilibria

The following proposition characterizes the set of pure-strategy equilibrium outcomes.

Proposition 7. *For $k = 0, 1, \ldots, \nu_{crit} - 1$, $\nu = k+1$ is the pure strategy equilibrium outcome for $\frac{C}{d_{R_0}L} \in (\pi_{k+1}, \pi_k]$.*

Proof. For any $i \in V(\boldsymbol{b})$, he/she has no incentive to deviate because $C \leqslant d_{R_0}L\pi_k$; for any $i \notin V(\boldsymbol{b})$, he/she has no incentive to deviate because $C > d_{R_0}L\pi_{k+1}$. $\nu < k+1$ cannot arise in pure strategy equilibrium because some $i \notin V(\boldsymbol{b})$ can always be better off by accepting vaccination. $\nu > k+1$ cannot arise in pure strategy equilibrium because some $i \in V(\boldsymbol{b})$ can always be better off by not accepting vaccination. Finally, note that $\nu > \nu_{crit}$ cannot arise in equilibrium because the infection probability vanishes. ■

In what follows, we offer a general characterization of all mixed-strategy equilibria and derive Proposition 3 as a corollary.

Proof of Proposition 3.

Let $\mathcal{M}$ be the set of players using mixed strategies, and $|\mathcal{M}| = m$. A mixed-strategy Nash equilibrium requires that every player in $\mathcal{M}$ is indifferent between vaccination and nonvaccination, i.e.,

$$EU_i(vc, \sigma^*_{-i}) = EU_i(nv, \sigma^*_{-i}) \quad \text{for any } i \in \mathcal{M}. \tag{22}$$

It follows that

$$\frac{1}{\mu}u(R) - C = \frac{1}{\mu}u(S) - d_{R_0}L\mathbb{E}[\pi_{P(\boldsymbol{\sigma})}].$$

where $\mathbb{E}[\pi_{P(\boldsymbol{\sigma})}]$ denotes the expected infection probability given $\nu = k$. Note that the additional vaccination arising from mixed-strategy follows the Poisson binomial distribution with success probabilities σ_{-i}, we

obtain

$$\frac{C}{d_{R_0}L} = \sum_{k=0}^{m-1} \pi_{\nu+k} \sum_{V \in \mathcal{P}(\mathcal{M}_{-i};k)} \prod_{j \in V} \sigma_j \prod_{l \in \mathcal{M}_{-i}/V} (1-\sigma_l) \tag{23}$$

for any $i \in \mathcal{M}$. Consider this system of equations (characterizing the indifference conditions for the players in set $\mathcal{M}$) where $\nu \leqslant \min\{\nu_{crit} - 1, n - m\}$; we claim that, for any mixed strategy equilibria with $m > 1$, the mixed-strategy profile σ is unique as shown by the following two lemmas.

Lemma 1. *There exists a solution to (23) for* $\frac{C}{d_{R_0}L} \in (\pi_{\nu+m-1}, \pi_\nu)$ *s.t.* $\sigma_i = \sigma^*$ *for any* $i \in \mathcal{M}$.

Proof. The system (23) reduces to

$$\frac{C}{d_{R_0}L} = \sum_{k=0}^{m-1} \pi_{\nu+k} \binom{m-1}{k} \sigma^{*k} (1-\sigma^*)^{m-1-k}. \tag{24}$$

By the intermediate value theorem, there exists $\sigma^* \in (0,1)$ such that the above equation holds. ■

Lemma 2. *(23) has at most one solution.*

Proof. Define vector-valued function $H\colon [0,1]^n \to \mathbb{R}^n$ where every component function

$$H_i := d_{R_0} r \sigma_i \sum_{k=0}^{m-1} \pi_{\nu+k} \sum_{V \in \mathcal{P}(\mathcal{M}_{-i};k)} \prod_{j \in V} \sigma_j \prod_{l \in \mathcal{M}_{-i}/V} (1-\sigma_l).$$

It is easy to check that H is continuously differentiable on $(0,1)^n$. The system of equations (23) is equivalent to $\sigma_i = H_i(\boldsymbol{\sigma})$ for all $i \in \mathcal{N}$. Suppose there exist two solutions $\boldsymbol{\sigma}^*$ and $\boldsymbol{\sigma}'$ such that $\|\boldsymbol{\sigma}^* - \boldsymbol{\sigma}'\| > 0$. By mean value inequality (Rudin, 1976), we have

$$\|\boldsymbol{\sigma}^* - \boldsymbol{\sigma}'\| = \|H(\boldsymbol{\sigma}^*) - H(\boldsymbol{\sigma}')\| \leqslant \|DH(\xi)\| \cdot \|\boldsymbol{\sigma}^* - \boldsymbol{\sigma}'\| \tag{25}$$

where $\xi \in (0,1)^n$ and $DH(\xi)$ is the Jacobian matrix evaluated at ξ. Since the row vectors of $DH(\xi)$ are linearly dependent, $DH(\xi)$ is not invertible and thus $\|DH(\xi)\| = 0$. It follows that $\|\boldsymbol{\sigma}^* - \boldsymbol{\sigma}'\| \leqslant 0$. This outcome requires a contradiction. ■

Combining the two lemmas, we reach the conclusion that, in any mixed strategy equilibrium with $m > 1$, mixing probabilities must be unique and identical across players. The next proposition characterizes all the mixed-strategy equilibria for $m > 1$,

Proposition 8. *For* $m > 1$ *and* $\nu \leqslant \min\{\nu_{crit} - 1, n - m\}$, $\langle \nu, m \rangle$ *arises as a mixed strategy equilibrium outcome*

for $\frac{C}{d_{R_0}L} \in (\pi_{\nu+m-1}, \pi_\nu)$ *with* $\boldsymbol{\sigma} = \boldsymbol{\sigma}^*$ *and is uniquely determined by*

$$\frac{C}{d_{R_0}L} = \sum_{k=0}^{m-1} \pi_{\nu+k} \binom{m-1}{k} \sigma^{*k}(1-\sigma^*)^{m-1-k}. \tag{26}$$

Proof. (26) implies the best response of any $i \in \mathcal{M}$. Any $i \in V(\boldsymbol{b})$ has no incentive to deviate since his/her incentive constraint $C < d_{R_0}L\mathbb{E}[\pi_{k-1}]$ can be simplified using (26) as

$$\frac{C}{d_{R_0}L} < \sigma^{m-1}\pi_{\nu+m-2} + (1-\sigma)^{m-1}\pi_{\nu-1},$$

which holds under $\frac{C}{d_{R_0}L} \in (\pi_{\nu+m-1}, \pi_\nu)$. Any $i \notin V(\boldsymbol{b})$ has no incentive to deviate since his/her incentive constraint $C > d_{R_0}L\mathbb{E}[\pi_{P(\boldsymbol{\sigma})}]$ can be simplified using (26) as

$$\frac{C}{d_{R_0}L} > \sigma^{m-1}\pi_{\nu+m} + (1-\sigma)^{m-1}\pi_{\nu+1},$$

which again holds under $\frac{C}{d_{R_0}L} \in (\pi_{\nu+m-1}, \pi_\nu)$. ■

It follows from this more general proposition that there exists a unique totally mixed strategy equilibrium, where $m = n$,

Corollary 2. *There exists a unique totally mixed strategy equilibrium, where* $\sigma_i^* = \sigma^*$*, and it is implicitly defined by*

$$\frac{1}{d_{R_0}r} = \sum_{k=0}^{\nu_{crit}} \left(1 - \frac{1}{R_0} - \frac{k}{n}\right) \binom{n}{k} \sigma^{*k}(1-\sigma^*)^{n-1-k}. \tag{27}$$

This step completes the proof. ■

Proof of Proposition 4.
Let ν and ν' be the equilibrium number of vaccinated people given the relative benefit r and r', respectively. For future use, denote $F_\nu(x)$ as the CDF of a random variable ν. We start by working directly with σ, as shown in the following lemma.

Lemma 3. ν *FOSD* ν' *if and only if* $\sigma > \sigma'$.

Proof. ν FOSD ν' if and only if $F_\nu(x) \leqslant F_{\nu'}(x)$ for any $x \in \{1, \ldots, n\}$. Since both ν and ν' follows

binomial distribution with n trials, it remains to show $\frac{dF_\nu(x)}{d\sigma} \leqslant 0$. Now consider for any x, the derivative

$$\begin{aligned}\frac{dF_\nu(x)}{d\sigma} &= \sum_{k=1}^{x} k\binom{n}{k}\sigma^{k-1}(1-\sigma)^{n-k} - \sum_{k=0}^{x}(n-k)\binom{n}{k}\sigma^{k}(1-\sigma)^{n-k-1} \\ &= n\left(\sum_{k=1}^{x}\binom{n-1}{k-1}\sigma^{k-1}(1-\sigma)^{n-k} - \sum_{k=0}^{x}\binom{n-1}{k}\sigma^{k}(1-\sigma)^{n-k-1}\right) \\ &= n\left(F_{\tilde{\nu}}(x-1) - F_{\tilde{\nu}}(x)\right) \leqslant 0\end{aligned}$$

where $\tilde{\nu} \sim Bin(n-1,\sigma)$. ■

It remains to be shown that σ^* is monotonically increasing in r. By the implicit function theorem, calculating the partial derivative of the vaccine uptake likelihood $\sigma^*(r,R_0)$ with respect to r gives us

$$\frac{\partial\sigma^*}{\partial r} = 1/\left[d_{R_0}r^2\sum_{k=0}^{v_{crit}}\left((n-1)\sigma^*-k\right)\left(1-\frac{1}{R_0}-\frac{k}{n}\right)\binom{n}{k}\sigma^{*k-1}(1-\sigma^*)^{n-2-k}\right],$$

which is positive for any $r > \frac{R_0}{(R_0-1)d_{R_0}}$.

Therefore, ν FOSD ν' and equivalently, $P(r,R_0)$ FOSD $P(r',R_0)$ if and only if $r > r'$. ■

Proof of Corollary 1.

By Proposition 4, $\Pr^*(P \geqslant P_{crit}) = 1 - F_\nu(v_{crit})$ is monotonically increasing in r. Furthermore, as $r \to \infty$, $\sigma^* \to 1$, and $F_\nu(v_{crit}) \to 0$. ■

Proof of Proposition 5.

By Proposition 8, the totally mixed-strategy equilibrium in the linearized environment $\sigma_i^L = \sigma^L$ is implicitly defined by

$$\frac{1}{d_{R_0}r} = \sum_{k=0}^{v_{crit}}\left(1-\frac{1}{R_0}-\frac{k}{n}\right)\binom{n-1}{k}{\sigma^L}^k(1-\sigma^L)^{n-1-k}. \tag{28}$$

By Pascal's rule, $\binom{n-1}{k} \leqslant \binom{n}{k}$ for any $k < n$ and thus $\sigma^L \leqslant \sigma^*$. By lemma 3, $P^*(r,R_0)$ FOSD $P^L(r,R_0)$. ■

Proof of Proposition 6.

By the implicit function theorem, calculating the partial derivative of the vaccine uptake likelihood $\sigma^*(r,R_0)$ with respect to R_0 gives us

$$\frac{\partial\sigma^*}{\partial R_0} = \frac{\sum_{k=0}^{v_{crit}} {R_0}^{-2}\binom{n}{k}\sigma^{*k}(1-\sigma^*)^{n-1-k} + \frac{\dot{d_{R_0}}}{d_{R_0}^2 r}}{\sum_{k=0}^{v_{crit}}\left((n-1)\sigma^*-k\right)\left(1-\frac{1}{R_0}-\frac{k}{n}\right)\binom{n}{k}\sigma^{*k-1}(1-\sigma^*)^{n-2-k}}, \tag{29}$$

which is positive for $r > \frac{R_0}{(R_0-1)d_{R_0}}$. By lemma 3, $P^*(r,R_0)$ FOSD $P^*(r,R_0')$ if and only if $R_0 > R_0'$. ■

Proposition 9. *For any* $r > \frac{R_0}{(R_0-1)d_{R_0}}$ *and* $R_0 > 1$, $\mathbb{E}[P^*] > P'$ *as* $n \to \infty$.

Proof. In the totally mixed-strategy equilibrium, letting n approach infinity, we have for any i,

$$\frac{1}{\mu}u_i(R) - C_i = \frac{1}{\mu}u_i(S) - d_{R_0}L_i\mathbb{E}[\pi_P].$$

where the expectation is over a Poisson distribution with mean σ. By Markov's inequality, it follows that, for any $\sigma \in (0,1]$,

$$\mathbb{E}[\pi_P] \geqslant \sigma \Pr[\pi_P \geqslant \sigma] = \sigma \Pr[1 - \frac{1}{R_0(1-P)} \geqslant \sigma] = \sigma F\left(1 - \frac{1}{R_0(1-\sigma)}\right).$$

Note that, for any $\sigma \in (0,1]$,

$$\sigma F\left(1 - \frac{1}{R_0(1-\sigma)}\right) \geqslant 1 - \frac{1}{R_0(1-\sigma)}.$$

Define the function

$$h(\sigma) := e^{1-\sigma-\frac{1}{R_0(1-\sigma)}}\left(\frac{\sigma}{1-\frac{1}{R_0(1-\sigma)}}\right)^{2-\frac{1}{R_0(1-\sigma)}}.$$

By Chernoff's formula, it is equivalent to show that $h(\sigma) > 1$ on $(0,1]$, which follows from $h(0) > 1$ and $h(\cdot)$ being continuous and increasing in σ. Now, by the definition of the market equilibrium, we have

$$\frac{1}{\mu}u_i(S) - d_{R_0}L_i\pi_{P'} = \frac{1}{\mu}u_i(S) - d_{R_0}L_i\mathbb{E}[\pi_P] < \frac{1}{\mu}u_i(S) - d_{R_0}L_i\pi_{\mathbb{E}[P]}.$$

since $P' < P_{crit}$ and π_P is decreasing in P, we have $\mathbb{E}[P] > P'$ as $n \to \infty$. ■

Appendix C. Figures and Tables

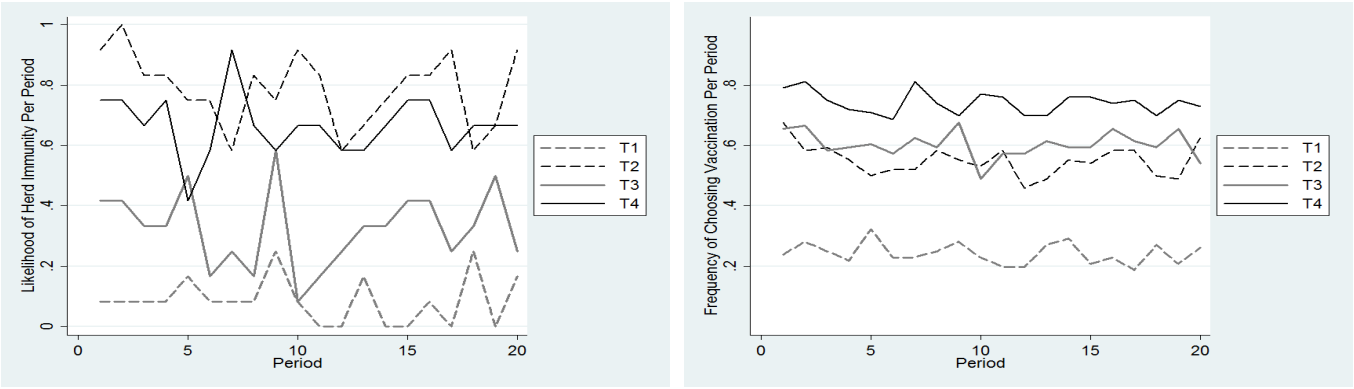

(a) Likelihood of Herd Immunity (b) Vaccination Coverage

Figure 10: Time Trend

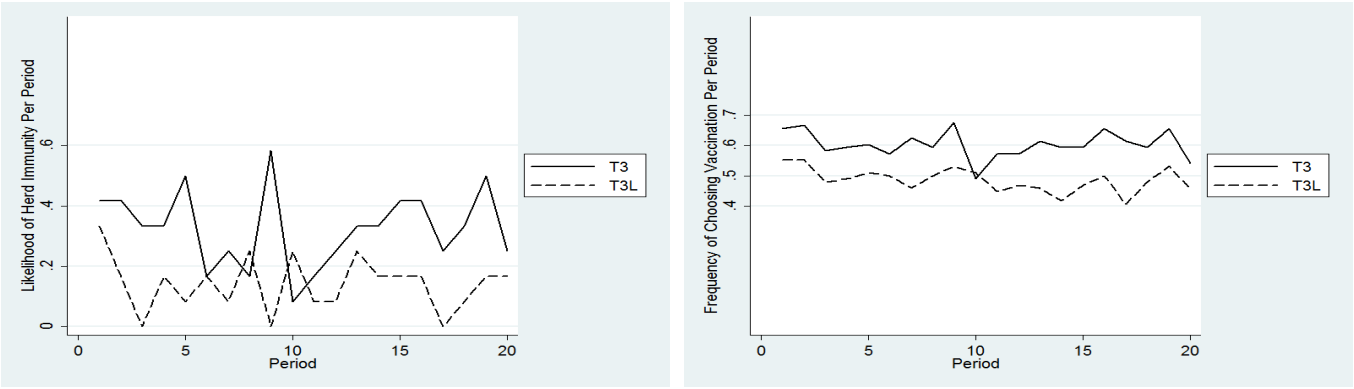

(a) Likelihood of Herd Immunity (b) Vaccination Coverage

Figure 11: Time Trend - Treatment 3L

Appendix D. For Online Publication: Experimental Instructions – Treatment 1

INSTRUCTION

Welcome to the experiment. This experiment studies decision making in groups of eight individuals. In the following two hours or less, you will participate in 20 rounds of decision making. Please read the instructions below carefully; the cash payment that you will receive at the end of the experiment depends on how well you make your decisions according to these instructions. If you have a question at any point, please raise your hand and wait for one of us to come over. We will then privately answer your question. We ask that you turn off your mobile phone and any other electronic devices. Communication of any kind with the other participants is not allowed.

Your Group

In each and every round, you will be *randomly* matched with seven other participants to form a group of eight. You will not be told the identity of the participants with whom you are matched, nor will these participants be told your identity—even after the end of the experiment. In each group, eight members are required to make decisions that will affect their earnings in the round. Participants will be randomly rematched after each round to form new groups.

Your Decision

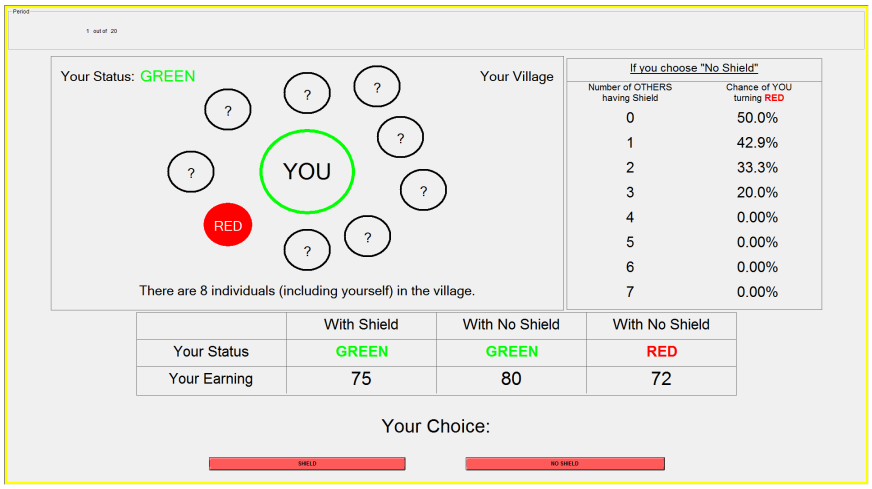

Figure 12: Screen Shot – Your Decision

There are eight individuals (including yourself) in your group. Suppose that these eight individuals are living in a village. Initially, every individual in the village begins with the same status of green. Figure 12 illustrates this situation. The green circle at the center represents yourself, and the other seven black circles represent other individuals in the village. The red circle carries the source of redness from which you want to protect yourself.

The only way for you to protect yourself from the redness is to purchase a shield. With the shield, you will be immune to the redness and stay green; Without the shield, you will either turn red or remain green. Without the shield, the chance of your turning ted depends on how many other individuals in the village have the shield. Table 5 below (or the table at the top-right corner of Figure 12) presents the probability of turning red. Note that, when four individuals or more in your village choose to buy the shield, everyone will stay green regardless of whether he/she buys the shield or not.

If you choose "No Shield"	
# of others having the shield	Chance of your turning red
0	50.0%
1	42.9%
2	33.3%
3	20.0%
4	0.00%
5	0.00%
6	0.00%
7	0.00%

Table 5: Probability of your turning red

Each individual independently and simultaneously decides whether to buy the shield or not. You will be prompted to make your decision by clicking one of the two buttons, "Shield" and "No Shield," presented at the bottom of your screen. Once you click one of the buttons, your decision in the round is completed.

Your Earnings

Your earning in each round depends on: 1) whether or not you buy the shield; and 2) your status. Table 6 below summarizes your earnings, in which the earnings are expressed in terms of the experimental currency units (ECU).

Your decision	Shield	No Shield	
Your status	Green	Green	Red
Your earning	75	80	72

Table 6: Your Earnings

Information Feedback

At the end of each round, the computer will provide a summary for the round: your decision (Shield or No Shield), your final status (Green or Red), your earnings, and the number of people having the shield in your group.

Your Cash Payment

The experimenter randomly selects 1 round to calculate your cash payment. Each round has an equal chance of being selected (so it is in your best interest to consider each round seriously.) Your final cash payment will be your earnings (1 ECU = 1 HKD) in the selected round, plus a HK$30 show-up fee.

Practice Rounds

To ensure your understanding of the instructions, we will provide you with a quiz and a practice round. We will go through the quiz after you answer it on your own. You will then participate in 1 practice round. The practice round is part of the instructions that are not relevant to your cash payment; its objective is to familiarize you with the computer interface and the flow of the decisions in each round. Once the practice round is over, the computer will tell you "The official rounds begin now!"

Administration

Your decisions, as well as your monetary payment, will be kept confidential. Remember that you must make your decisions entirely on your own; please do not discuss your decisions with any of the other participants. Upon finishing the experiment, you will receive your cash payment. You will be asked to sign your name to acknowledge your receipt of the payment (which will not be used for tax purposes). You are then free to leave. If you have any questions, please raise your hand now. We will answer your questions individually. If there are no questions, we will proceed to the quiz.

Quiz

1. True or False: I will be matched with the same seven other players in all 20 rounds.
2. True or False: I will turn red for sure if I do not have the shield.
3. Suppose that you decide to buy the shield. Calculate your earnings. _____
4. Suppose that you decide not to buy the shield. It turns out that there are five other individuals in your group who buy the shield. Calculate your earnings. _____
5. Suppose that you decide not to buy the shield. It turns out that there are two other individuals in your group who buy the shield. What is the chance of your turning red? _____

References

Anderson, Roy M and Robert M May, *Infectious diseases of humans: dynamics and control*, Vol. 28, Wiley Online Library, 1992.

Andreoni, James, "Cooperation in public-goods experiments: kindness or confusion?," *The American Economic Review*, 1995, pp. 891–904.

Bagnoli, Mark and Michael McKee, "Voluntary contribution games: Efficient private provision of public goods," *Economic Inquiry*, 1991, *29* (2), 351–366.

Barrett, Scott, "Global disease eradication," *Journal of the European Economic Association*, 2003, *1* (2-3), 591–600.

Bauch, Chris T and David JD Earn, "Vaccination and the theory of games," *Proceedings of the National Academy of Sciences of the United States of America*, 2004, *101* (36), 13391–13394.

Boulier, Bryan L, Tejwant S Datta, and Robert S Goldfarb, "Vaccination externalities," *The BE Journal of Economic Analysis & Policy*, 2007, *7* (1).

Bracha, Anat, Michael Menietti, and Lise Vesterlund, "Seeds to succeed?: Sequential giving to public projects," *Journal of Public Economics*, 2011, *95* (5), 416–427.

Brito, Dagobert L, Eytan Sheshinski, and Michael D Intriligator, "Externalities and compulsary vaccinations," *Journal of Public Economics*, 1991, *45* (1), 69–90.

Cambridge Dictionary, "Cambridge dictionaries online," 2015.

Casari, Marco and Charles R Plott, "Decentralized management of common property resources: experiments with a centuries-old institution," *Journal of Economic Behavior & Organization*, 2003, *51* (2), 217–247.

Cason, Timothy N and Lata Gangadharan, "Promoting cooperation in nonlinear social dilemmas through peer punishment," *Experimental Economics*, 2015, *18* (1), 66–88.

Chen, Frederick and Flavio Toxvaerd, "The economics of vaccination," *Journal of theoretical biology*, 2014, *363*, 105–117.

Diekmann, Andreas, "Volunteer's dilemma," *Journal of Conflict Resolution*, 1985, pp. 605–610.

Dixit, Avinash and Mancur Olson, "Does voluntary participation undermine the Coase Theorem?," *Journal of Public Economics*, 2000, *76* (3), 309–335.

Draief, Moez and Laurent Massoulié, "Epidemics and rumours in complex networks, volume 369 of London Mathematical Society Lecture Notes," *Cambridge University Press, Cambridge*, 2010, *1* (52), 103.

Farrell, Joseph and Garth Saloner, "Standardization, compatibility, and innovation," *The RAND Journal of Economics*, 1985, pp. 70–83.

Fischbacher, Urs, "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental economics*, 2007, *10* (2), 171–178.

Francis, Peter J, "Dynamic epidemiology and the market for vaccinations," *Journal of Public Economics*, 1997, *63* (3), 383–406.

Galeotti, Andrea and Brian W Rogers, "Strategic Immunization and Group Structure," *American Economic Journal: Microeconomics*, 2013, *5* (2), 1–32.

Geoffard, Pierre-Yves and Tomas Philipson, "Disease eradication: private versus public vaccination," *The American Economic Review*, 1997, pp. 222–230.

Gersovitz, Mark and Jeffrey S Hammer, "The Economical Control of Infectious Diseases," *The Economic Journal*, 2004, *114* (492), 1–27.

Goeree, Jacob K and Charles A Holt, "An explanation of anomalous behavior in binary-choice games: Entry, voting, public goods, and the volunteers dilemma," *Unpublished paper, University of Virginia*, 2000.

Goeree, Jacob K., Charles A. Holt, and Angela M. Smith, "An experimental examination of the volunteer's dilemma," *Games and Economic Behavior*, 2017, *102*, 303 – 315.

Heal, Geoffrey and Howard Kunreuther, "The vaccination game," *Risk Management and Decision Processes Center Working Paper*, 2005, (05-10).

Ibuka, Yoko, Meng Li, Jeffrey Vietri, Gretchen B Chapman, and Alison P Galvani, "Free-riding behavior in vaccination decisions: An experimental study," *PloS one*, 2014, *9* (1), e87164.

Isaac, R Mark and James M Walker, "Nash as an organizing principle in the voluntary provision of public goods: Experimental evidence," *Experimental Economics*, 1998, *1* (3), 191–206.

Jackson, Matthew O, *Social and economic networks*, Princeton university press, 2010.

Just, Winfried, Hannah Callender, M Drew LaMar, and R Robeva, "Disease transmission dynamics on networks: Network structure vs. disease dynamics," *Algebraic and Discrete Mathematical Methods for Modern Biology*, 2015, pp. 217–235.

Katz, Michael L and Carl Shapiro, "Network externalities, competition, and compatibility," *The American economic review*, 1985, *75* (3), 424–440.

Keeling, Matt J and Pejman Rohani, *Modeling infectious diseases in humans and animals*, Princeton University Press, 2008.

Kermack, William O and Anderson G McKendrick, "A contribution to the mathematical theory of epidemics," in "Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences," Vol. 115 The Royal Society 1927, pp. 700–721.

Keser, Claudia, "Voluntary contributions to a public good when partial contribution is a dominant strategy," *Economics Letters*, 1996, *50* (3), 359–366.

Kretzschmar, Mirjam and Martina Morris, "Measures of concurrency in networks and the spread of infectious disease," *Mathematical biosciences*, 1996, *133* (2), 165–195.

Kureishi, Wataru, "Partial vaccination programs and the eradication of infectious diseases," *Economics Bulletin*, 2009, *29* (4), 2758–2769.

Laury, Susan K, James M Walker, and Arlington W Williams, "The voluntary provision of a pure public good with diminishing marginal returns," *Public Choice*, 1999, *99* (1), 139–160.

Laxminarayan, Ramanan, Anne J Mills, Joel G Breman, Anthony R Measham, George Alleyne, Mariam Claeson, Prabhat Jha, Philip Musgrove, Jeffrey Chow, Sonbol Shahid-Salles et al., "Advancement of global health: key messages from the Disease Control Priorities Project," *The Lancet*, 2006, *367* (9517), 1193–1208.

Mas-Colell, Andreu, Michael Dennis Whinston, Jerry R Green et al., *Microeconomic theory*, Oxford university press New York, 1995.

Ostrom, Elinor, Roy Gardner, and James Walker, *Rules, games, and common-pool resources*, University of Michigan Press, 1994.

Palfrey, Thomas R and Howard Rosenthal, "Participation and the provision of discrete public goods: a strategic analysis," *Journal of public Economics*, 1984, *24* (2), 171–193.

Pastor-Satorras, Romualdo and Alessandro Vespignani, "Epidemic dynamics and endemic states in complex networks," *Phys. Rev. E*, May 2001, *63*, 066117.

__ **and** __ , "Epidemic spreading in scale-free networks," *Physical review letters*, 2001, *86* (14), 3200.

__ **and** __ , "Immunization of complex networks," *Physical review letters*, Feb 2002, *65*, 036104.

Reuben, Ernesto and Arno Riedl, "Enforcement of contribution norms in public good games with heterogeneous populations," *Games and Economic Behavior*, 2013, *77* (1), 122–137.

Rosenthal, Robert W, "Sequences of games with varying opponents," *Econometrica: Journal of the Econometric Society*, 1979, pp. 1353–1366.

Rudin, Walter, "Principles of Mathematical Analysis," 1976.

Sattenspiel, Lisa and Carl P Simon, "The spread and persistence of infectious diseases in structured populations," *Mathematical Biosciences*, 1988, *90* (1-2), 341–366.

Sefton, Martin and Richard Steinberg, "Reward structures in public good experiments," *Journal of Public Economics*, 1996, *61* (2), 263–287.

Shiller, Robert J, Stanley Fischer, and Benjamin M Friedman, "Stock prices and social dynamics," *Brookings papers on economic activity*, 1984, *1984* (2), 457–510.

Stoneman, Paul, *The economic analysis of technological change*, Oxford [Oxfordshire]; New York: Oxford University Press, 1983.

WHO Ebola Response Team, "West African Ebola epidemic after one year-slowing but not yet under control," *N Engl J Med*, 2015, *372* (6), 584–7.