Economic epidemiology in the wake of Covid-19*

David McAdams†

August 27, 2020

Abstract

Infectious diseases, ideas, new products, and other "infectants" spread in epidemic fashion through social contact. The Covid-19 pandemic, the proliferation of "fake news," and the rise of antibiotic resistance have thrust economic epidemiology into the forefront of public-policy debate and re-invigorated the field. Focusing for concreteness on disease-causing pathogens, this paper provides a taxonomy of economic-epidemic models, emphasizing both the biology / immunology of the disease and the economics of the social context. An *economic epidemic* is one whose diffusion through the agent population is generated by agents' endogenous behavior. I highlight properties of the Nash-equilibrium epidemic trajectory and discuss ways in which public-health authorities can change the game for the better, (i) by imposing restrictions on agent activity to reduce the harm done during a viral outbreak and (ii) by enabling diagnostic-informed interventions to slow or even reverse the rise of antibiotic resistance.

Keywords: economic epidemic, epidemic limbo, Covid-19, lockdown, antibiotic resistance, diagnostics

*When citing this paper, please use the following: McAdams D. 2021. Economic epidemiology in the wake of Covid-19. *Annual Review of Economics* 13: Submitted. https://doi.org/10.1146/annurev-economics-082120-122900

†Fuqua School of Business and Economics Department, Duke University, Durham, North Carolina, USA. Email: david.mcadams@duke.edu. I thank David Argente, Chris Avery, Troy Day, Mike Hoy, Gregor Jarosch, Philipp Kircher, Anton Korinek, Ramanan Laxminarayan, Tomas Philipson, Elena Quercioli, Steve Redding, Tim Reluga, Bob Rowthorn, Yangbo Song, Arjun Srinivasan, and Flavio Toxvaerd for helpful comments and encouragement.

The spring of 2020 will long be remembered for the loss of life and widespread economic disruption due to Covid-19, the disease caused by the novel coronavirus SARS-CoV2. Yet something constructive came out of those awful months: many economists discovered infectious-disease epidemiology. The volume of new work was so great that *Covid Economics*, an online journal of the Centre for Economic Policy Research launched in April 2020, published twelve issues in May alone. Avinash Dixit, a renowned economic theorist, wittily remarked: "If any pandemic spread faster than Covid-19, it is that of research about Covid-19" (Dixit 2020).[1]

In fact, this was the second wave of infectious interest among economists in infectious disease. The first came in the 1990s, motivated by the global HIV/AIDS pandemic and adding an economic dimension to the classic epidemiological models used to chart the course of a viral outbreak. Because HIV spreads primarily through sexual intercourse, people's decisions around sex clearly impact HIV's spread. Geoffard and Philipson (1996, 1997), Kremer (1996), and others therefore argued that the transmission rate of the virus needed to be treated as a time-varying endogenous variable, derived as a Nash-equilibrium outcome of a dynamic game.

The new generation of economists studying SARS-CoV2 fits the same basic mold but, much like a superbug returning with new genetic machinery, today's economic epidemiologists come with new tools and perspectives drawn from other subfields of economics. The intellectual connectedness between economic epidemiology and other subfields was readily apparent in the various online workshops that sprung up in the virus' wake. For instance, the *Covid-19 Search and Matching Workshop* series (hosted by the labor economist Simon Mongey) had "an emphasis on understanding how the economics of search and matching models can be useful for understanding economic and virological aspects of the coronavirus epidemic."

Of course, viruses aren't the only things that spread infectiously, and SARS-CoV2 isn't the only parasite currently burdening our society. False information and hateful beliefs are colonizing our minds, spreading much like viruses but accelerated by social-media platforms and amplified by partisan outlets and foreign adversaries. Interestingly, the 1990s also saw the first substantial wave of interest in this other form of infection. Social-learning models emerged in which *infectious behavior* played a central

[1]After the National Bureau of Economic Research (NBER) released more than a dozen pandemic-related working papers on April 13th, the MIT economist Jonathan Parker quipped, "Do we need to flatten the curve so we don't exceed NBER WP capacity?" (https://twitter.com/ProfJAParker/status/1249739129962876928).

role, most notably Bikhchandani et al. (1992) and Banerjee (1992) on "information cascades / herding" and Banerjee (1993) on "rumors." New work has evolved these foundational early models in directions that draw even closer parallels with infectious-disease epidemiology, for instance, by re-framing social learning in an epidemic context (McAdams and Song 2020a) and highlighting the possibility of mutation and selection during an information outbreak (Jackson et al. 2018).

The general topic of economic infection intersects with enormous literatures in several fields, from parasitology and public health (global pandemics, antibiotic resistance, evolution of virulence) to finance and information systems (information diffusion, media). Rather than attempting to provide a comprehensive review, I have decided to focus on two central thematic questions: *how an economic epidemic unfolds over time* and *whether economic infectants can "survive" in the long run.* Moreover, I restrict attention here to epidemics of biological pathogens—leaving information epidemics as fertile ground for a future review.

I focus here on recent developments, but credit is due to the handful of economists who pushed economic epidemiology forward during the 2000s and 2010s,[2] a time when most economists showed little interest in the field. A steady trickle of notable empirical contributions appeared in leading economics outlets (e.g., Lakdawalla et al. (2006), Adda (2016), Chan et al. (2016), Greenwood et al. (2019)) but, with a few exceptions (e.g., Auld (2003)), the best new theoretical work by economists found its home in biology journals (e.g., Chen (2004, 2006, 2012), Chen and Toxvaerd (2014)) or remained unpublished for years (e.g., Rowthorn and Toxvaerd (2012)). Fortunately, theoretical biology was a welcoming space for economic theorists, as mathematical epidemiologists and evolutionary ecologists had already embraced game-theoretic methods; see e.g., Bauch et al. (2003), Bauch and Earn (2004), Cressman et al. (2004), and Reluga (2010, 2013). What they did in those years, economists and biologists together, laid the groundwork for the blossoming of economic epidemiology that we see today.

The rest of the paper is organized as follows. Section 1 provides a taxonomy of economic-epidemic models, based on the immunology of infection, manner of transmission, agent decision-making, and economic impacts of agent behavior. Section 2 discusses key features of the equilibrium epidemic trajectory, accounting for agents' behavioral response. Section 3 then examines "lockdown policies" that restrict agents' ability to remain socially active. Section 4 concludes by exploring the possibility of

[2]Useful surveys include Philipson (2000), Gersovitz (2011), and Manfredi and D'Onofrio (2013).

eradicating a disease-causing pathogen through treatment (Section 4.1) and of eradicating the antibiotic-resistant strains of a pathogen—thereby restoring the effectiveness of existing antibiotics—through diagnostic-informed interventions (Section 4.2).

1 A taxonomy of economic-epidemic models

This section provides a taxonomy of economic-epidemiological models of a viral epidemic, categorizing these models along four main dimensions: immunology; transmission; agent decision-making; and economic impacts. Along the way, I introduce notation and preliminary analysis used throughout the rest of the paper.

1.1 Immunology

The epidemiological dynamics of infection hinge on how the virus interacts with the host immune system. Is it possible to recover from infection? If so, does recovery confer subsequent immunity from re-infection? Does transmission begin immediately after infection? How about harmful symptoms? Is it possible to spread the virus without showing any symptoms? Is infection deadly? Are some hosts more prone to be infected, experience symptoms, transmit the virus, or die? Because there are so many possibilities, there is no single benchmark model of a viral epidemic. There is rather an array of benchmark models, what I will refer to as the "SI-X models."[3] See Hethcote et al. (2002) for an epidemiological review and Avery et al. (2020) for a useful critical survey of early models of the SARS-CoV-2 epidemic from an economic perspective.

SI model. The simplest variation is the "Susceptible-Infected (SI) model." A pathogen circulates among a unit-mass population of hosts, each of whom is either uninfected (i.e., "susceptible," state S) or infected (state I) at each point in time t. Let $S(t)$ and $I(t)$ be the mass of susceptible and infected agents at time t. Each susceptible host becomes infected upon meeting an infected host, with such meetings occurring at rate

[3]These models build on an intellectual foundation laid over a century ago by Ronald Ross and Hilda Hudson (Ross 1916, Ross and Hudson 1917) and further systematized by Kermack and McKendrick (1927). For more on the history of the theory of epidemics, see Serfling (1952) and the citation tree on Tim Reluga's website (http://personal.psu.edu/tcr2/post20150624.html).

$\beta I(t)$, where $\beta > 0$ is the "transmission rate."[4]

In the SI model, hosts never recover from infection but may[5] be born into the susceptible state and may die due to infection and/or for other reasons. In the simplest case when the host population is fixed, epidemic dynamics are characterized by the differential equation

$$I'(t) = \beta I(t)S(t) \tag{1}$$

and the adding-up condition $S(t) + I(t) = 1$. In this case, everyone in the population will eventually be infected. More generally, suppose that there is an equal flow $z \geq 0$ of births and deaths across the population, and assume for simplicity that each host dies at constant rate z. Equation (1) then becomes $I'(t) = \beta I(t)S(t) - zI(t)$, and the steady-state mass of infection $I^\infty \equiv \lim_{t\to\infty} I(t) = 1 - \frac{z}{\beta}$.

What if, in addition, infected hosts die at some rate $x > 0$? The host population, typically denoted $N(t) = S(t) + I(t)$, is no longer fixed:

$$I'(t) = \beta I(t)S(t) - (z + x)I(t) \tag{2}$$

$$N'(t) = z - z(S(t) + I(t)) - xI(t) \tag{3}$$

For simplicity, I henceforth focus on models with a fixed host population, an assumption that is most appropriate when the epidemic is fast-moving and the disease is not deadly.

SIRS/SIR/SIS model. Suppose next that infected hosts recover at rate $\gamma > 0$ and, after recovery, are initially immune but lose their immunity at some rate $\iota \geq 0$, after which they become susceptible to re-infection. In addition to the susceptible and infected states, let R denote the "recovered with acquired immunity" state and let $R(t)$ denote the mass of hosts in this state. The special case with permanent immunity ($\iota = 0$) is called the "SIR model," while that with no immunity ($\iota = \infty$) is the "SIS model." The more general case spanning both possibilities is the "SIRS model."

Epidemic dynamics in the SIRS model (with a fixed host population) are governed

[4]If each host meets another randomly-selected host at rate β, then each susceptible host meets an infected host at rate $\beta I(t)$. Note that such "meetings" corresponds to exposure *plus* successful infection. If a susceptible person exposed to the virus only becomes infected with probability $y \in (0, 1)$, then the rate of infection for susceptible agents is $\widehat{\beta} I(t)$, where $\widehat{\beta} = \beta y$.

[5]If transmission occurs mainly within a single age cohort (as might approximately be the case, say, for sexually-transmitted diseases), then the relevant host population consists of all those in the same age cohort, with death but no birth.

by the following system of differential equations

$$S'(t) = -\beta I(t)S(t) + \iota R(t) \tag{4}$$

$$I'(t) = \beta I(t)S(t) - \gamma I(t) \tag{5}$$

plus the adding-up condition $S(t) + I(t) + R(t) = 1$.

Each infected person on average exposes $R_0 = \beta L$ others during the course of their infection, where $L = 1/\gamma$ is average for length of time until recovery. R_0 (pronounced R-naught) is the pathogen's "basic reproduction number." An epidemic with $R_0 \leq 1$ is self-extinguishing, the prevalence of infection falling over time toward zero. By contrast, an epidemic with $R_0 > 1$ grows explosively and, so long as $\iota > 0$, persists with long-run steady state prevalence of infection $I^\infty = \frac{1-\gamma/\beta}{1+\iota/\gamma}$.[6]

In the SIR model, equation (4) simplifies to $S'(t) = -\beta I(t)S(t)$. When $R_0 > 1$, the prevalence of infection increases until the number of previously infected hosts $1 - S(t)$ reaches $1 - \gamma/\beta$, the level required for "herd immunity." More hosts become infected after that point, but at a decreasing rate, and some escape infection entirely. The fraction of hosts who are eventually infected is known as the "attack rate" and is always less than one; see Brauer et al. (2012) and Katriel and Stone (2012).

In the SIS model, equation (4) simplifies to $S'(t) = -\beta I(t)S(t)+\gamma I(t)$. When $R_0 > 1$, the prevalence of infection increases monotonically from approximately zero (when the pathogen first enters the host population) to a steady-state level $I^\infty \equiv 1 - \gamma_0/\beta$.

SCIRS/SCIR/SCIS model. Many bacterial pathogens colonize hosts for an extended period of time, an asymptomatic infection phase referred to as "carriage" (C) during which they may also be transmitted to new hosts.[7] For instance, enteric pathogens are transmitted through feces, whether or not they are currently causing harmful symptoms. Some viruses, including SARS-CoV-2 and HIV, can also transmit from carriage. To avoid confusion, I refer to a pathogen as "colonizing" its host while in carriage and "infecting" the host while causing symptomatic infection.

Suppose for simplicity that the transmission rate β and recovery rate γ are the same during carriage and infection, and let $\psi > 0$ be the rate at which the pathogen

[6] In the SIRS model, the system oscillates around this steady state while converging toward it. Other more complex variations may never reach a steady state (Hethcote et al, 2002).

[7] For pathogens that are unable to transmit during an initial quiescent phase, the SIRS model is typically extended to include a non-transmitting "exposed" state (E) prior to infection.

proceeds from carriage to infection. All other variables and parameters are the same as in the SIRS model.

Epidemic dynamics are governed by the following system:

$$S'(t) = -\beta(C(t) + I(t))S(t) + \iota R(t) \tag{6}$$

$$C'(t) = \beta(C(t) + I(t))S(t) - (\psi + \gamma)C(t) \tag{7}$$

$$I'(t) = \psi C(t) - \gamma I(t) \tag{8}$$

plus the adding-up condition $S(t) + C(t) + I(t) + R(t) = 1$.

From an economics point of view, it is useful to divide the recovered state R into two substates: R_C, for those who recovered most recently from carriage (without experiencing any symptoms); and R_I, for those who recovered from symptomatic infection. Note that $R'_C(t) = \gamma C(t) - \iota R'_C(t)$ and $R'_I(t) = \gamma I(t) - \iota R'_I(t)$.

The SCIRS model differs qualitatively from the others discussed so far, in that hosts do not immediately observe when they have been colonized. For instance, consider the special case of the "SCIR model" with permanent immunity after recovery. An agent who has not yet experienced disease by time t might currently be (a) susceptible (state S), (b) colonized (state C), or (c) recovered from carriage (substate R_C). On the other hand, agents know when they begin to experience disease (state I) and when they recover from disease (substate R_I).

Absent diagnostic testing, the epidemiological states $\{S, C, R_C\}$ constitute an *information set*, referred to as "not-yet-sick" and denoted by N, with $S(t)+C(t)+R_C(t)$ being the mass of not-yet-sick agents. For each state (or "health status") $h \in N$, let $p_h(t)$ denote agents' belief about the likelihood that their health status is h, conditional on being not-yet-sick. By Bayes' Rule, $p_h(t) = \frac{h(t)}{S(t)+C(t)+R_C(t)}$.[8]

Agent heterogeneity. Agents naturally differ in many ways that impact infection and transmission. For example: older people and those with co-morbidities may be more likely to die of infection; those with access to health care will receive supportive care (and curative treatment, if available) that reduces their subsequent transmissibility; those who have been vaccinated are less likely to become infected after exposure;

[8]The resulting belief dynamics are non-trivial. For instance, although fewer agents remain susceptible over time in the SCIR model, their likelihood $p_S(t)$ of being susceptible *conditional* on being not-yet-sick—a key consideration in the "social distanicng" game-theoretic models considered later—may rise or fall over time.

those with wider social networks are more likely to be exposed and to expose others; and so on.

Such heterogeneity is typically captured by defining "sub-compartments" (i.e., substates) of each of the basic epidemiological states and modeling the epidemic as following a Markov process with respect to this enriched state space. For instance, suppose that some of the population is vaccinated and that vaccination cuts in half the likelihood of developing infection each time that an agent is exposed to the virus.[9] This can be incorporated by dividing susceptible agents into two classes—those who are vaccinated (substate S_V) and those who are not (substate S_0)—with vaccinated people becoming infected at half the rate. In particular, in the SIR model, the differential equation $S'(t) = -\beta I(t)S(t)$ would be replaced by the pair of equations $S_0'(t) = -\beta I(t)S_0(t)$ and $S_V'(t) = -\beta I(t)S_V(t)/2$, with $S_V(t) + S_0(t) = S(t)$.

From an economic-theory perspective, an especially intriguing (and understudied) source of heterogeneity is *information*, especially: information about the epidemic, which itself may spread infectiously; information about one's own health status, creating new options for targeted treatment and control; and information about others' health status, enabling people to avoid infectious contact.

1.2 Manner of transmission

How a virus circulates among the host population, and what agents know about transmission, is essential to the trajectory of an epidemic.

"Fully mixed" vs. network models of transmission. In 1999, an American psychiatric facility was struck with an outbreak of *Mycoplasma pneumonia*, a leading cause of "walking pneumonia." The bacterium spread widely through the facility, but not through random meetings. Each patient was confined to a single ward, and hence unable to transmit the bacterium directly to those in other wards. However, some caregivers worked in multiple wards and, as such, served as links in a *transmission network* over which the bacterium spread throughout the facility. Meyers et al. (2003) modeled this network as a directed graph, based on detailed data collected by the Centers for Disease Control and Prevention (CDC) (Hyde et Hyde et al. (2001)), and estimated the rate of transmission along each edge of the graph—finding, for instance,

[9] If vaccination reduces the harm of infection and increases the degree or duration of immunity after recovery, then one would also want to divide the infection and recovery classes in an SIRS model.

that the bacterium was more likely to pass from caregivers to patients than vice versa.

Models with random meetings (referred to as "fully-mixed") are especially easy to analyze, in terms of a system of differential equations tracking how many hosts are in each epidemiological state at each point in time. Due to this simplicity, many applied-theory papers in the theoretical biology literature (and most of the recent Covid-inspired literature within economics) assume that transmission occurs via random meetings, or slight variations thereof with a small number of agent types.

Richer models with transmission over a network might seem hopelessly complex but, in fact, infection that spreads over a network can also be tractably analyzed. Newman (2002) characterizes epidemic dynamics for an arbitrary directed graph in terms of an adjacency matrix capturing exposure/transmission intensities between different agents or types of agents. Jackson and López-Pintado (2013) builds on this analysis, providing conditions on the adjacency matrix under which a new infectant ("an idea, a product, a disease, a cultural fad, or a technology") will spread from a small seed of initially-infected agents to a significant fraction of the population. See also Prakash et al. (2012), who provide thresholds for epidemic spread over a network.

Network models of transmission are appealing given their generality and tractability, and I expect the literature to shift in the near future more in this direction, especially given the increasing availability of individual-level data on physical mobility; see e.g., Fang et al. (2020) and Glaeser et al. (2020). However, in this review, I will follow the bulk of the existing literature and focus on models in which the infectant, here a biological virus, is spread through random meetings.

Awareness of contagious contact. The SARS outbreak of 2003 was quickly brought under control, with only about 8,000 people infected, in large part because the SARS-CoV-1 virus is (mostly) unable to transmit itself to new hosts until after causing severe symptoms, at which point those hosts are in the hospital and out of the general population. By contrast, SARS-CoV-2 can transmit from an asymptomatic state, making containment much more challenging absent diagnostics capable of determining who is carrying the virus.

This critical difference between SARS-CoV-1 and SARS-CoV-2 highlights an important modeling distinction in the economic epidemiology literature, regarding what hosts know about their own and/or others' health status. In particular, models differ on (i) whether there is an asymptomatic phase before symptomatic infection and (ii)

whether others can detect whether a host is infected, e.g., by measuring their temperature or performing a rapid diagnostic. It also matters whether people can credibly *disclose* their health status to others. As Paul Romer explained in a *New Yorker* article featuring his advocacy for greater testing: "I don't want to go back to the dentist's office in New York City until I know that he can show me a recent negative test, and he doesn't want me to come into his office until I can show him that I've got a recent negative test" (Chotiner 2020).

1.3 Agent decision-making process

Hosts (also called "agents") make many sorts of decisions that impact the trajectory of an epidemic, such as how frequently to wash their hands, whether to stay at home, whether to get tested, and so on. The way in which hosts are assumed to make decisions varies across the literature, falling into three main categories:

1. *mechanistic behavior*: agents' actions are determined by the current state of the epidemic

2. *rule-of-thumb behavior*: agents act to maximize an objective different (and typically simpler) than maximizing their own welfare

3. *individually-optimal behavior*: agents' actions are individually optimal given others' current and future behavior

All three approaches have their merits. Mechanistic models allow one to gain insight into the epidemiological properties of infection phenomena and lay the groundwork for future research that seeks to endogenize behavior. (Indeed, this is how the literature on infectious-disease dynamics has progressed, with about a century of work in mostly mechanistic models now growing in new directions that account for the dynamics of agent intention.) Models with individually-optimal behavior are useful as a fully-rational benchmark but, of course, may fail to predict actual outcomes if people are not the sophisticated reasoners that such models assume them to be. If so, rule-of-thumb models may come closer to capturing how people actually reason and process information and hence do a better job at predicting epidemic outcomes.

Example: social distancing. Consider a SCIR model in which not-yet-sick hosts decide at each point in time whether to reduce their likelihood of exposure to the virus

by staying away from others. Bootsma and Ferguson (2007) model such decisions by assuming that "individuals reduce their contacts as a function of the number of deaths occurring in the population in the previous time period."[10] However, because the risk of infection is not tied directly to the number of recent deaths, it is difficult to construct a reasonable objective function maximized by such a rule. Thus, in my phraseology, Bootsma and Ferguson (2007) assumes "mechanistic behavior."

Keppo et al. (2020) (and its predecessor Quercioli and Smith (2006)) approach behavioral adaptation during an epidemic in a different way, assuming that hosts' social distancing choices constitute a Nash equilibrium of a game in which each host acts *as if* maximizing an objective that depends only on the current epidemic state, their own distancing choice, and others' choices, i.e., agents are strategic yet also myopic. Since agents maximize an objective, but this objective does not correspond to their actual payoffs, Keppo et al. (2020) assumes "rule-of-thumb behavior." Having a simpler objective makes it easier to characterize the epidemic trajectory, relative to models that assume agents maximize the expected present value of their lifetime payoffs; see e.g., Reluga (2010), Farboodi et al. (2020), Toxvaerd (2020), McAdams (2020), and McAdams and Song (2020b), discussed in more depth later.

To convey ideas as clearly and simply as possible, I will focus the mathematical exposition of ideas here mostly in models with rule-of-thumb behavior and SIR or SIS transmission. However, with a few exceptions, most papers in the economics literature assume that agents are forward-looking optimizers, and some work within other transmission models. Table 1 categorizes several papers highlighted in this review, depending on (i) how agents make decisions (mechanistic, rule-of-thumb, or forward-looking), and (ii) the transmission model (SI, SIR/SIS/SIRS,[11] SCIR, or SCIS).

1.4 Economic impacts

Infectious disease creates economic harm directly through sickness, and indirectly as people take costly steps to avoid becoming sick.

[10] The economist John Cochrane took a similar approach in a May 2020 blog post, assuming that distancing varies with current infection prevalence or current death rate (Cochrane 2020).

[11] Some of these papers consider just the SIR model or just the SIS model, while others consider both separately or the more general SIRS model that encompasses both as special cases.

	mechanistic	rule-of-thumb	forward-looking
SI			Geoffard Philipson (1996) Kremer (1996) Geoffard Philipson (1997) Auld (2003) Chen (2004, 2006) Chan et al. (2016)
SIR / SIS / SIRS	Del Valle et al. (2005) Bootsma Ferguson (2007) Cochrane (2020)		Rowthorn et al. (2009) Reluga (2010) Chen (2012) Rowthorn Toxvaerd (2012) Farboodi et al. (2020) Alvarez et al. (2020) Bethune Korinek (2020) Toxvaerd (2020) Brotherhood et al. (2020)
SCIR		Keppo et al. (2020) McAdams Day (2020)	McAdams (2020) McAdams Song (2020)
SCIS	McAdams et al. (2019)		

Table 1: A selection of papers with dynamic economic-epidemiological models, categorized by their assumptions about decision-making and pathogen transmission.

Standard framework: Geoffard and Philipson (1996). The economic literature on infectious disease has for the most part followed Geoffard and Philipson (1996) in modeling the economic impacts of infection. In their approach, agents get instantaneous flow utility of the form $u(h_t, a_t)$, where h_t is an agent's health status and $a_t \in [0, 1]$ is their chosen level of "social activity," and seek to maximize the expected present value of their lifetime utility stream. (Equivalently, one can describe agents as choosing their "social distance" $d_t = 1 - a_t$.) A recent paper that takes this modeling approach is Farboodi et al. (2020). As they explain:

> "The assumptions that preferences u depend on social *activity* while disease transmission depends on social *interactions* are central to our view of social distancing. The former captures the idea that individuals value social activity (going to a restaurant, going for a walk, going to the office) and, absent health issues, are indifferent about whether other people are also engaging in social activity. On the other hand, if an individual goes for a walk and doesn't encounter anybody, they cannot get sick. Thus interactions are critical for disease transmission."

Under these assumptions, the "social-distancing game" at time t exhibits both positive

externalities (agents benefit when others distance more, due to reduced exposure risk) and strategic substitutes (agents have less incentive to be active themselves when others are more active). These properties of the game have significant theoretical implications, such as uniqueness of the epidemic trajectory, and policy implications, such as that a social planner always finds it optimal to tax social activity; see Rowthorn et al. (2009) and Rowthorn and Toxvaerd (2012).

Extension: multi-dimensional actions. Agents are typically modeled as making a one-dimensional choice—*either* (i) how much to curtail their public / social activities ("self-isolation") *or* (ii) how much to protect themselves during such activities ("vigilance")—but both sorts of decisions are relevant. For instance, a person might reduce how frequently they visit with friends *and* take precautionary steps such as wearing a mask when doing so. Monotone equilibrium comparative statics can be unintuitive in games with strategic substitutes (Roy and Sabarwal (2010)), especially with multidimensional actions, and social distancing is no exception. Salanié and Treich (2020) examine this issue in a static-game context. If self-isolation protects others but vigilance does not, they show that a social planner can increase social welfare by taxing vigilance. Why? Slightly reducing each agent's vigilance from its equilibrium level has a negligible (second-order) welfare effect due to the Envelope Theorem, but induces agents to increase their self-isolation, creating a first-order indirect benefit.

Extension: Complementarities in social activity. In the standard framework, social activity creates an *infection spillover* as more active agents are more likely to infect others with the virus, but there are no *economic spillovers* associated with activity. This seems reasonable if "activity" in akin to going for a walk. But what if "activity" is going to work in an office or playing a team sport? The risk of infection due to social activity is highest when others are active, but so is the benefit of being active yourself. Consequently, the social-activity game may exhibit strategic complements—and perhaps have multiple Nash equilibria—and it might sometimes be optimal to subsidize social activity. McAdams (2020), discussed later, is to my knowledge the first paper to extend the standard framework to allow for economic complementarities associated with economic activity.

Extension: Impact on search and matching. In the standard framework, social activity has no impact on who "matches" with whom but does impact the likelihood

of viral transmission due to each match. As Geoffard and Philipson (1996) explains:

> "Agents continuously meet one another over time, and upon each meeting, they must decide whether to engage in transmissive or protective behavior. If a susceptible agent chooses [transmissive behavior], he runs a risk of contracting the disease.."

This seems reasonable if "protective behavior" corresponds, say, to wearing a condom, but less so if it corresponds to abstaining from sexual activity altogether. A person looking for a sexual partner will find, not a random person, but someone else who is also looking. In a seminal contribution discussed more later, Kremer (1996) shows how such selection effects create the potential for multiple equilibrium epidemic trajectories, driven by a positive feedback between the composition of those looking for sex and the riskiness of doing so.

2 Equilibrium epidemic trajectory

This section examines how an epidemic unfolds over time, when agents decide for themselves whether to incur a cost to "distance" themselves from others. The analysis here synthesizes ideas in Toxvaerd (2020) and McAdams (2020), while also drawing on ideas in several other papers, especially Reluga (2010), Farboodi et al. (2020), and Keppo et al. (2020). A common theme in all these papers is that behavioral adaptation can have a profound impact on the epidemic trajectory.

2.1 Epidemic limbo

As the SARS-CoV-2 virus ripped through the United States in May 2020, two hair stylists in Springfield, Missouri continued working for several days despite having Covid symptoms. They saw 139 clients in total during that time but, when public-health officials scrambled to trace those contacts, they were surprised to find that none tested positive for the virus, and none developed symptoms. A subsequent field report published in the CDC's *Morbidity and Mortality Weekly Report* attributed this lack of transmission to the fact that the hair stylists and their clients wore face coverings throughout their interactions (Hendrix et al. 2020). Citing this report, CDC Director Robert Redfield said that "If we could get everybody to wear a mask right now, I really think in the next four, six, eight weeks, we could bring this epidemic under control."

Del Valle et al. (2005) examines the impact of behavioral change on the course of a viral epidemic. Within the context of a mathematical model of a biological attack resulting in a smallpox outbreak, they computed (i) how many people in a population of one million are ultimately infected and (ii) how long it takes until 99% of all infections have occurred, under various medical interventions and behavioral responses. In the baseline case with no intervention and no behavioral response, over 966,000 are infected and the outbreak lasts 307 days. By contrast, quick adoption of a behavioral response that reduces transmission by 90% reduces the number of cases to 306 over 208 days, while slower adoption of this response leads to 1,647 cases over 274 days.

Sustained and effective behavioral response speeds the end of the epidemic by driving down the basic reproductive number (R_0) of the virus, the average number of people exposed to the virus by each infected person. R_0 naturally changes over time, depending on public-health interventions and voluntary behavioral change. Smallpox's R_0 is estimated at being between 3 and 6; so, a behavioral change that reduces transmission by 90% will drive R_0 down to less than one and result in an exponentially decreasing number of cases. But as the number of cases falls, people's incentive to continue to "distance" themselves from others also naturally declines. Indeed, as the outbreak is squashed and almost no one in the community is infected, people have an incentive to relax, in which case the outbreak could flare up once again.

Game-theoretic models of social distancing have emerged to account for this feedback between the state of the epidemic and people's behavior. These models differ in several important respects, but a common feature emerges in many of them, what I refer to as *epidemic limbo*. People have an incentive to adopt precautionary measures once the epidemic has become sufficiently severe; so, the epidemic turns out to be not as bad as one would have predicted without accounting for behavioral response. However, as the epidemic wanes and there is less risk of being exposed, people eventually have an incentive to return to their usual behavior. Due to this self-limiting feedback, the epidemic can remain for an extended period of time in a limbo of intermediate severity: not so bad that *all* people take it seriously enough to distance themselves, but remaining enough of a threat that *some* people do so.

Fine and Clarkson (1986) was the first to provide a game-theoretic analysis of agents' incentives to take precautionary measures to avoid infection during an epidemic. More sophisticated dynamic analysis followed in the 1990s, with pioneering work by Philipson and Posner (1993), Kremer (1996), and Geoffard and Philipson

(1996), among others.[12] A recurring theme of this literature is that there is a limit to what can be achieved through voluntary precautionary measures, because of the negative feedback between infection prevalence and the incentive to take precautions. For instance, diseases that spread through random meetings[13] cannot be eradicated by costly vaccination alone, since the benefit of vaccination vanishes as the disease comes close to being eradicated; see e.g., Geoffard and Philipson (1997).

In the same way, there is a limit to how much voluntary social distancing can reduce the overall harm done during an epidemic. Reluga (2010) provides a game-theoretic model of a viral epidemic with forward-looking agents, in which agents decide at each point in time how intensively to distance themselves from others. Numerically solving the equilibrium epidemic trajectory for a relatively wide range of parameters, he finds that voluntary distancing reduces the overall harm done during the epidemic by at most 30%, relative to a no-distancing benchmark. That's a far cry from the 99.9% reduction in infection cases found by Del Valle et al. (2005), when assuming that agents engage in quick and sustained social distancing.

SIR model with rule-of-thumb vigilance. Consider an SIR model with transmission rate β and recovery rate γ, and hence basic reproduction number $R_0 = \frac{\beta}{\gamma}$. Suppose that agents have the option at each point in time to take an action (referred to by Keppo et al. (2020) as "vigilance") that has no effect on who they meet but reduces the likelihood of viral transmission during each given meeting. In particular, for simplicity, suppose that vigilance is a zero-one decision that reduces the instantaneous risk of transmission during a meeting (being infected or infecting others) to zero, at flow cost $c > 0$. Moreover, suppose that agents are rule-of-thumb decision-makers who act *as if* willing to pay $H > 0$ to avoid becoming infected. H is their "perceived harm" from being infected. (The actual economic harm associated with being infected varies over the course of the epidemic, as discussed later.)

Those who are infected have no individual incentive to be vigilant. Any susceptible agent who is not vigilant will therefore become infected whenever meeting an infected agent, which happens at rate $\beta I(t)$, creating expected perceived harm of $H\beta I(t)$ per unit time. So, each susceptible agent strictly prefers to be vigilant when $I(t) > \overline{I} \equiv \frac{c}{\beta H}$,

[12] The economists working on the game theory of infection prevention in the 1990s and 2000s were apparently unaware that epidemiologists had beat them to the punch. Yamin and Gavious (2013) were the first to cite Fine and Clarkson (1986) in an economics journal.

[13] Perisic and Bauch (2009) show that equilibrium eradication may be possible for diseases that spread over a persistent network.

strictly prefers not to be vigilant when $I(t) < \overline{I}$, and is indifferent when $I(t) = \overline{I}$.

The resulting equilibrium epidemic trajectory is uniquely determined and easily characterized, with three distinct phases, as illustrated in Figure 1 below.

Phase #1: Rising epidemic. Let t_1 be the first time at which $I(t_1) = \overline{I}$, i.e., $t_1 = \sup\{t : I(t) \leq \overline{I}\}$. If $\overline{I}$ is sufficiently high that $t_1 = \infty$, then no one is ever vigilant and the epidemic progresses as in a standard SIR model without behavioral adaptation. Otherwise, no one is vigilant and infection prevalence is strictly increasing up until time t_1, at which point the epidemic transitions to Phase #2.

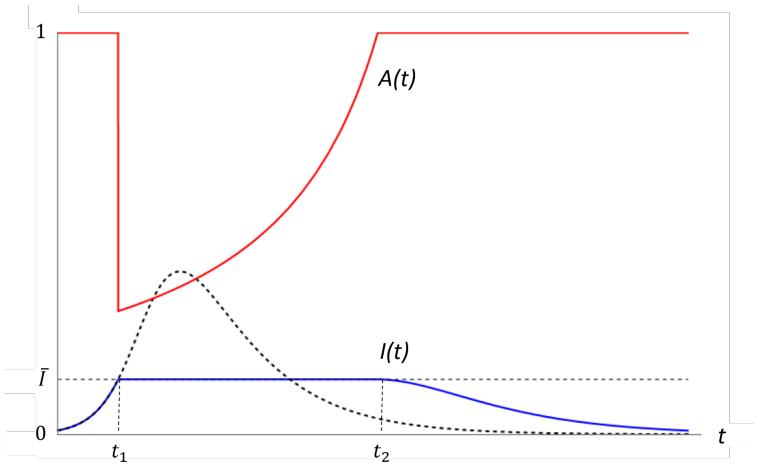

Figure 1: Infection prevalence in standard SIR model without any behavioral change (dotted line), equilibrium infection prevalence $I(t)$ (blue line) and fraction of agents $A(t)$ who are not vigilant (red line) in SIR model with rule-of-thumb vigilance. This is a slightly modified version of Figure 3 in Toxvaerd (2020).

Phase #2: Epidemic limbo. Once the mass of infections hits $\overline{I}$, some but not all susceptible agents must choose to be vigilant, just enough so that the mass of infections remains equal to $\overline{I}$. This requires exactly fraction $1 - \frac{\gamma}{\beta S(t_1)}$ of susceptible agents to be vigilant, meaning that fraction $A(t) = \frac{\gamma}{\beta S(t_1)}$ are not vigilant (A is mnemonic for "active"); note that $A(t)\beta S(t_1) = \gamma$. If more susceptible agents than this were vigilant, the mass of infections would fall and none would want to be vigilant, a contradiction. Similarly, if fewer were vigilant, the mass of infections would rise and all would want to be vigilant, another contradiction.

The resulting epidemic dynamics are characterized by the system:

$$S'(t) = -\beta \overline{I} S(t) A(t) = -\gamma \overline{I} \tag{9}$$

$$I(t) = \overline{I} \tag{10}$$

$$A(t) = \frac{\gamma}{\beta S(t_1)} \tag{11}$$

plus the usual adding-up condition $S(t) + I(t) + R(t) = 1$. Since infections clear at rate γ, the flow of agents out of the infected state is $\gamma \overline{I}$. Equilibrium social distancing $1 - A(t)$ is just enough so that the flow of new infections also equals $\gamma \overline{I}$. Note that, since the mass of susceptible agents $S(t)$ falls over time, agents distance less and less throughout the "limbo" phase of the epidemic, i.e., activity $A(t)$ is increasing.

Let t_2 be the time at which $S(t) = \frac{\gamma}{\beta}$. This is moment at which the population as a whole achieves "herd immunity," in the sense that the mass of infected agents will henceforth fall over time *even if* no one distances.[14]

Phase #3: Declining epidemic. After time t_2, no one is vigilant and the mass of infections declines over time, with $I'(t_2) = 0$, $I'(t) < 0$ for all $t > t_2$, and $\lim_{t \to \infty} I(t) = 0$.

Vigilance versus self-isolation. Suppose that, instead of deciding whether to take protective actions (such as wearing a mask) to prevent transmission during each given meeting, agents decide whether to avoid such meetings altogether. The effect of such "self-isolation" on others depends on whether isolating oneself reduces the number of encounters that *others* experience.

The most common assumption in the literature, following Geoffard and Philipson (1997), is that isolating yourself causes transmission events that would have happened not to happen at all. For instance, suppose that a susceptible person decides to go for a walk in the park, and that half of all susceptible people stay home (but all infected and recovered people go out). That person will "cross paths" with half as many susceptible people but the same number of infected people, and hence be at the same amount of risk as if no one had stayed home. The overall flow of new exposures in this case when half of susceptible agents stay home is therefore the same as in the earlier "vigilance

[14] Herd immunity is achieved in the SIR model with random meetings once mass $1 - \frac{\gamma}{\beta}$ of hosts have been exposed, leaving mass $\frac{\gamma}{\beta}$ still susceptible. In an uncontrolled epidemic, herd immunity is achieved at the moment when infection prevalence is at its peak. Distancing reduces the overall number of infections by reducing how many people are infected *after* herd immunity is achieved.

model" with half of all susceptible agents choosing to be vigilant.

This modeling approach has been widely adopted in the recent Covid-inspired theoretical literature, but misses a key feature of the social context of disease transmission captured originally by Kremer (1996). In many settings, people get *partner-unspecific benefit* from social interactions and, because of this, will seek out alternative partners if the person they would have otherwise matched with is absent. To see the point, suppose that the person going to the park is there to play a game of pickup basketball.[15] Fewer teams will form, but whoever is there will still form teams and play. Moreover, because only susceptible people stay home, the people playing will be more likely to be infected than if no one had stayed at home. In this way, social distancing by susceptible people makes other susceptible people *more* likely to be infected when not distancing themselves.

In Kremer (1996)'s pioneering model, each agent in an atomless population decides how many interactions they want to have, and then agents are randomly matched in a way so that each agent has the number of interactions that they desire. For example, suppose that the host population consists of two equally-likely types—"high-activity agents (H type)" who are fully active ($a_i = 1$) and "low-activity agents (L type)" who cut back by half ($a_i = 1/2$)—and that, if everyone were fully active, everyone would encounter 6 H types and 6 L types per unit time on average. In Kremer (1996), H types encounter 8 H types and 4 L types per unit time, while L types encounter 4 H types and 2 L types. By contrast, in the more commonly-used approach, H types encounter 6 H types and 3 L types, while L types encounter 3 H types and 1.5 L types.

The key difference is that, in Kremer (1996), people's distancing decisions impact not just how many matches occur, but who matches with whom. In particular, one type of agent staying out of the "matching market" makes it more likely that market participants will match with other types. Social distancing by susceptible agents therefore creates a positive feedback: the more that susceptible people distance, the more that other susceptible people want to distance. The game among susceptible agents therefore exhibits strategic complements and, as such, can possess multiple equilibria.

[15]The same issues arise in many other contexts. For instance, in Kremer (1996)'s original example, someone going to a brothel for sex is going to have sex with *someone*, but the odds that that person is HIV-positive depends on the relative likelihood that HIV-positive and HIV-negative women will being working at the brothel.

Extension: economic complementarities. Building on a model introduced in McAdams (2020), McAdams and Song (2020b) explores the impact of economic complementarities on the equilibrium epidemic trajectory. Each agent i is assumed to get flow economic payoff of the form

$$u(a_i; A) = \alpha_0 + \alpha_1 a_i + \alpha_2 a_i A, \tag{12}$$

where a_i is agent i's activity level, A is the population-wide average activity level, and $\alpha_0, \alpha_1, \alpha_2 \geq 0$ are parameters capturing the importance of isolated, public non-social, and public social activities, respectively, for agent welfare.[16] (Those who are sick incur an additional cost and may or may not be incapacitated.)

Interpretation of parameters: α_0 captures the baseline level of benefits that a well agent gets while quarantined in the home; α_1 captures the extra benefits associated with being able to leave the home, e.g., the extra pleasure and health benefit of walking outside; and α_2 captures the extra benefits associated with sharing the same physical space with others, e.g., hugging a friend rather than just talking on the phone. These parameters can be changed in many ways. For instance, a restaurant service that delivers safely-prepared fresh-cooked meals would increase α_0 and reduce α_2, as would improved virtual-meeting technology that enhances remote collaboration.

The presence of economic complementarities ($\alpha_2 > 0$) changes the qualitative features of the social-distancing game played by agents throughout the epidemic, in two main ways. First and most importantly, there can be multiple equilibrium trajectories. The course of the epidemic may therefore depend on coordinating mechanisms (e.g., public announcements) that induce agents to play one equilibrium rather than another. Second, as people begin to distance, there is a positive feedback as others' inactivity reduces agents' incentive to be active themselves. For instance, entrepreneurs who share an incubator space might have a strong incentive to work in their office so long as everyone else is doing so, to share ideas during impromptu encounters, but not once most other people are working from home. Similarly, there is less reason to go to a shopping area when most stores are closed, less benefit from operating a production facility if suppliers and shut down, and so on.

McAdams and Song (2020b) has forward-looking agents but, to gain intuition, it is helpful to consider the impact of economic complementarities in the SIR model

[16] Assuming linear payoffs simplifies equations but is not essential. The analysis can also be easily modified to allow for congestion effects ($\alpha_2 < 0$).

considered above, with rule-of-thumb decision-makers who have perceived harm H from being infected. However, now assume that the cost of self-isolation takes the form $c(A) = \alpha_1 + \alpha_2 A$, where A is the fraction of the overall population that remains socially active. In this context, each agent has a dominant strategy to self-isolate whenever infection prevalence $I(t)$ is greater than $\overline{I} \equiv \frac{c(1)}{\beta H} = \frac{\alpha_1+\alpha_2}{\beta H}$ and a dominant strategy *not* to self-isolate whenever $I(t)$ is less than $\underline{I} \equiv \frac{c(0)}{\beta H} = \frac{\alpha_1}{\beta H}$. When $I(t) \in \left(\underline{I}, \overline{I}\right)$, then are multiple equilibria, including one in which all susceptible agents isolate and another in which no one isolates.

In the equilibrium with the most infection, $I(t)$ increases until time t_1 at which $I(t_1) = \overline{I}$, when agents are indifferent whether to self-isolate. Immediately after time t_1, at least fraction $1 - \frac{\gamma}{\beta S(t_1)}$ of susceptible agents must isolate (by the same argument as before). But then the economic benefit of activity falls, from $\alpha_1 + \alpha_2$ to $\alpha_1 + \alpha_2 \frac{\gamma}{\beta S(t_1)}$, causing agents to strictly prefer to self-isolate. Everyone isolates and the equilibrium prevalence of infection falls precipitously right after time t_1—very unlike the "epidemic limbo" that prevails in models without complementarities.

The most extreme version of this phenomenon arises when all of the benefit of public activity comes from social activity, i.e., when $\alpha_1 = 0$ but $\alpha_2 > 0$. Once infection prevalence hits $\overline{I}$ at time t_1, each susceptible agent is indifferent whether to isolate when no one else is doing so. But then as some people start isolating, all agents strictly prefer to isolate and the unique equilibrium has everyone hunkered down in isolation, getting flow utility $u(0;0) = \alpha_0$ from isolated activities alone. Such sudden collective voluntary isolation stops the virus in its tracks but, so long as there is even a small amount of virus in circulation, it remains an equilibrium for everyone to remain at home. In this context, a social planner can increase welfare by subsidizing some agents to re-engage in social activity, to prod them out of the no-activity equilibrium.

Extension: altruism. Altruism can also have a dramatic effect on equilibrium epidemic outcomes. Suppose that people are willing to pay $B \geq 0$ to avoid causing someone else to be infected, and recall that $d_S(t)$ is the share of susceptible agents who distance. Each infected agent encounters a susceptible agent at rate $\beta S(t)(1 - d_S(t))$, and hence gets expected altruistic benefit $B\beta S(t)(1 - d_S(t))$ when isolating themselves from others. Since self-isolation costs $c > 0$, infected agents strictly prefer to isolate at time t if and only if $S(t)(1 - d_S(t)) > \overline{S} \equiv \frac{c}{\beta B}$.

Early during an outbreak, susceptible agents choose not to isolate because infection

is rare (shown earlier); that is, $d_S(t) = 0$. Altruistic agents who become infected early on while infection is rare therefore choose to isolate if and only if $B > \frac{c}{\beta}$. The outbreak will therefore die out with only a few people infected ... and never reach the "epidemic limbo" phase.

Extension: asymptomatic infection. Consider the SCIR model described in Section 1.1. At time t, each susceptible agent who does not distance becomes infected at rate $\beta(C(t) + I(t))$, where $C(t)$ and $I(t)$ are, respectively, the mass of agents with asymptomatic infection ("carriage") and symptomatic infection ("sickness"). Each agent who has not yet gotten sick by time t therefore has an incentive to distance so long as $p_S(t)H\beta(C(t) + I(t)) > c$, where $p_S(t) = \frac{S(t)}{S(t)+C(t)+R_C(t)}$ is the likelihood of being susceptible conditional on being not-yet-sick at time t.

For any given prevalence of infection, not-yet-sick agents have less incentive to distance in the SCIR model than susceptible agents do in the SIR model, due to their uncertainty about whether they remain susceptible, i.e., due to the fact that $p_S(t) < 1$. Consequently, (i) a smaller fraction of not-yet-sick agents distance in the SCIR model for any given prevalence of infection, and (ii) agents wait longer in the SCIR model before they begin distancing, i.e., they do not distance until $C(t) + I(t)$ exceeds a threshold strictly higher than $\overline{I}$. This implies, as one would expect, that more people ultimately become infected when the pathogen has asymptomatic spread.

2.2 Forward-looking behavior

The analysis thus far has assumed that agents are (myopic) rule-of-thumb decision-makers, whose behavior depends only on the current prevalence of infection, their own perceived harm from being infected, and, in the SCIR model, their own likelihood of being susceptible. How does agent behavior and the epidemic trajectory change when agents are forward-looking optimizers?

Consider first an SIR model as in Farboodi et al. (2020) and Toxvaerd (2020), in which agents know once they have become infected and there are no economic complementarities. Moreover, for simplicity and to highlight key ideas as clearly as possible, assume that agents make a discrete choice whether to isolate themselves fully or not distance at all. In particular, suppose that agents seek to maximize the expected present value of their future lifetime payoff stream ("continuation welfare"), incur flow cost $s > 0$ while sick, incur flow cost $c > 0$ while self-isolating, and use interest rate

$r > 0$ to discount future payoffs. As in the analysis surrounding Figure 1, let $\beta > 0$ be the transmission rate absent any distancing, let $\gamma > 0$ be the infection recovery rate, and assume that agents are not altruistic and self-isolation reduces the risk of transmission to zero.

Equilibrium social distancing. For each epidemiological state $\omega \in \{S, I, R\}$, let $\Pi_\omega(t)$ be the continuation welfare of agents in state ω at time t. Susceptible agents are willing to pay $H(t) \equiv \Pi_S(t) - \Pi_I(t)$ in order to avoid becoming infected. Since they become infected at rate $\beta I(t)$ when not distancing, susceptible agents strictly prefer to distance if and only if $I(t) > \overline{I}(t) \equiv \frac{c}{\beta H(t)}$, much as in the previous rule-of-thumb analysis but now with an endogenous time-varying cost $H(t)$ of being infected.

Infected agents' welfare. Once someone has become infected, they will choose thereafter not to distance themselves, earn flow payoff $-s < 0$ while infected, and then earn zero flow payoff once recovered. Since recovery occurs at rate γ, each infection has likelihood $e^{-\gamma L}$ of lasting longer than length of time L. Given discounting at interest rate r, the expected present value of the sickness costs incurred during a given infection is therefore

$$\Pi_I = -s \int_0^\infty e^{-(r+\gamma)L} \mathrm{d}L = \frac{s}{r+\gamma} \tag{13}$$

and does not depend on the time t; in particular, $\Pi_I(t) = \frac{s}{r+\gamma}$ for all t.

Susceptible agents' welfare. The continuation welfare of a susceptible agent varies over time, and in a non-monotone fashion. Early in the epidemic while infection is rare, susceptible agents do not distance and face little immediate risk of exposure. However, as time passes, the risk of soon being infected grows exponentially and the epidemic looms larger in agents' welfare considerations. Over this timeframe, susceptible agents' welfare is declining over time. On the other hand, near the end of the epidemic when infection is once again relatively rare, agents will once again choose not to distance. The difference is that now, as time passes, susceptible agents' *remaining* risk of becoming infected falls as the epidemic continues to fade, causing their continuation welfare to increase.

Equilibrium trajectory. An equilibrium epidemic with forward-looking agents typically follows a similar[17] three-part trajectory as in the previous rule-of-thumb analysis: (i)

[17]Other patterns are possible. For instance, if distancing is only partially effective at limiting

first, a period of uncontrolled growth in the prevalence of infection until time t_1; (ii) second, an intermediate period until time t_2 in which some but perhaps not all agents self-isolate; and (iii) a final period after t_2 in which no one distances, but the prevalence of infection continues to fall because "herd immunity" has been achieved.

The main difference is that the prevalence of infection $I(t)$ is no longer constant but *falls* over time during the intermediate phase.[18] To gain intuition, note that agents are indifferent whether to incur the cost c to self-isolate. A susceptible agents' continuation welfare at time $t \in (t_1, t_2)$, $\Pi_S(t)$, is therefore the expected present value associated with incurring cost c all the way from time t until t_2 and then getting lump-sum payment $\Pi_S(t_2)$ at time t_2. Moreover, because agents strictly prefer not to distance after time t_2, they are obviously better off than if they had to pay c in perpetuity. Consequently, $\Pi_S(t)$ is strictly *increasing* from time t_1 until time t_2. That implies that the harm of infection $H(t) = \Pi_S(t) - \Pi_I$ is also increasing in t. In order for agents to be indifferent whether to self-isolate, the risk of infection must therefore be decreasing in t, which requires that fewer people are infected over time.

Impact of a vaccine or treatment. Those who are vaccinated are less likely to become infected for any given level of activity and hence will choose to be more active than otherwise. If the vaccine is imperfect, the overall effect of such "risk compensation" can be to increase the amount of infection; see e.g. Hoy and Polborn (2015) and Talamas and Vohra (2018). Similarly, treatments that reduce the harm of infection may lead to greater transmission, as people are less cautious about avoiding infection.

Even before a vaccine or treatment becomes available, the anticipation of its arrival can change behavior. Suppose that agents are forward-looking optimizers and that they expect a *perfect vaccine* to become available at time $T > 0$. Anyone exposed at or after time T will not become infected; so, susceptible agents have no reason to distance and will not become sick, i.e., $\Pi_S(t) = 0$ for all $t \geq T$. Just before time T, the harm of being infected, $H(t) = \Pi_S(t) - \Pi_I(t) \approx \frac{-s}{r+\gamma}$, is therefore as large as it can ever be. This gives agents a relatively strong incentive to distance just before the vaccine becomes available—the intuition being that they have nearly "made it" to the point when they won't need to distance any longer.

transmission events, then there can be periods in which all agents distance, interspersed with periods in which some but not all distance.

[18] Toxvaerd (2020) states that the prevalence of infection is constant over the intermediate phase. This is incorrect, as I have confirmed through an email correspondence with the author. Fortunately, the underlying error is easily corrected and his other main qualitative findings remain.

What if, instead, a *perfect treatment* becomes available at time $T > 0$. Anyone infected at or after time T will not suffer; so, as with a perfect vaccine, susceptible agents have no reason to distance and $\Pi_S(t) = 0$ for all $t \geq T$. The main difference is that those who are infected shortly before time T now also do not suffer much at all. In particular, $\Pi_I(t) \approx -s(T-t) \approx 0$ for all t slightly less than T, implying that the harm of infection is approximately zero. Thus, agents respond quite differently to news of a coming vaccine versus a coming treatment.

3 Lockdown policies

In early April 2020, the *Wall Street Journal* reported that "U.S. counties under lockdown orders ... represent nearly 96% of national output" and that "at least one-quarter of the U.S. economy has suddenly gone idle ... an unprecedented shutdown that economists say has never occurred on such a wide scale." Many Americans chafed under these restrictions and called for them to be eased, including President Donald Trump,[19] but the case for government intervention of some kind was strong. States like Florida and Texas whose governors initially resisted public-health measures to slow the virus' spread were soon overwhelmed and forced to follow suit.

Consider the simple example analyzed in Section 2.1. During the "limbo phase" of the epidemic, when some but not all agents distance, agents have a choice between (i) suffering economically by distancing or (ii) facing the risk of infection by not distancing. A key feature of the equilibrium epidemic during this period is that agents are indifferent whether to distance. Indeed, they suffer more than if everyone had been completely locked down since, with a lockdown, there would at least have been less infection when activity is allowed to resume.

What type of lockdown policy is socially optimal? The answer depends on three key questions about the economic-epidemiological environment: (i) can infections be identified?, (ii) is social distancing voluntary?, and (ii) are there economic complementarities associated with economic activity?

[19] After relatively small protests against state-ordered lockdowns in mid-April, Trump tweeted "LIBERATE MINNESOTA" and "LIBERATE MICHIGAN" and criticized the Democratic governors of these and other states. Two weeks later, protesters armed with assault rifles occupied the Michigan statehouse (DeBrabander 2020).

What if infections can be identified? If public-health authorities can identify who is infected, then the epidemic can potentially be halted in its tracks by imposing a *targeted lockdown* that only restricts the activity of infected agents.[20]

Suppose for the moment that lockdown reduces an agent's social activity to fraction $0 \leq \psi < 1$ of its normal level and that, absent lockdown, all agents would engage in their normal level of activity. If $\psi < \frac{\gamma}{\beta}$, then locking down all infected agents reduces the basic reproduction number of the virus from $R_0 = \frac{\beta}{\gamma}$ to $R_0 = \frac{\psi\beta}{\gamma} < 1$, extinguishing the epidemic. What if $\psi > \frac{\gamma}{\beta}$? Locking down infected hosts is then not enough to stop the virus from infecting a substantial fraction of the population. However, since the virus spreads *as if* it has transmission rate $\psi\beta$ rather than β, fewer people will be infected and, if recovery does not confer immunity, the steady-state prevalence of infection will be lower.

Bethune and Korinek (2020) characterize the socially-optimal targeted lockdown policy, assuming that the social planner can dictate each agent's level of activity $a_i \in [0, 1]$ and that agents get concave flow utility $u(a)$ from activity-level a, with $u'(1) = 0$ so that small reductions in activity have a negligible impact on agent welfare. Figure 2 illustrates their findings in the SIS model, the panels on the left showing agents' equilibrium activity as a function of current infection prevalence, for a high-cost disease such as Covid (top) and a low-cost disease such as the common cold (bottom), and the panels on the right showing the corresponding socially-optimal activity levels.

In the decentralized equilibrium without any forced lockdown, only susceptible agents reduce their activity and much more so in the case of a high-cost disease, as expected. For the high-cost disease, the social planner always imposes a sufficiently severe lockdown policy so that the virus' basic reproduction number R_0 is held below one, so that the prevalence of infection falls over time. However, this is achieved through a *combination* of restrictions on infected and susceptible agents, with restrictions on infected and susceptible agents growing more and less severe, respectively, as the prevalence of infection falls over time. In the long run, as infection prevalence vanishes to zero, susceptible agents are unrestricted and only infected agents are restricted.

The most interesting case is when the disease is less severe, as illustrated in the lower panels of Figure 2. If infection is sufficiently rare—for instance, if there is an

[20]Locking down all infected hosts and no susceptible hosts has the same effect as locking down all susceptible hosts and no infected hosts. Thus, even after most people have become infected, an epidemic's *further* expansion can be stopped if the uninfected hunker down in isolation—a key plot point in most zombie-apocalypse movies.

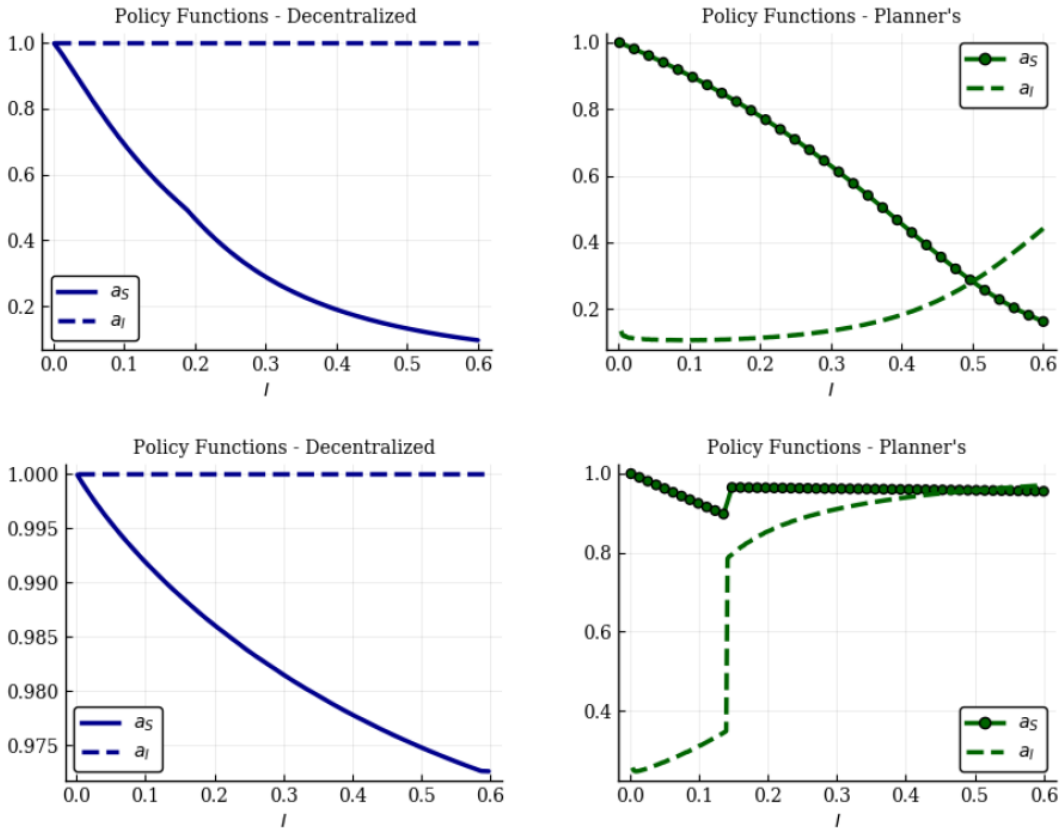

Figure 2: Equilibrium ("decentralized") and socially-optimal activity levels for susceptible and infected agents in the SIS model, for a severe disease (top panels) and a mild disease (bottom panels). Courtesy of Zachary Bethune and Anton Korinek.

outbreak of a novel virus—then the social planner finds it optimal to target infected agents with severe restrictions, enough to reduce R_0 below one and drive the prevalence of infection to zero. Although these restrictions are burdensome on those who become infected, the mass of agents who ever become infected is small; so, the overall burden of the lockdown policy is small, relative to the harm that would have been done if the outbreak had been allowed to grow into an epidemic. However, once infection prevalence passes a threshold (about 0.15 in Figure 2), the lockdown "cure" is worse than the disease and the social planner finds it optimal to blunt the epidemic but not stop it.

What if infections cannot be identified? Suppose that the disease spreads mainly during an early asymptomatic phase and that there is no way to determine who is

asymptomatically infected. In this case, the social planner is constrained to uniform lockdown policies. Alvarez et al. (2020), Bethune and Korinek (2020), and Rowthorn and Maciejowski (2020) characterize the optimal uniform lockdown in an SIR model, under various parameter conditions motivated by the Covid-19 outbreak. In Alvarez et al. (2020), "the optimal policy prescribes a lockdown starting four weeks after the outbreak, and covering 45% of the population after 8 weeks." Similarly, in Rowthorn and Maciejowski (2020), "the lockdown lasts 5.3 weeks and brings the disease under control quite soon, although not before millions of people have been infected and many thousands have died." On the other hand, Bethune and Korinek (2020) find that it is optimal not to allow the infection ever to become common. In all three cases, continued "relaxed" restrictions are maintained for an extended period, to keep R_0 close to one and prevent a resurgence of the epidemic.

Two features[21] of optimal lockdown here are worth emphasizing. First, the social planner waits until the outbreak is sufficiently large before imposing any restrictions. To see why, suppose that the social planner were to severely constrain activity while infection is extremely rare. Infection would grow even more rare, but would not disappear entirely. Once the lockdown is eased, the virus will then come roaring back and, since few have been exposed, the population as a whole will then be in the same position as if they had not just locked down their economy.

Second, the social planner applies relatively intense restrictions and then eases up over time. Why? Because lockdowns are costly and must be applied across the board, it is never optimal to maintain such restrictions forever. But then that means that the virus must eventually infect enough people for the population to achieve herd immnuity, after which the rate of new infection declines over time. In an uncontrolled epidemic, the number of people who are infected reaches its peak at the point in time when herd immunity is achieved, and the epidemic blows through the herd-immunity threshold; in the end, many more are infected than needed to be. By constraining activity while the epidemic would otherwise be raging, the social-planner can ensure that relatively few people are infected after herd immunity is reached.

Of course, both of these findings hinge critically on the assumption that there is no way to detect infection. Should testing for asymptomatic infection be available, the social planner typically finds it optimal to isolate those who are found to be infected—

[21]These features are evident in Alvarez et al. (2020) and Rowthorn and Maciejowski (2020), but also present in Bethune and Korinek (2020), just with a very low prevalence threshold for lockdown and a very long time until herd immunity is reached.

not just to slow or smooth out the epidemic trajectory, but to eliminate the virus entirely.

Discussion: nuances around testing and heterogeneity. Testing to discover asymptomatic carriers is essential, as it empowers public-health authorities to get the epidemic under control without a large number of people being infected. However, as Acemoglu et al. (2020b) and Deb et al. (2020) have noted, testing without appropriate incentives can lead to perverse outcomes. As Acemoglu et al. (2020b) explains:

> "Testing enables the isolation of infected individuals, slowing down the infection. But greater testing also reduces voluntary social distancing or increases social activity, exacerbating the spread of the virus. We show that the effect of testing on infections is non-monotone. This non-monotonicity also implies that the optimal testing policy may leave some of the testing capacity of society unused."

Further nuances arise when agents have heterogeneous types. Brotherhood et al. (2020) and Acemoglu et al. (2020a) analyze an extension of the SIR model allowing for multiple agent types that differ in their health status (e.g., "old" are more likely to die if infected) and/or in their connections with others (e.g., "young" interact mostly with themselves). A shared insight that emerges in these papers is that it may be optimal to allow infection to spread more widely among the young.[22] This can be good for the young, since they are not as burdened economically, but can also be good for the old as herd immunity can be reached with fewer of the old becoming sick.[23]

Are there economic complementarities associated with economic activity? McAdams and Day (2020) explores the implications of economic complementarities on lockdowns. Consider an SCIR model in which a lockdown can be imposed on not-yet-sick agents. In this context, restricting others' activity benefits agents by reducing pathogen transmission ("health spillover") but also harms agents economically, by reducing the benefit that agents get from social activity ("economic spillover"). McAdams and Day (2020) find that, if transmission only occurs during carriage, then

[22] An earlier literature grapples with the question of how to optimally devote limited infection-prevention resources. In a multi-population SIS model, Anderson et al. (2012) show that it can be optimal to focus entirely on just one population.

[23] As these authors note, this conclusion hinges on the assumption of unlimited hospital capacity. If the young take up limited hospital beds, then more of the old could die even as fewer are infected.

the health spillover always dominates the economic spillover, in the sense that agents benefit by having a lockdown imposed *before* they would begin to distance voluntarily. On the other hand, if transmission occurs primarily during infection, then not-yet-sick agents may sometimes benefit by being forced to be active when they would not otherwise voluntarily choose to do so.

4 Pathogen eradication

When can a pathogen be eradicated from a host population? An extensive literature explores the potential to eradicate diseases through vaccination; see e.g., Geoffard and Philipson (1996) and Chen and Toxvaerd (2014) on vaccination in fully-mixed models and Perisic and Bauch (2009) in a model with transmission over a network. Here I focus on complementary approaches that leverage *treatment* and/or *infection control*.

Section 4.1 follows Rowthorn and Toxvaerd (2012), focusing on the special case of their model in which a highly-effective treatment is available and there are no prevention options. The main finding is that an infectious disease *can* potentially be eradicated through treatment alone, but only if the treatment is introduced while infection remains sufficiently rare.

Section 4.2 follows McAdams et al. (2019) in considering a context in which multiple strains of the same pathogen co-exist, some of which are resistant to antibiotic treatment. The main finding is that resistant strains *can* potentially be eradicated, thereby restoring the effectiveness of existing antibiotics to which resistance has already emerged, but only if (i) diagnostics are available to detect resistant infection and (ii) public-health resources can be deployed specifically to reduce transmission of resistant strains.

4.1 Eradicating a disease

Consider an SIS-model infectious disease for which a costly treatment is available, and suppose that this treatment is sufficiently effective that the disease could be eradicated if every infection were treated. Here I focus on two central questions. First, is it socially optimal to eradicate the disease? Second, can the disease be eradicated in a decentralized (Nash equilibrium) setting in which each infected person incurs the cost of treatment and decides whether to receive treatment?

Epidemiological model. A pathogen circulates among a unit-mass population of hosts according to a standard SIS model. While infected, each host may or may not receive treatment that speeds recovery. Let $S(t)$ and $I(t)$ denote the mass of susceptible and infected hosts, respectively, with $S(t) + I(t) = 1$ because the population has unit mass. Let $I_Y(t)$ and $I_N(t)$ denote the mass of infected hosts who are treated (Y for "yes") or not treated (N for "no"). Infected hosts recover at baseline rate γ_0 if untreated or at faster rate $\gamma_A > \gamma_0$ if treated. The overall flow of newly-recovered agents therefore equals $\gamma_0 I_N(t) + \gamma_A I_Y(t) = \gamma_0 I(t) + (\gamma_A - \gamma_0) I_Y(t)$. The resulting epidemiological dynamics are characterized by the differential equation

$$S'(t) = -\beta I(t) S(t) + \gamma_0 I_N(t) + \gamma_A I_Y(t) \tag{14}$$

with $I(t) = I_Y(t) + I_N(t)$ and $S(t) + I(t) = 1$. If all infections are treated, equation (14) simplifies to $S'(t) = I(t)(\gamma_A - \beta(1 - I(t)))$.

Eradication through treatment is possible when $\gamma_A > \beta$, since then $S'(t) > 0$ (and hence $I'(t) < 0$) no matter how many agents are currently infected. I will focus here on the case when $\gamma_A > \beta > \gamma_0$, meaning that the disease will be eradicated if all infections are treated but not if no infections are treated.

Untreated infections last on average for length of time $L_0 = \frac{1}{\gamma_0}$. Each untreated infected person therefore on average exposes $R_0 = \frac{\beta}{\gamma_0} > 1$ others during the course of their infection. Should all infections be left untreated, the prevalence of infection will increase from approximately zero (when the pathogen first enters the host population) to a steady-state level $I_0^\infty \equiv 1 - \frac{\gamma_0}{\beta}$. On the other hand, if all infections are treated, then $R_0 = \frac{\beta}{\gamma_A} < 1$ and the long-run steady state prevalence of infection is zero.

To close the model, it remains to characterize when infected hosts receive treatment. To do so, one needs to overlay an economic model on top of the epidemiological model.

Economic model. Each host i receives flow payoff $\pi_i(t) = 0$ when susceptible, $\pi_i(t) = -s$ when sick and untreated (s is "sickness cost"), or $\pi_i(t) = -s - c$ when sick and treated (c is used here for "treatment cost"), and discounts payoffs with respect to interest rate $r > 0$. While infected, each agent decides whether or not to receive treatment.

Maximizing social welfare. Across the entire population, the disease does harm $sI(t) + cI_Y(t)$ at each time t. Treating more infections at time t raises the costs that

are immediately incurred, but reduces the future prevalence of infection. Rowthorn and Toxvaerd (2012) show that the socially-optimal treatment policy is bang-bang. In particular, a threshold $\widehat{I}$ exists such that it is socially optimal to treat all infections whenever $I(t) < \widehat{I}$ but to leave all infections untreated whenever $I(t) > \widehat{I}$.

The bang-bang nature of optimal treatment arises because treatment is more socially valuable when a disease is rarer. To gain intuition, note that the prevalence of infection grows exponentially during the early phase of an outbreak while infection is relatively rare, but that the rate of growth declines as more people become infected. Treating any given infection therefore prevents more infections when the overall prevalence of infection is lower.

Nash-equilibrium outcomes. What if individuals decide for themselves whether or not to be treated? The resulting game exhibits strategic complements, with each agent having more incentive to be treated if they believe that others (now and in the future) are more likely to seek out treatment as well. To gain intuition, note that agent i's benefit of recovering from infection depends on the risk of re-infection: the lower the rate of re-infection, the longer that agent i expects to remain infection-free and hence the more valuable it is to recover. If others are more likely to be treated (now and in the future), then fewer people will be spreading infection once agent i eventually recovers, increasing agent i's incentive to be treated herself.

As in any game with strategic complements, there is a maximal and a minimal Nash equilibrium (Milgrom and Roberts 1990). Rowthorn and Toxvaerd (2012) characterize these maximal and minimal equilibria in terms of two additional infection-prevalence thresholds, denoted here as $\underline{I}$ and $\overline{I}$, with $0 \leq \underline{I} \leq \overline{I} \leq \widehat{I}$.

In the maximal equilibrium, all infections are treated at time t if $I(t) < \overline{I}$ but none are treated if $I(t) > \overline{I}$. Similarly, in the minimal equilibrium, all infections are treated at time t if $I(t) < \underline{I}$ but none are treated if $I(t) > \underline{I}$. Note that the disease is eradicated in *all* equilibria if its initial prevalence $I(0)$ is less than $\underline{I}$, and is eradicated in *some* equilibrium if $I(0) \leq \overline{I}$.

A necessary condition for equilibrium eradication. Suppose that agent i is infected but that the disease is exceedingly rare, so that agent i faces negligible re-infection risk. When deciding whether to seek out treatment, agent i will compare the expected present value of the costs associated with her current infection, with and without treatment.

When untreated, agent i incurs flow cost s until recovery, which occurs at rate γ_0. When treated, agent i incurs flow cost $s + c$ until recovery at rate γ_A. Let C_0 and C_A be the expected present value of the costs incurred during an untreated and treated infected, respectively:

$$C_0 = \int_0^\infty s e^{-rt} e^{-\gamma_0 t} dt = \frac{s}{r + \gamma_0} \tag{15}$$

$$C_A = \int_0^\infty (s + c) e^{-rt} e^{-\gamma_A t} dt = \frac{s + c}{r + \gamma_A} \tag{16}$$

If $\gamma_A \leq \frac{s+c}{s}\gamma_0 + \frac{c}{s}r$, then $C_0 \leq C_A$ and agent i is strictly better off being left untreated. In this case, the disease cannot be eradicated in any equilibrium. Otherwise, if $\gamma_A > \frac{s+c}{s}\gamma_0 + \frac{c}{s}r$, then $C_0 > C_A$ and agent i is strictly better off being treated so long as the infection is sufficiently rare. In this case, the disease will be eradicated so long as the treatment becomes available sufficiently early, while infection remains sufficiently rare.

4.2 Restoring antibiotic effectiveness

> "Some experts say we are moving back to the pre-antibiotic era. No. This will be a post-antibiotic era ... an end to modern medicine as we know it."
>
> – Margaret Chan, Director-General of the World Health Organization, 2012

Staphylococcus aureus ("staph"), a bacterium that commonly colonizes the nasal passage, has numerous strains that dwell peacefully within the human microbiome. But some strains long ago acquired the genetic machinery to produce toxins that cause disease and, more recently, to survive exposure to the antibiotics used to treat bacterial infection ("antibiotic resistance"; see Laxminarayan et al. 2013). These staph strains are in a competition[24] that naturally favors the antibiotic-resistant strains, since they are more likely to survive when a person receives antibiotic treatment. For instance, methicillin-resistant staph (MRSA) first emerged in the 1960s but by 2014 accounted for over half of hospital-associated staph infections in the United States (WHO 2014, pg. 118).

[24]Dall'Antonia et al. (2005) found that colonization by a methicillin-sensitive strain reduced a person's likelihood of subsequently being colonized by a MRSA strain by 78%. Yang et al. (2018) found in mice that a commensal strain suppressed the growth of a MRSA strain and elicited both innate and adaptive immunity against MRSA skin infection.

Since antibiotic exposure is more effective at killing antibiotic-sensitive bacteria, increased *indiscriminate* antibiotic use advantages resistant strains and hence hastens the rise of untreatable infection. In that context, antibiotics are "exhaustible resources" whose value is diminished by use (Laxminarayan and Brown 2001). As the CDC and the Review on Antimicrobial Resistance (AMR Review) explained: "Because antibiotic resistance occurs as part of a natural evolution process, it can be significantly slowed but not stopped" (CDC 2013); and "Any use of antimicrobials, however appropriate and conservative, contributes to the development of resistance" (AMR Review 2014).

While indiscriminate antibiotic use promotes resistance, increased *targeted* use of any given antibiotic can slow or even reverse the rise of resistance to *other* antibiotics. This point was first made in McAdams (2014) and subsequently elaborated in McAdams (2017) for an obligate pathogen (SIS model) and in McAdams et al. (2019) for an opportunistic pathogen (SCIS model) with incidental antibiotic exposure in carriage. With targeted interventions, antibiotics become *preservable resources* whose value can be maintained even as all patients receive the best-available treatment.[25]

SIS model with competing strains. Consider an SIS model with multiple strains having different antibiotic-resistance profiles. In particular, suppose that there are two antibiotics, drug A and drug B, and four strains in circulation: "strain 0," sensitive to both drugs; "strain A," resistant to drug A but sensitive to drug B; "strain B," resistant to drug B but sensitive to drug A; and "strain AB," resistant to both drugs. For each resistance profile $X \in \{0, A, B, AB\}$, let $I_X(t)$ be the mass of hosts with strain-X infection at time t; so, $I(t) = \sum_X I_X(t)$, and $S(t) + I(t) = 1$, where $S(t)$ is the mass of uninfected hosts.

Let β_X be the transmission rate of strain X. Being resistant to an antibiotic can sometimes disadvantage bacteria in other ways. Such "fitness costs" can be captured by assuming that $\beta_0 > \max\{\beta_A, \beta_B\}$ and $\min\{\beta_A, \beta_B\} > \beta_{AB}$. When there are fitness costs, *withholding treatment* is enough to put resistant bacteria at disadvantage and cause their (relative and eventually absolute) number to dwindle over time. I focus here on the more challenging case without fitness costs, i.e., $\beta_X = \beta > 0$ for all X.

Treatment is assumed to have two sorts of effects: (i) *speedier recovery*, from baseline

[25]In their final report, citing an early version of McAdams et al. (2019), the AMR Review acknowledged the game-changing potential of rapid diagnostics to reverse the rise of resistance: "The information garnered from rapid diagnostics might eventually allow doctors to improve treatment and infection control to such an extent that this places negative selective pressure on resistance pathogens, thus reducing resistance in older drugs" (AMR Review 2016, pg. 35).

rate $\gamma_0 > 0$ to $\gamma_A > \gamma_0$ for A-sensitive infections treated with drug A, or $\gamma_B > \gamma_0$ for B-sensitive infections treated with drug B; and (ii) *resistance emergence*,[26] with sensitive infections becoming resistant to the drug being used for treatment at rate $\eta \geq 0$.

Case #1: indiscriminate treatment. Suppose for a moment that all patients are treated with drug A. The resulting flows among epidemiological states are as follows. *New infection ($S \to I_X$)*: Each X-infected host meets a susceptible host at rate $\beta S(t)$, creating flow $\beta S(t) I_X(t)$ from the susceptible state S to the X-infected state I_X. *Treated recovery ($I_0, I_B \to S$)*: 0- and B-infections clear at treated rate γ_A, creating flows $\gamma_A I_0(t)$ from I_0 to S and $\gamma_A I_B(t)$ from I_B to S. *Untreated recovery ($I_A, I_{AB} \to S$)*: A- and AB-infections clear at untreated rate γ_0, creating flows $\gamma_0 I_A(t)$ from I_A to S and $\gamma_0 I_{AB}(t)$ from I_{AB} to S. *Emergence of drug-A resistance ($I_0 \to I_A$, $I_B \to I_{AB}$)*: 0- and B-infections acquire A-resistance at rate η, creating flows $\eta I_0(t)$ from I_0 to I_A and $\eta I_B(t)$ from I_B to I_{AB}. The resulting epidemiological dynamics, expressed as percentage rates of change, are given by the following system of equations:

$$\frac{I_0'(t)}{I_0(t)} = \frac{I_B'(t)}{I_B(t)} = \beta S(t) - \gamma_A - \eta \tag{17}$$

$$\frac{I_A'(t)}{I_A(t)} = \beta S(t) - \gamma_0 + \eta \frac{I_0(t)}{I_A(t)} \tag{18}$$

$$\frac{I_{AB}'(t)}{I_{AB}(t)} = \beta S(t) - \gamma_0 + \eta \frac{I_B(t)}{I_{AB}(t)} \tag{19}$$

Note that strains $0, B$ grow at a slower percentage rate than strains A, AB. The percentage of infections that can be effectively treated with drug A therefore falls over time until, eventually, all infections are resistant to drug A.

Case #2: targeted treatment and no control. A rapid resistance diagnostic enables doctors to identify each infection's resistance profile and prescribe the best-available antibiotic treatment. Suppose that doctors prefer to prescribe drug A when both are effective, perhaps because drug A is less expensive or induces milder side effects. 0- and B-infections will continue to be treated with drug A, but now A-infections will be treated with drug B while AB-infections will be left untreated (to avoid harmful side

[26] "Resistance emergence" is shorthand for all the various ecological pathways by which a host who is initially infected (only or primarily) by sensitive bacteria can transition to being infected by resistant bacteria. Notably: (i) *treatment-induced mutation*, whereby antibiotic exposure triggers accelerated mutation (via stress response, chemical signaling by other bacteria, etc.), increasing the likelihood of a resistance-conferring mutation; and (ii) *competitive release*, whereby killing off sensitive bacteria promotes explosive growth of resistant bacteria that were already colonizing the host in small numbers.

effects). This leads to quicker clearance and hence less transmission of A-infection, but also faster emergence of multidrug resistance as AB-infections now emerge both from A-infections being treated with drug B and B-infections being treated with drug A. In particular, equations (18,19) become

$$\frac{I'_A(t)}{I_A(t)} = \beta S(t) - \gamma_B - \eta + \eta \frac{I_0(t)}{I_A(t)} \tag{20}$$

$$\frac{I'_{AB}(t)}{I_{AB}(t)} = \beta S(t) - \gamma_0 + \eta \frac{I_A(t) + I_B(t)}{I_{AB}(t)} \tag{21}$$

If both drugs are equally effective ($\gamma_B = \gamma_A$), then $\frac{I'_{AB}(t)}{I_{AB}(t)} > \frac{I'_A(t)}{I_A(t)} > \frac{I'_B(t)}{I_B(t)} > \frac{I'_0(t)}{I_0(t)}$. Strains $0, B$ dwindle over time the fastest, and strain A is also eventually overwhelmed by strain AB—a "post-antibiotic world" dominated by untreatable infection.

Case #3: targeted treatment and targeted control. While targeted treatment alone is insufficient to prevent a post-antibiotic world, targeted treatment combined with targeted infection control can be enough to put resistant strains at a disadvantage. "Targeted infection control" can take many forms depending on the pathogen and relevant host population, e.g., requiring a child with resistant pneumoccal infection to stay home from school (Ekdahl et al. 1998), providing skilled wound care (Solberg 2000) or free needles (Bassetti and Battegay 2004) for those with resistant skin staph infection, etc.

For simplicity, I assume here that infection control takes the form of *perfect isolation* of up to mass $\Delta > 0$ of hosts, with prioritized isolation first of those with untreatable AB-infection, then the B-infected, and then the A-infected, but no isolation of the 0-infected or of the uninfected. Moreover, I will assume that (i) isolation capacity is very limited, i.e., $\Delta \approx 0$, (ii) untreatable AB-infections are sufficiently rare *at first* to all be isolated, i.e., $I_{AB}(0) < \Delta$, and (iii) resistance emergence is sufficiently rare[27] that $\eta < \Delta\gamma_0$ and $\eta < \Delta\frac{\gamma_0}{\beta}$.

Under these assumptions, all AB-infections can be isolated in perpetuity. Targeted isolation reduces strain AB's *average transmission rate* from β to $\widehat{\beta}_{AB}(t) \equiv \beta \max\left\{\frac{I_{AB}(t)-\Delta}{I_{AB}(t)}, 0\right\}$. If AB-infection were already sufficiently widespread at time 0 that $\frac{I_{AB}(0)}{\Delta} \gg 1$, then $\widehat{\beta}_{AB}(t) \approx \beta$ and targeted isolation would have a negligible effect

[27]Resistance-conferring mutation is rare in bacteria and, so long as resistance to an antibiotic remains rare, competitive release will also tend to be rare as few hosts are colonized by even small numbers of the resistant strain. The assumption here that $\eta \approx 0$ therefore appears reasonable while resistance remains rare, but perhaps not once resistance has become common.

on epidemiological dynamics. However, because $I_{AB}(0) < \Delta$, all untreatable infections can be isolated and $\widehat{\beta}_{AB}(0) = 0$. Equation (21) at time 0 then becomes

$$\frac{I'_{AB}(0)}{I_{AB}(0)} = -\gamma_0 + \eta \frac{I_A(0) + I_B(0)}{I_{AB}(0)} \tag{22}$$

Note that $I'_{AB}(0) < 0$ if and only if $I_{AB}(0) < \widehat{I}_{AB}(0) \equiv \frac{\eta}{\gamma_0}(I_A(0) + I_B(0))$. Since $\eta < \gamma_0 \Delta$ (by assumption) and $I_A(0) + I_B(0) < 1$ (obviously), we have $\widehat{I}_{AB}(0) < \Delta$. Consequently, *either* $I_{AB}(0) < \widehat{I}_{AB}(0)$ and $I'_{AB}(0) > 0$ *or* $I_{AB}(0) \in \left(\widehat{I}_{AB}(0), \Delta\right)$ and $I'_{AB}(0) < 0$. Either way, AB-infection prevalence remains strictly less than Δ, allowing all AB-infections to continue to be isolated. Because this logic continues to apply after time 0, all AB-infections can be isolated in perpetuity. Moreover, $\lim_{t\to\infty} \frac{I_{AB}(t)}{I_A(t)+I_B(t)} = \frac{\eta}{\gamma_0} < \Delta \approx 0$ and $\lim_{t\to\infty} I_{AB}(t) \leq \Delta(1 - \gamma_0/\beta)$ (straightforward details omitted), implying that at least $\frac{\Delta\gamma_0}{\beta}$ isolation capacity is available to target other infections.

Since B-infections are equally-well treated as 0-infections (with drug A) and some B-infections are isolated, $\frac{I'_B(t)}{I_B(t)} < \frac{I'_0(t)}{I_0(t)}$. The prevalence of B-infection must therefore eventually fall to zero. What about A-infections? Let $\Delta_A(t) = \max\{\Delta - I_{AB}(t) - I_B(t), 0\}$ denote the isolation capacity available at time t to isolate A-infected hosts. Isolation reduces the flow of new A-infection from $\beta S(t) I_A(t)$ to $\max\{\beta S(t)(I_A(t) - \Delta_A(t)), 0\}$; modifying equation (20) and comparing to (17) yields

$$\frac{I'_0(t)}{I_0(t)} - \frac{I'_A(t)}{I_A(t)} = (\gamma_B - \gamma_A) + \beta S(t) \frac{\Delta_A(t) - \eta I_0(t)}{I_A(t)} \tag{23}$$

Consider the case in which both drugs are equally effective, so that $\gamma_B - \gamma_A = 0$.[28] As discussed previously, $\lim_{t\to\infty} \Delta_A(t) \geq \frac{\Delta\gamma_0}{\beta}$. Since $\eta < \frac{\Delta\gamma_0}{\beta}$ (by assumption) and $I_0(t) < 1$ (obviously), we have $\Delta_A(t) - \eta I_0(t) > 0$ and hence $\frac{I'_0(t)}{I_0(t)} > \frac{I'_A(t)}{I_A(t)}$ for all large t. The prevalence of A-infection must therefore also eventually fall to zero—and, with A-infection and B-infection each vanishing, AB-infection must also vanish.

In the end, the effectiveness of both antibiotics is completely restored, even as all patients receive the best-available antibiotic treatment—a complete *turning of the tables* on resistant bacteria!

What made this possible? First, doctors and public-health officials had access to a rapid resistance diagnostic, to know what treatment to prescribe and where specifically

[28] If $\gamma_A > \gamma_B$, then strain A may grow in number and eventually dominate the bacterial population. However, drug B remains effective in this case and can be used to treat A-resistant infections.

to deploy public-health resources. Second, a highly-effective infection-control option ("isolation") is available that specifically disrupts transmission from targeted infections. Finally, the pathogen in question spreads according to an SIS model, i.e., it is an *obligate pathogen* that causes disease immediately after colonizing the host.

Infection-oriented interventions such as isolation are much less effective against *opportunistic pathogens* that dwell in carriage for extended periods, such as *Streptococcus pneumoniae* (pneumococcus) and *Escherichia coli*. However, other microbiome-oriented interventions can potentially reverse the rise of resistance among such pathogens. For instance, McAdams et al. (2019) argues that (i) an annual "microbiome checkup" to detect resistant bacteria currently colonizing a patient plus (ii) a moderately effective intervention aimed at clearing these bacteria from carriage (e.g., seeding or promoting the growth of competitor bacteria) may be enough to select against resistance.

References

D. Acemoglu, V. Chernozhukov, I. Werning, and M. D. Whinston. A multi-risk SIR model with optimally targeted lockdown. NBER Working Paper 27102, 2020a.

D. Acemoglu, A. Makhdoumi, A. Malekian, and A. Ozdaglar. Testing, voluntary social distancing and the spread of an infection. NBER Working Paper 27483, 2020b.

J. Adda. Economic activity and the spread of viral diseases: Evidence from high frequency data. *Quarterly Journal of Economics*, 131(2):891–941, 2016.

F. E. Alvarez, D. Argente, and F. Lippi. A simple planning problem for Covid-19 lockdown. NBER Working Paper 26981, 2020.

S. T. Anderson, R. Laxminarayan, and S. W. Salant. Diversify or focus? spending to combat infectious diseases when budgets are tight. *Journal of health economics*, 31 (4):658–675, 2012.

M. C. Auld. Choices, beliefs, and infectious disease dynamics. *Journal of health economics*, 22(3):361–377, 2003.

C. Avery, W. Bossert, A. Clark, G. Ellison, and S. F. Ellison. Policy implications of models of the spread of coronavirus: Perspectives and opportunities for economists. NBER Working Paper 27007, 2020.

A. V. Banerjee. A simple model of herd behavior. *Quarterly Journal of Economics*, 107(3):797–817, 1992.

A. V. Banerjee. The economics of rumours. *Review of Economic Studies*, 60(2):309–327, 1993.

S. Bassetti and M. Battegay. *Staphylococcus aureus* infections in injection drug users: risk factors and prevention strategies. *Infection*, 32(3):163–169, 2004.

C. T. Bauch and D. J. Earn. Vaccination and the theory of games. *Proceedings of the National Academy of Sciences*, 101(36):13391–13394, 2004.

C. T. Bauch, A. P. Galvani, and D. J. Earn. Group interest versus self-interest in smallpox vaccination policy. *Proceedings of the National Academy of Sciences*, 100 (18):10564–10567, 2003.

Z. A. Bethune and A. Korinek. Covid-19 infection externalities: Trading off lives vs. livelihoods. NBER Working Paper 27009, 2020.

S. Bikhchandani, D. Hirshleifer, and I. Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5): 992–1026, 1992.

M. C. Bootsma and N. M. Ferguson. The effect of public health measures on the 1918 influenza pandemic in us cities. *Proceedings of the National Academy of Sciences*, 104(18):7588–7593, 2007.

F. Brauer, C. Castillo-Chavez, and C. Castillo-Chavez. *Mathematical models in population biology and epidemiology*, volume 2. Springer, 2012.

L. Brotherhood, P. Kircher, C. Santos, and M. Tertilt. An economic model of the Covid-19 epidemic: The importance of testing and age-specific policies. CESifo Working Paper, 2020.

Centers for Disease Control and Prevention. Antibiotic resistance threats in the United States, 2013. http://www.cdc.gov/drugresistance/pdf/ar-threats-2013-508.pdf, April 2013.

T. Y. Chan, B. H. Hamilton, and N. W. Papageorge. Health, risky behaviour and the value of medical innovation for infectious disease. *Review of Economic Studies*, 83 (4):1465–1510, 2016.

F. Chen. Rational behavioral response and the transmission of STDs. *Theoretical Population Biology*, 66(4):307–316, 2004.

F. Chen. A susceptible-infected epidemic model with voluntary vaccinations. *Journal of Mathematical Biology*, 53(2):253–272, 2006.

F. Chen. A mathematical analysis of public avoidance behavior during epidemics using game theory. *Journal of Theoretical Biology*, 302:18–28, 2012.

F. Chen and F. Toxvaerd. The economics of vaccination. *Journal of Theoretical Biology*, 363:105–117, 2014.

I. Chotiner. Paul Romer's case for nationwide coronavirus testing. *New Yorker*, May 2013.

J. C. Cochrane. An SIR model with behavior. https://johnhcochrane.blogspot.com/2020/05/an-sir-model-with-behavior.html, May 2020.

R. Cressman, V. Křivan, and J. Garay. Ideal free distributions, evolutionary games, and population dynamics in multiple-species environments. *The American Naturalist*, 164 (4):473–489, 2004.

M. Dall'Antonia, P. Coen, M. Wilks, A. Whiley, and M. Millar. Competition between methicillin-sensitive and-resistant *Staphylococcus aureus* in the anterior nares. *Journal of Hospital Infection*, 61(1):62–67, 2005.

R. Deb, M. Pai, A. Vohra, and R. Vohra. Testing alone is insufficient. Working paper (Toronto), 2020.

F. DeBrabander. The great irony of America's armed anti-lockdown protesters. *The Atlantic*, page May 13, 2020.

S. Del Valle, H. Hethcote, J. M. Hyman, and C. Castillo-Chavez. Effects of behavioral changes in a smallpox attack model. *Mathematical Biosciences*, 195(2):228–251, 2005.

A. Dixit. R0 for Covid research: An early estimate and policy implications. Working paper (Princeton), 2020.

K. Ekdahl, H. B. Hansson, S. Mölstad, M. Sóderström, M. Walder, and K. Persson. Limiting the spread of penicillin-resistant *Streptococcus pneumoniae*: experiences from the South Swedish Pneumococcal Intervention Project. *Microbial Drug Resistance*, 4(2):99–105, 1998.

H. Fang, L. Wang, and Y. Yang. Human mobility restrictions and the spread of the novel coronavirus (2019-nCoV) in China. NBER Working Paper 26906, 2020.

M. Farboodi, G. Jarosch, and R. Shimer. Internal and external effects of social distancing in a pandemic. NBER Working Paper 27059, 2020.

P. Fine and J. Clarkson. Individual versus public priorities in the determination of optimal vaccination policies. *American Journal of Epidemiology*, 124(6):1012–1020, 1986.

P.-Y. Geoffard and T. Philipson. Rational epidemics and their public control. *International Economic Review*, pages 603–624, 1996.

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

M. Gersovitz. The economics of infection control. *Annual Review Resource Economics*, 3(1):277–296, 2011.

E. L. Glaeser, C. S. Gorback, and S. J. Redding. How much does covid-19 increase with mobility? evidence from new york and four other us cities. NBER Working Paper 27519, 2020.

J. Greenwood, P. Kircher, C. Santos, and M. Tertilt. An equilibrium model of the African HIV/AIDS epidemic. *Econometrica*, 87(4):1081–1113, 2019.

M. J. Hendrix, C. Walde, K. Findley, and R. Trotman. Absence of apparent transmission of SARS-CoV-2 from two stylists after exposure at a hair salon with a universal face covering policy—Springfield, Missouri, may 2020. *Morbidity and Mortality Weekly Report*, 69, 2020.

H. Hethcote, M. Zhien, and L. Shengbing. Effects of quarantine in six endemic models for infectious diseases. *Mathematical Biosciences*, 180(1-2):141–160, 2002.

M. Hoy and M. K. Polborn. The value of technology improvements in games with externalities: A fresh look at offsetting behavior. *Journal of Public Economics*, 131: 12–20, 2015.

T. B. Hyde, M. Gilbert, S. B. Schwartz, E. R. Zell, J. P. Watt, W. L. Thacker, D. F. Talkington, and R. E. Besser. Azithromycin prophylaxis during a hospital outbreak of *Mycoplasma pneumoniae* pneumonia. *Journal of Infectious Diseases*, 183(6):907–912, 2001.

M. O. Jackson and D. López-Pintado. Diffusion and contagion in networks with heterogeneous agents and homophily. *Network Science*, 1(1):49–67, 2013.

M. O. Jackson, S. Malladi, and D. McAdams. Learning through the grapevine: the impact of message mutation, transmission failure, and deliberate bias. *arXiv preprint arXiv:1812.03354*, 2018.

G. Katriel and L. Stone. Attack rates of seasonal epidemics. *Mathematical Biosciences*, 235(1):56–65, 2012.

J. Keppo, M. Kudlyak, E. Quercioli, L. Smith, and A. Wilson. The behavioral SIR model, with applications to the swine flu and COVID-19 pandemics. Presentation available at https://www.lonessmith.com/wp-content/uploads/2020/06/BSIRslides-June.pdf, 2020.

W. O. Kermack and A. G. McKendrick. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London*, 115(772):700–721, 1927.

M. Kremer. Integrating behavioral choice into epidemiological models of AIDS. *Quarterly Journal of Economics*, 111(2):549–573, 1996.

D. Lakdawalla, N. Sood, and D. Goldman. Hiv breakthroughs and risky sexual behavior. *Quarterly Journal of Economics*, 121(3):1063–1102, 2006.

R. Laxminarayan and G. M. Brown. Economics of antibiotic resistance: A theory of optimal use. *Journal of Environmental Economics and Management*, 42:183–206, 2001.

R. Laxminarayan, A. Duse, C. Wattal, A. K. Zaidi, H. F. Wertheim, N. Sumpradit, E. Vlieghe, G. L. Hara, I. M. Gould, H. Goossens, et al. Antibiotic resistance—the need for global solutions. *Lancet Infectious Diseases*, 13(12):1057–1098, 2013.

P. Manfredi and A. D'Onofrio. *Modeling the interplay between human behavior and the spread of infectious diseases*. Springer Science & Business Media, 2013.

D. McAdams. *Game-Changer: Game Theory and the Art of Transforming Strategic Situations*. W.W. Norton, 2014.

D. McAdams. Resistance diagnosis and the changing epidemiology of antibiotic resistance. *Annals of the New York Academy of Sciences*, 1388(1):5–17, 2017.

D. McAdams. Nash SIR: An economic-epidemiological model of strategic behavior during a viral epidemic. *Covid Economics*, 16:115–134, 2020.

D. McAdams and T. Day. The political economy of a viral epidemic. Working paper (Duke), 2020.

D. McAdams and Y. Song. Viral social learning. Working paper (Duke), 2020a.

D. McAdams and Y. Song. Strategic social distancing with coordination motives. Working paper (Duke), 2020b.

D. McAdams, K. Wollein Waldetoft, C. Tedijanto, M. Lipsitch, and S. P. Brown. Resistance diagnostics as a public health tool to combat antibiotic resistance: A model-based evaluation. *PLoS Biology*, 17(5):e3000250, 2019.

L. A. Meyers, M. E. Newman, M. Martin, and S. Schrag. Applying network theory to epidemics: control measures for *Mycoplasma pneumoniae* outbreaks. *Emerging Infectious Diseases*, 9(2):204, 2003.

P. Milgrom and J. Roberts. Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica*, pages 1255–1277, 1990.

M. E. Newman. Spread of epidemic disease on networks. *Physical Review E*, 66(1): 016128, 2002.

A. Perisic and C. T. Bauch. Social contact networks and disease eradicability under voluntary vaccination. *PLoS Computational Biology*, 5(2):e1000280, 2009.

T. Philipson. Economic epidemiology and infectious diseases. *Handbook of Health Economics*, 1:1761–1799, 2000.

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

B. A. Prakash, D. Chakrabarti, N. C. Valler, M. Faloutsos, and C. Faloutsos. Threshold conditions for arbitrary cascade models on arbitrary networks. *Knowledge and Information Systems*, 33(3):549–575, 2012.

E. Quercioli and L. Smith. Contagious matching games. Technical report, Technical Report. Working paper, 2006.

T. C. Reluga. Game theory of social distancing in response to an epidemic. *PLoS Computational Biology*, 6(5), 2010.

T. C. Reluga. Equilibria of an epidemic game with piecewise linear social distancing cost. *Bulletin of Mathematical Biology*, 75(10):1961–1984, 2013.

Review on Antimicrobial Resistance. Antimicrobial resistance: Tackling a crisis for the future health and wealth of nations. http://amr-review.org/Publications, December 2014.

Review on Antimicrobial Resistance. Tackling drug-resistant infections globally: final report and recommendations. http://amr-review.org/Publications, May 2016.

R. Ross. An application of the theory of probabilities to the study of a priori pathometry—Part I. *Proceedings of the Royal Society of London*, 92(638):204–230, 1916.

R. Ross and H. P. Hudson. An application of the theory of probabilities to the study of a priori pathometry—Part II. *Proceedings of the Royal Society of London*, 93(650): 212–225, 1917.

R. Rowthorn and J. Maciejowski. A cost-benefit analysis of the covid-19 disease. *Oxford Review of Economic Policy*, 2020.

R. E. Rowthorn and F. Toxvaerd. The optimal control of infectious diseases via prevention and treatment. CEPR Discussion Paper No. DP8925, 2012.

R. E. Rowthorn, R. Laxminarayan, and C. A. Gilligan. Optimal control of epidemics in metapopulations. *Journal of the Royal Society Interface*, 6(41):1135–1144, 2009.

S. Roy and T. Sabarwal. Monotone comparative statics for games with strategic substitutes. *Journal of Mathematical Economics*, 46(5):793–806, 2010.

F. Salanié and N. Treich. Public and private incentives for self-protection. *The Geneva Risk and Insurance Review*, pages 1–10, 2020.

R. E. Serfling. Historical review of epidemic theory. *Human Biology*, 24(3):145–166, 1952.

C. O. Solberg. Spread of *Staphylococcus aureus* in hospitals: causes and prevention. *Scandinavian Journal of Infectious Diseases*, 32(6):587–595, 2000.

E. Talamas and R. Vohra. Go big or go home: a free and perfectly safe but only partially effective vaccine can make everyone worse off. PIER Working Paper 18-006, 2018.

F. Toxvaerd. Equilibrium social distancing. Cambridge-INET Working Paper 2020/08, 2020.

World Health Organization. Antimicrobial resistance: Global report on surveillance. 2014. https://www.who.int/antimicrobial-resistance/publications/surveillancereport/en/, 2014.

D. Yamin and A. Gavious. Incentives' effect in influenza vaccination policy. *Management Science*, 59(12):2667–2686, 2013.

J.-J. Yang, T.-W. Chang, Y. Jiang, H.-J. Kao, B.-H. Chiou, M.-S. Kao, and C.-M. Huang. Commensal *Staphylococcus aureus* provokes immunity to protect against skin infection of methicillin-resistant *Staphylococcus aureus*. *International Journal of Molecular Sciences*, 19(5):1290, 2018.