Review of Economic Analysis 14 (2022) 3–25 1973-3909/2022003

The Economic Consequences of $\hat{\mathcal{R}} = 1$: Towards a Workable Behavioural Epidemiological Model of Pandemics

JOSHUA S. GANS[*†]
Rotman School of Management, University of Toronto and NBER

This paper reviews the literature on incorporating behavioural elements into epidemiological models of pandemics. While modelling behaviour by forward-looking rational agents can provide some insight into the time paths of pandemics, the non-stationary nature of Susceptible-Infected-Removed (SIR) models of viral spread makes characterisation of resulting equilibria difficult. Here I posit a shortcut that can be deployed to allow for a tractable equilibrium model of pandemics with intuitive comparative statics and also a clear prediction that effective reproduction numbers (that is, $\mathcal{R}$) will tend towards 1 in equilibrium. This motivates taking $\hat{\mathcal{R}} = 1$ as an equilibrium starting point for analyses of pandemics with behavioural agents. The implications of this for the analysis of widespread testing, tracing, isolation and mask-use is discussed.

Keywords: epidemiology, basic reproduction number, SIR model, rational behaviour, pandemic policy

JEL Classifications: D01, I11, I18

1 Introduction

The workhorse model for the modelling of epidemics is the SIR (Susceptible-Infected-Removed) model of (Kermack and McKendrick 1927). It has been adopted to inform policy-makers in the management of the COVID-19 pandemic. The model is mechanistic in that people in the model do not make decisions that are reactive to current and predicted prevalence of an infectious disease in the population. As a key parameter, the basic reproduction number, $\mathcal{R}_0$, (a measure of the expected number of infections generated by a single infected person) is driven by people's choices regarding physical interactions. For this reason, the lack of behavioural elements has been a persistent source of criticism of such models.

This paper argues that, while a full behavioural model of pandemics is difficult to analyse as there is an element of non-stationarity in dynamic outcomes, there is value to be gained by analysing models that generate predictions that, for considerable lengths of time, the equilibrium reproduction number, $\hat{\mathcal{R}}$ is equal to 1 implying that the prevalence of an infectious dis-

[*]Thanks to Barry Nalebuff and Stephen King for helpful discussions

[†]Corresponding author: joshua.gans@utoronto.ca

© 2022 Joshua S. Gans. Licenced under the Creative Commons Attribution-Noncommercial 4.0 Licence (http://creativecommons.org/licenses/by-nc/4.0/). Available at http://rofea.org.

ease/virus is constant over time with the number of those newly infected approximately equaling the number of those newly recovered in a given time period. For COVID-19, such outcomes have been observed empirically beyond the initial stages of outbreaks across many regions (see Figure 1; (Atkeson 2021) and (Atkeson, Kopecky and Zha 2021)).[1]

Models that can generate an $\hat{\mathcal{R}} = 1$ equilibrium exist in the literature. For the SIR model whereby infectious individuals who recover are removed from the susceptible pool, I show that an $\hat{\mathcal{R}} = 1$ outcome requires a special set of assumptions that are unlikely to generally hold. This is because individuals may base their behaviour on prevalence (i.e., the number of infected people they are likely to encounter) rather than on the ever falling set of susceptibles. That set, however, does impact on the reproduction number. Nonetheless, for the SIS model, whereby infectious individuals who recover remain susceptible to future infections, the $\hat{\mathcal{R}} = 1$ outcome is a natural equilibrium. This suggests that, when prevalence is relatively low, even for the SIR model, the number of susceptibles will not change at a rapid pace and thus, an $\hat{\mathcal{R}} = 1$ outcome provides an approximate outcome that may explain observed behaviour.

In what follows, I first present the standard (non-behavioural) SIR model. I then review various behavioural models that have been utilised in the literature deriving. I provide a graphical approach to describe the resulting equilibrium outcomes. A final section offers some predictions from this approach.

2 The Standard SIR Model

Let $\{S(t), I(t), R(t)\}$ denote the shares (and levels) of the population (normalised to be of size 1 over a continuum of agents) who are either susceptible to the virus, infected with the virus or removed (i.e., recovered or dead) from the virus at time $t \geq 0$. It is assumed that time is discrete. In the SIR model, these variables are assumed to evolve according to the following dynamic equations:

$$S(t+1) - S(t) = -\beta S(t) I(t)$$

$$I(t+1) - I(t) = (\beta S(t) - \gamma) I(t)$$

$$R(t+1) - R(t) = \gamma I(t)$$

Here γ is the probability that an infected person will be removed in any given period while β is the probability that a susceptible person will become infected by an infected person in a given period. Observe that the number of infections in the population will be falling (i.e., $I(t+1) < I(t)$) if $\frac{\beta}{\gamma} S(t) < 1$ and will be rising (i.e., $I(t+1) > I(t)$) if $\frac{\beta}{\gamma} S(t) > 1$. The LHS of these inequalities is the effective reproduction number, $\mathcal{R}_t$. Since $S(0) \approx 1$, then $\mathcal{R}_0 = \frac{\beta}{\gamma}$. $\mathcal{R}_0$

[1] Moreover, there is plenty of evidence that people act to mitigate their own infection risk apart from those mandated by governments. See (Farboodi, Jarosch and Shimer 2021), (Goolsbee and Syverson 2021) and (Allcott, Boxell, Conway, Ferguson, Gentzkow and Goldman 2020)).

Figure 1: Estimated $\mathcal{R}_t$ for US States (COVID-19) from epiforecasts.io - part I

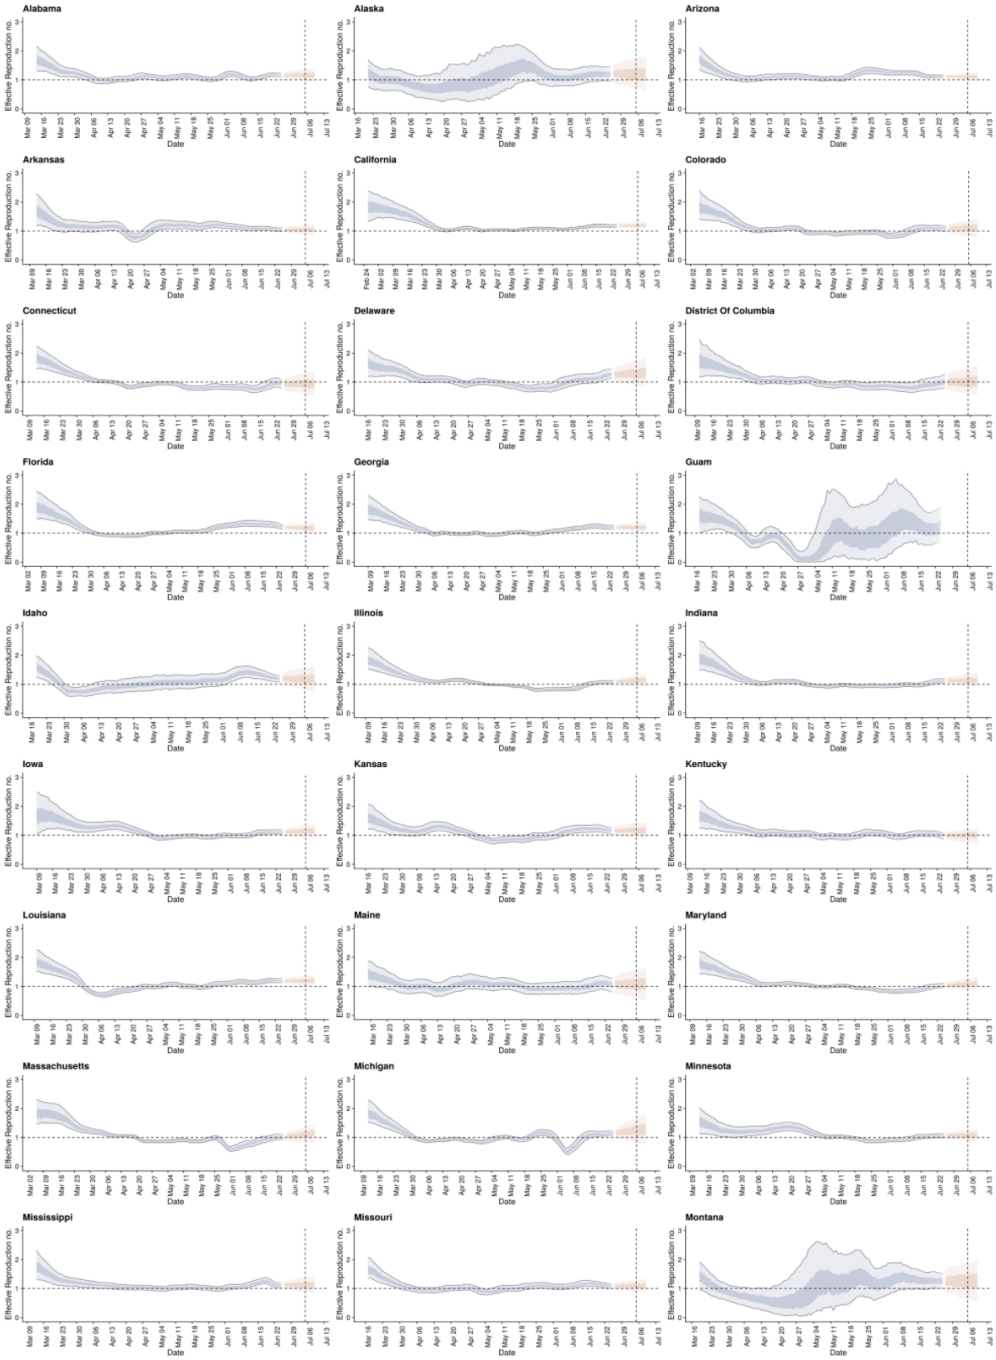

Figure 2: Estimated $\mathcal{R}_t$ for US States (COVID-19) from epiforecasts.io - part II

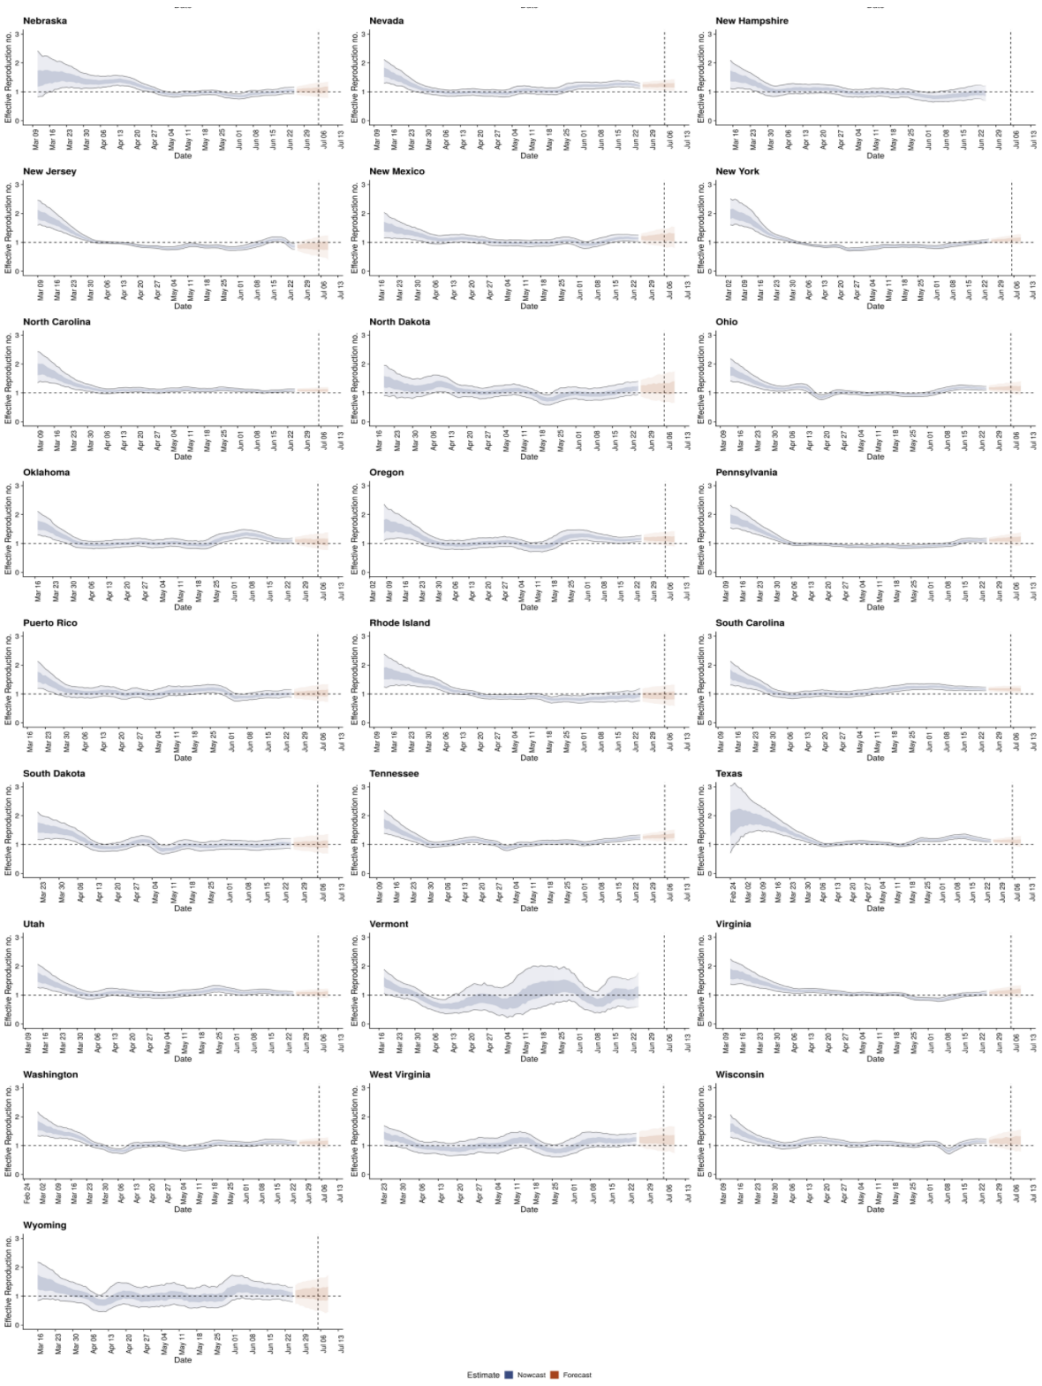

is the basic reproduction number which has the interpretation as the total expected number of infections one infectious person will create over the life of their infection.

A few remarks about this model. First, there are two relevant state variables $\{I(t), S(t)\}$ and they co-evolve according to:

$$I(t) = 1 - S(t) + \frac{1}{\mathcal{R}_0} \log(S(t))$$

.

where it is assumed that $\{I(0), S(0)\} = \{0, 1\}$. Second, the share of the population that is eventually infected, $i \equiv 1 - S(\infty)$, is given by:

$$\mathcal{R}_0 = -\frac{\log(1-i)}{i}$$

.

Third, temporary changes to β can influence the eventual share of infected people, i, although regardless $i \geq \frac{1}{\mathcal{R}_0}$; the 'herd immunity' threshold. Fourth, the peak prevalence arises when $\bar{S} = \frac{1}{\mathcal{R}_0}$ and involves, at that point:

$$\bar{I} \equiv 1 - \frac{1 + \log(\mathcal{R}_0)}{\mathcal{R}_0}$$

This all implies that, for $\mathcal{R}_0 > 1$, (a) an equilibrium with $S = 1$ is locally unstable and (b) with temporary interventions that decrease β or increase γ, the absorbing states for i are characterised by $S(\infty) = [\underline{S}, \frac{1}{\mathcal{R}_0}]$ where $\underline{S}$ is defined by $\mathcal{R}_0 = -\frac{\log(\underline{S})}{1-\underline{S}}$; that is, either infections are kept at zero or they evolve to a point beyond the 'herd immunity' threshold.[2]

The standard SIR model is useful in that it relates the evolution of a pandemic according to $\mathcal{R}_0$ and how the underlying parameters associated with it may be impacted upon over the life of the pandemic. This can be useful for analysing the impact of non-pharmaceutical interventions that impact those underlying parameters. However, if, as is likely, those underlying parameters are not fixed but vary according in ways that relate to the underlying state variables, the standard SIR model will face challenges in being of predictive value.

3 The Behavioural SIR Model

The fact that the standard SIR model lacked behavioural elements has not been lost on epidemiologists. In particular, it has been recognised that people might observe current prevalence (that is, $I(t)$) and modify their own behaviour so as to reduce infection risk. However, the mathematical epidemiologists have typically taken what economists would call a 'reduced-form' approach to this. For instance, they might posit a variable, $x \in [0, 1]$, that is a filter reducing the impact

[2] See (Rachel 2020) for details.

of β on new infections. That variable is then assumed to be a decreasing function of $I(t)$; e.g., $x(I(t))$.[3] A similar approach was used by (Cochrane 2020).[4]

3.1 Literature Review

Work in economics to include behavioural elements in models of epidemics started in earnest with the study of the spread of AIDS. Following (Philipson and Posner 1993), (Geoffard and Philipson 1996) examined an SI model, whereby people can transition from susceptible to infected but cannot recover or become non-infectious, and examined the way in which increased prevalence would change the behaviour of a forward-looking rational agent. They showed that the incentives of infected agents – e.g., whether they altruistic or not – played an important role.[5] This line of research has continued with a mapping to empirical models by (Greenwood, Kircher, Santos and Tertilt 2019).

The pioneering treatment that first introduced forward-looking, rational economic agents into epidemiological models that could provide insights on COVID-19 was provided by (Gersovitz and Hammer 2004). They examined SIR (in addition to SIS and SID models) to explore the different effects that prevention versus a treatment might have on the dynamics of epidemics. In doing this, they were able to clarify the externalities that may be present and the efficacy of various forms of interventions (including taxes and subsidies) to improve social welfare.[6] This approach inspired other analyses developing variants of their behavioural model including (Reluga 2010) who showed that agents will socially distance more when $\mathcal{R}_0$ is high (as they fear becoming infected) and (Fenichel 2013) who showed that non-targeted lockdown policies may be worse than a decentralised behavioural outcome in terms of overall utility. There is a literature that has examined behavioural SIS models with rational agents including (Chen 2009) looks at how the provision of information impacts on agent's incentives to minimise risks of infection in an SIS model. When prevalence is low, agents may take more risks and make eradication impossible. (Toxvaerd 2019) uses an SIS model where agents bear costs of reducing interactions. Agents are forward looking and understand the SIS dynamics. He examines the impact on a treatment that reduces transmission rates on social welfare and finds potential for a welfare-reducing rebound effect. (Rowthorn and Toxvaerd 2020) examines the appropriate mix of prevention and treatment while (Goodkin-Gold, Kremer, Snyder and Williams 2020)

[3]See for example, (Eksin, Paarporn and Weitz 2019) who also explore assumptions where $x(I(t), R(t))$ is decreasing in both variables, that they argue is a model of 'long-term awareness' in contrast to 'short-term awareness' where x is a function of $I(t)$ alone.

[4]This might be termed an 'old-timey' macro approach.

[5](Kremer 1996) also included behavioural elements in an SI model but his focus was on equilibrium outcomes in a broader matching game. See also an early contribution by (Auld 2003).

[6](Chen 2012) uses an SIR model where agents can reduce their physical interactions and infected agents may be debilitated and so interact less. He focuses on myopic agents and analyses the impact of different matching functions on the resulting equilibrium outcomes.

looks at vaccine pricing where epidemiological effects are anticipated and influenced. There is also a literature that focuses on incentives to be vaccinated using behavioural foundations. (Francis 1997) uses an SIR model to consider an agent's choice of when to vaccine and finds that the market is efficient. (Gersovitz 2003) and (Chen and Toxvaerd 2014) relaxes those conditions and finds inefficiency especially if individuals can independently acquire immunity.[7]

A recent literature on COVID-19 has similarly built on these behavioural foundations with forward-looking rational agents including (Eichenbaum, Rebelo and Trabandt 2020), (Makris 2021) and (Boppart, Harmenberg, Hassler, Krusell and Olsson 2020) (and by extension (Krueger, Uhlig and Xie 2020)) who provide a model of endogenous social distancing in a macroeconomic model; (Farboodi et al. 2021) who examine how altruistic preferences (that capture the degree to which individuals choose to self-isolate if they know they are infected) impact on behaviour. (Jones, Philippon and Venkateswaran 2020) use a macro-model and highlight a 'fatalism' effect whereby, when prevalence is high, people do not socially distance as they are likely to become infected anyway; (Bethune and Korinek 2020) who look at what optimal policies look like when the planner has a high degree of information regarding who is infected and who has recovered. (Bisin and Moro 2020) examine behavioural elements combined with frictions in spatial diffusion something also done by (Aguirregabiria, Gu, Luo and Mira 2020) using a structural model; (McAdams 2020) provides a finite time model but focuses on the case where agent value from economic activity depends on the activity of others introducing a complementarity and the possibility of multiple equilibria; (Di Guilmi, Galanis and Baskozos 2020) who look at the impact of limited information provided to agents and (Brotherhood, Kircher, Santos and Tertilt 2020) do a variety of policy experiments.[8] The most careful analyses in this regard of the microeconomic foundations of the SIR model come from (Toxvaerd 2020), (Rachel 2020), (Dasaratha 2020) and (Fukuda, Kos and Wolf 2021) who provide analyses that show the conditions under which endogenous social distancing will be too little, and potentially, too much compared with what might be socially optimal. For a recent review see (McAdams 2021).

3.2 Model Setup

At the core of each of these models is a conception of a behavioural agent. An agent, n, chooses their level of activity, $x_n \in [0, 1]$, which can be interpreted as their risk of interacting with another agent or preventative measures (such as wearing a mask). That activity gives them value in utility terms of $u_n(x_n)$ in each period where $u_n(.)$ is increasing, concave and independent across time periods. Agents have a common discount factor of $\delta < 1$. If an agent becomes infected, they incur an additional loss, L, in utility unless they die in which case they can incur no utility thereafter. An infected agent has a probability, γ of becoming no longer infectious in each period they are infected. At that point, with probability ρ, they survive and become

[7] See (Philipson 2000) and (Gersovitz 2011) for reviews.

[8] See (Gans 2020b) for further discussion of this literature.

immune. Otherwise, they die. Either way they are part of R, the set of removed agents.

An agent's activity choices at t are determined by the condition, $\{S, I, R\}$, they are in at that time. If they are part of R and have not died, they are no longer infectious or at risk. Hence, they will set their activity, $x_{n,R} = 1$ and will earn an expected present discounted payoff of $\frac{u_n(1)}{1-\delta}$. In this, there is an implicit assumption that a recovery means a full recovery to the utility they would earn had the epidemic not emerged.

3.3 Infected Agent Activity

For an infected agent (a member of I), they are infectious and sick. Their instantaneous utility is $u_n(x_{n,I}) - L$ and their expected discounted payoff is:

$$V_{n,I}(t) = u_n(x_{n,I}(t)) - L + \delta(\gamma V_{n,R} + (1-\gamma)V_{n,I}(t+1))$$

where here $V_{n,R} = \rho\frac{u_n(1)}{1-\delta}$. Note that, being self-interested, infected agents set $x_{n,I}(t) = 1$ in each period and, thus, their expected discounted payoff becomes:

$$V_{n,I} = \frac{u_n(1) - L + \delta(1-\gamma)\rho\frac{u_n(1)}{1-\delta}}{1-\delta\gamma}$$

This captures, in a stark way, a key externality that arises for infectious diseases when an infected person does not perceive a personal risk from social interactions. Of course, various factors could alter this stark result including that infected people may not be capable of or desire the same level of activity if they were healthy and that such activity may not be as valuable because others may avoid them if they knew they were infectious. For COVID-19, this was complicated by the fact that many of the infected were asymptomatic or pre-symptomatic and did not know they were infectious. In this situation, an agent may act as if they were still susceptible.

3.4 Susceptible Agent Activity

For both the infected and recovered, their choice of economic activity is not impacted upon by the state variables, $\{I(t), S(t)\}$. Thus, the key to the behavioural approach to epidemiology are the choices of the susceptible. Their instantaneous utility is $u_n(x_(n, S)(t))$ and their expected discounted payoff is:

$$V_{n,S}(t) = u_n(x_{n,S}(t)) + \delta(p(x_{n,S}(t), I(t))V_{n,I}(t+1) + (1 - p(x_{n,S}(t), I(t)))V_{n,S}(t+1)$$

where $p(x_{n,S}(t), I(t))$ is probability that n becomes infected at time t (the consequences of which are felt at time $t+1$). $p(.)$ is generally increasing in both of its arguments; i.e., a higher rate of infection in the population as well as a higher rate of activity by n raises the probability that n becomes infected. If $V_{n,I}(t+1) < V_{n,S}(t+1)$ this is not something that n wants and, thus, the increased risk of becoming infected will constrain the agent's choice of activity.

The structure of $p(x_{n,S}(t), I(t))$ depends upon how activity translates into an individual's risk of infection. The standard SIR model assumes that susceptible individuals face a probability, β, of becoming infected if they interact with an infected individual. What an 'interaction' precisely is, however, is potentially rich. For instance, if an agent visits a location where a number of other people are present, then β would be interpreted as the probability that at least one those people are infected. If a virus lingers or is spread on surfaces, then the probability that an agent becomes infected relates to the number of infected people who may be at a place in the past.[9]

Typically, the standard epidemiological models consider simpler environments. The simplest case assumes that an individual agent encounters one other member of the population at random in each period. In this situation, $x_{n,S}(t)$, is interpreted as the probability that n is matched with another person in period t who is infected with probability $I(t)$. Thus, the probability that n becomes infected is:

$$p(x_{n,S}(t), I(t)) = x_{n,S}(t)\beta I(t)$$

Of course, it is possible to imagine a slightly richer model whereby a susceptible understands that β might differ between alternative activities or that they can choose different populations with different $I(t)$ probabilities to interact with.[10] This structure presumes that $x_{m,I}(t) = 1$ for infected agents, $m \in I(t)$. If, for reasons of altruism or regulation, $x_{m,I}(t) < 1$, then the probability that n encounters an infected agent is $\frac{1}{I(t)} \int_0^{I(t)} x_{m,I}(t)dm$ so that $p(x_{n,S}(t), I(t)) = x_{n,S}(t)\beta \int_0^{I(t)} x_{m,I}(t)dm$.

3.5 First-Order Effects

A susceptible individual, n, will choose $x_{n,S}(t)$ to maximise $V_{n,S}(t)$ holding the state variables and their future path as given. This gives rise to the marginal condition for the optimal choice $\hat{x}_{n,S}(t)$:

$$u_n'(\hat{x}_{n,S}(t)) = \beta I(t)\delta(V_{n,S}(t+1) - V_{n,I}) \tag{OPT}$$

This leads to a myriad of insights.

- (**Greater prevalence reduces susceptible activity**) Holding $V_{n,S}(t+1)$ fixed, as $I(t)$ increases $\hat{x}_{n,S}(t)$ falls. That is, the first-order effect of greater prevalence reduces an agent's activity as they forgo utility to reduce the risk of becoming infected.

[9](Acemoglu, Chernozhukov, Werning and Whinston 2022) explore these issues by considering a variety of matching functions between susceptible agents and infecteds in an SIR model.

[10]See (Ellison 2020) for a review of these richer environments.

- (**A more infectious virus reduces susceptible activity**) Holding $V_{n,S}(t+1)$ fixed, if the infectiousness of the virus (β) rises then $\hat{x}_{n,S}(t)$ falls. As will be noted below, this can reduce the rate of growth of the epidemic which stands in contrast to the clear prediction of the standard SIR model that a higher β will lead to faster epidemic spread and higher long term infections ((Toxvaerd 2020)).

- (**Greater activity from infecteds reduces susceptible activity**) If, for some infecteds, $x_{m,I}(t) < 1$, it can be seen that $\hat{x}_{n,S}(t)$ may be higher. Thus, there is a strategic substitute between the activity choices of infected agents and susceptible agents (as noted by (Keppo, Quercioli, Kudlyak, Wilson and Smith 2020)).

- (**Activity is slower to return to normal as pandemic eases**) The future path of the epidemic is captured in the term, $V_{n,S}(t+1) - V_{n,I}$. Note, in particular, that if $I(t+1) > I(t)$, then $\hat{x}_{n,S}(t) \leq \hat{x}_{n,S}(t-1)$ while the opposite is true if $I(t+1) < I(t)$. This, as (Rachel 2020) shows, implies that a susceptible agent is going to engage in a smaller reduction in activity at the beginning of an epidemic than at the end for the same level of prevalence.[11] That is, for, $\underline{T} < \bar{T}$ where $I(\underline{T}) > I(\underline{T}-1)$, $I(\bar{T}) < I(\bar{T}-1)$ and $I(\underline{T}) = I(\bar{T})$, $\hat{x}_{n,S}(\underline{T}) > \hat{x}_{n,S}(\bar{T})$. Individuals will be more cautious at the end of a pandemic as the relative on-going value of being susceptible is higher.

- (**Complementarity between activity of susceptibles**) The interaction between a susceptible agent's decision on their own activity and the activity of other susceptible agents is potentially subtle. As will be described below, if susceptible agents reduce their activity at t, then this will reduce the share of the population infected at $t+1$. For an individual agent, therefore, a reduction in expected activity by other susceptibles increases $V_{n,S}(t+1)$ and hence, decreases their own choice of activity at time t as there is a greater value to not being infected. Thus, for susceptibles, their activity are strategic complements while at the same time constituting a negative externality on one another.

- (**Prospects for a vaccine or treatment have opposite effects on susceptible activity**) If a vaccine is expected at a future time, this increases $V_{n,S}(t+1)$ and hence, causes susceptible agents to reduce their activity; becoming more cautious so as to obtain the vaccine and not become infected. By contrast, if a treatment is expected at a future time, this, by either increasing ρ or decreasing L, causes $V_{n,I}$ to be higher and, thus, susceptibles to be less cautious of becoming infected and so increase their activity.

These insights are all implications of the first-order effects of changes in the environment on the behaviour of susceptible individuals. However, the full equilibrium effects can be harder to derive.

[11]The notion that at the onset of a pandemic, agents who expect a higher growth in infections tend to increase their activity and risk of infection is called the *fatalism* effect by (Jones et al. 2020).

3.6 Equilibrium Analysis

To see this, we need to explore the evolution of the state variables under the behavioural assumptions that individual agents can influence their individual infection risk. Fortunately, the simple specification for $p(.)$ used above provides a natural way of aggregating into the expected path for the state variables, $\{I(t), S(t)\}$.

Let $X_S(t) \equiv \int_0^{S(t)} x_{n,S}(t)dn$. The expected number of new infecteds is equal to $\beta X_S(t)I(t)$ while each period $\gamma I(t)$ infecteds are removed. Thus,

$$I(t+1) - I(t) = (\beta X_S(t) - \gamma)I(t)$$

By construction, this also means the total number of susceptibles declines by:

$$S(t+1) - S(t) = -\beta X_S(t)I(t)$$

Note that if $x_{n,S}(t) = 1$ for all $n \in S(t)$, then $X_S(t) = S(t)$ and the above two equations become the same as the standard SIR model.

It can be seen here that the time path of $\{X_S(t), ...\}$ determines the net presented expected value of continuing to be susceptible and, thus, the incentives to undertake activity at time t. Thus, the equilibrium outcome would require solving for a multi-dimensional fixed even with commonly used simplifying assumptions such as all agents being symmetric in preferences. Moreover, the set of susceptibles is being reduced in size over time at a rate that is endogenous to the activity choices of susceptibles themselves. This means that there is unlikely to be stationary equilibrium outcome that we usually look for in order to conduct comparative statics. For this reason, most studies of behavioural SIR models have used simulations to demonstrate potential outcomes rather than analytical solutions. For this reason, I propose here, instead, taking a shortcut that will permit an analytical solution albiet at the expense of not (usually) satisfying our usual equilibrium requirements.[12]

4 An Analytical Shortcut

The analytical shortcut I propose here is to establish conditions under which $I(t+1) = I(t)$ for an interval of time. The condition is a simple one: $S(t+1) = S(t) = S$ for all t. It is immediately is apparent that this condition violates the laws of motion of the SIR model whenever $\gamma > 0$.

[12]This non-stationarity makes this analysis difficult and has been a challenge to forward-looking SIR models (see (McAdams 2021) for a discussion). (Rachel 2020) characterises this equilibrium and shows that infection rates are such that $\mathcal{R}_t$ is just below 1 and so fall over time. This implies that $\hat{x}_S(t)$ should also fall over time. However, (Rachel 2020) uses various approximations in order to analyse policies in the model that have the effect of making $\hat{x}_S(t)$ constant over time. An alternative would be not to use an SIR model but an SI model as in (Gersovitz 2003). However, as the present paper is motivated by the Covid-19 pandemic it is useful to keep closer to the epidemiology of the SIR model (and as will be demonstrated below, this does matter for the policy analysis that follows).

As an accounting measure, it simply cannot be the case that some infected individuals are recovered (or strictly speaking) removed and $S(t)$ is not falling over time. Of course, this state of affairs is possible for the SIS model which is perhaps why much of the initial work integrating behavioural assumptions into epidemiology examined that environment. However, because we want the incentives of agents to reflect the possibility that they can be removed following an infection, I cannot simply follow the SIS model here. Instead, I have just been inspired by it.[13]

4.1 Equilibrium Solution

The focus is on the equations governing the relationship between $X_S(t)$ and $I(t)$. The first equation is behavioural.

$$\hat{X}_{n,S}(I(t)) = \int_0^S \hat{x}_{n,S}(I(t))dn \tag{BEH}$$

This equation is how the aggregate activity of susceptible agents (now fixed at size S) is a function of $I(t)$ when individual agents are optimising. Note that $\hat{X}_{n,S}(I(t))$ is a non-decreasing function of $I(t)$ as discussed earlier.

The second equation comes from the SIR laws of motion.

$$I(t+1) = I(t) + (\beta X_S(t) - \gamma)I(t)$$

The number of infected agents is an increasing function of the aggregate activity, $X_S(t)$, of those agents.

Essentially, these two equations describe a dynamic aggregate game involving choices of susceptible agents but under an assumption that the set of those agents is now fixed. The goal will be to characterise stationary Markov perfect equilibria of this game using a dynamic programming approach.

From the law of motion, we have:

$$X_S(t) = \frac{\frac{I(t+1)-I(t)}{I(t)} + \gamma}{\beta} \tag{EPI}$$

Setting this equal to $\hat{X}_S(I(t))$ equilibria in which $I(t+1) = I(t)$ for all t can be explored. When this condition is satisfied then $\hat{x}_{n,S}(t+1) = \hat{x}_{n,S}(t)$ for all t which carries over to, $\hat{X}_{n,S}(I(t))$. Importantly, this means that:

$$I(t+1) - I(t) = 0 = (\beta\hat{X}_S(I(t)) - \gamma)I(t) \implies \hat{X}_S(I^*) = \frac{1}{\mathcal{R}_0} \tag{EQM}$$

[13]The idea is to focus on conditions under which $I(t+1) = I(t)$ for an interval of time (which is also what (Rachel 2020)'s approximations do). There are other ways to ensure stationarity. One path, discussed in depth by (McAdams 2021) is to not assume forward looking behaviour beyond an additional period. Another would be to introduce some randomness associated with the point at which herd immunity is reached – say, because of endemic factors in the disease – that will have a similar analytical effect.

Importantly, this implies that the equilibrium effective reproduction number,

$$\hat{\mathcal{R}} = \hat{X}_S(I^*)\mathcal{R}_0 = 1$$

Thus, prevalence will neither rise nor decline in equilibrium and this pins down that equilibrium steady state of infected agents.[14]

4.2 Graphical Analysis

The analytical shortcut has the advantage that it permits a (familiar to economists) graphical analysis. Figure 3 shows the EPI and BEH lines in (X_S, I) space. BEH shows how the aggregate choice of activity level is determined by the prevailing share of infected agents and, as shown, earlier is typically downward sloping as agent's reduce activity more when there is a greater chance of encountering an infected agent. EPI shows how the number of infected agents relates to the aggregate choice of activity level by susceptibles. It is upward sloping as a higher X_S directly increases $I(t+1)$ in a linear fashion in the SIR model. Where the two curves intersect is the equilibrium outcome under the assumption that S is held fixed.

This graphical approach also shows why the equilibrium is stable. Suppose that $I(t) < I^*$. Then $\hat{X}_S(t) > \frac{1}{S\mathcal{R}_0}$ and $I(t+1) > I(t)$. By contrast, if $I(t) > I^*$, $\hat{X}_S(t) < \frac{1}{S\mathcal{R}_0}$ and $I(t+1) < I(t)$. These processes only stop as $I(t) = I^*$.

This approach allows for intuitive comparative static analysis. Figure 4 shows what happens if there is an increase in baseline infectiousness, β. Firstly, as is well known in epidemiology, an increase in β means that more activity translates into higher infections at a faster rate; shifting the EPI line to the right. Second, from the analysis of behavioural responses, an increase in β causes susceptible agents to choose to be more careful and reduce their activity. Thus, the BEH curve shifts to the left. An increase in β has a negative equilibrium impact on aggregate activity from susceptibles but an ambiguous impact on the equilibrium number of infected agents. Figure 4 is drawn to show the case where an increase in β leads to a lower rate of infection in contrast to the standard epidemiological prediction. However, if the behavioural response is weaker, then the opposite comparative static is possible.

Interestingly, there are some unambiguous monotone comparative static results that can be derived. For instance, a change that impacts on $V_{n,I}$ only without changing anything else – e.g., a treatment that increases ρ or a measure that makes being infected less costly (i.e., reduces L) only impacts on BEH; shifting it up and to the right. Thus, the availability of a treatment causes both a higher equilibrium activity and a higher equilibrium level of infection.

[14] It can readily be seen that this equilibrium exists if $\mathcal{R}_0 > 1$. When $I(t) = 0$, all agents set $\hat{x}_{n,S} = 1$ so that $\hat{X}_S(0) = 1$. At this point $X_S(0) = \frac{1}{\mathcal{R}_0}$ which is less than 1. On the other hand, if $I(t) = 1$, $X_S(1) = \frac{\frac{1-I(t-1)}{I(t-1)}+\gamma}{\beta} > 0$ while $\hat{X}_S(I(t)) \to 0$. As all of the relevant functions are continuous, there is a fixed point where $I(t) = I^*$.

Figure 3: Equilibrium

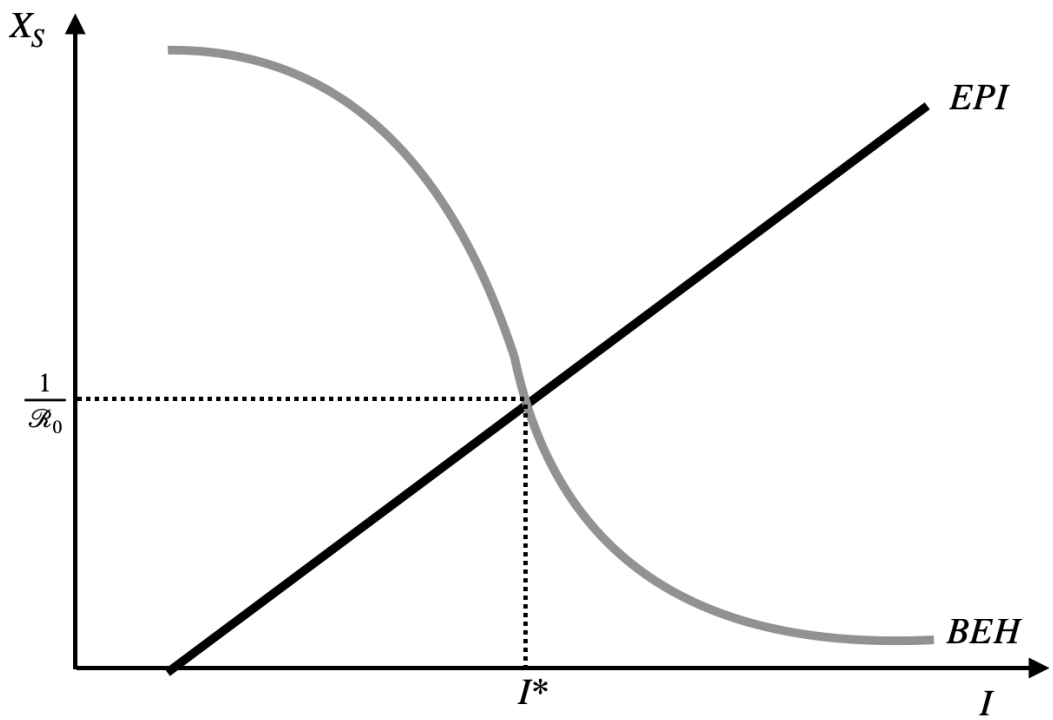

4.3 Impact of Testing/Isolation

One policy that has received attention in COVID-19 is the increased use of testing (and contact tracing) to identify infected individuals earlier and isolate them to prevent them spreading the virus. This approach was adopted as a standard practice by many countries and appeared to be successful in reducing the scale of the COVID-19 pandemic. (See (Gans 2020a), Chapter 7 for more details). However, some recent work in economics has raised the possibility of unintended behavioural consequences from increased testing including testing giving infected people confidence to engage in activity because they can't get more infected ((Taylor 2020) and (Deb, Pai, Vohra and Vohra 2020)), a reluctance to be tested for fear of being quarantined ((Eichenbaum et al. 2020) to the potential for a rebound effect that increases activity choices ((Acemoglu, Makhdoumi, Malekian and Ozdaglar 2020)). The model presented here permits the examination of these consequences.

The focus here is on the situation where tests immediately trigger isolation (say, because they are done by a public authority with enforcement power or have subsidies that induce iso-

Figure 4: Increase in Infectiousness

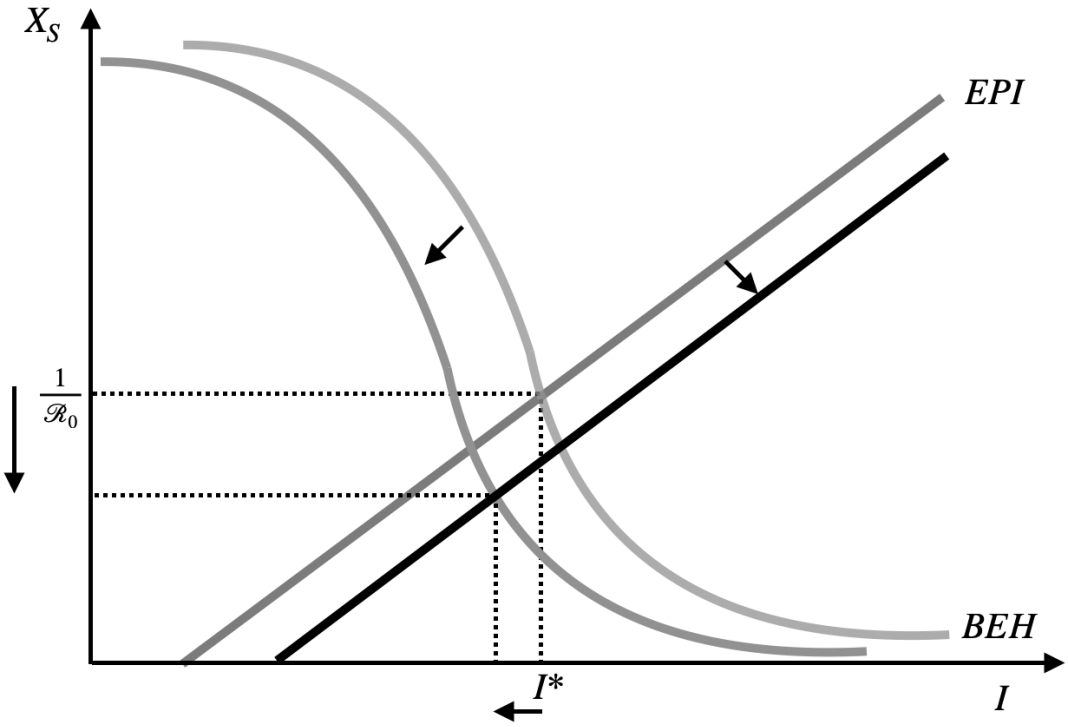

lation).[15] The first impact of testing (along with isolation) is one that is intended: it reduces the probability that a susceptible encounters an infected agent. This impacts on both the EPI and BEH equations. The epidemiological response is to shift the EPI curve to left as it directly reduces the probability that a susceptible agent will encounter and infected agent. However, this also leads to a shift outwards of the BEH curve. The reduction in the probability of encountering an infected agent, increases the incentives of susceptible agents to engage in activity for given level of prevalence. This is the effect identified by (Acemoglu et al. 2020) and (Hellmann and Thiele 2020).[16] Thus, examining this impact alone, we would find a similar ambiguous comparative static as that for infectious but in the opposite direction to the movements depicted in Figure 4.

[15]The situation where people may keep test outcomes private or not obtain tests is captured by the α below but the analysis does not inform on the issue created by that possibility as to whether it is desirable to have a testing regime relative to leaving individuals uninformed as to their infectiousness. The approach here could be used to analyse such cases but that is left to future work.

[16]They use a network rather than SIR model and so implicitly adopt the analytical shortcut proposed here.

There is, however, a second impact of testing – and specifically, isolation – that has not been examined in the literature. Testing followed by isolation reduces the utility from becoming infected as an agent would not expect to be able to freely choose their activity level in that event. Formally, their utility becomes $u_n(0)$ rather than $u_n(1)$ in that case. While (Eichenbaum et al. 2020) focused on how agents may avoid tests altogether, if agents are tested, those tests themselves will cause the impact identified here. Specifically, with a reduction in the utility of becoming infected, agents will become more cautious. This will shift the BEH curve to the left countering the impact of increased testing on the likelihood of encountering an infected person.

Putting the two impact mechanisms together, we can explore further whether the ambiguity may be removed if BEH, on net, shifted to the left. To explore this, let α be the probability that an infected agent is isolated as a result of testing regime. Given this, we have:

$$p(x_{n,S}(t), I(t)) = x_{n,S}(t)\beta(1-\alpha)I(t)$$

$$V_{n,I} = \frac{(1-\alpha)u_n(1) + \alpha u_n(0) - L + \delta(1-\gamma)\rho\frac{u_n(1)}{1-\delta}}{1-\delta\gamma}$$

Note that:

$$V_{n,S} - V_{n,I} = \frac{(1-\delta\gamma)u(x_{n,S}) - (1-\delta)((1-\alpha)u_n(1) + \alpha u_n(0) - L) - \delta(1-\gamma)\rho u_n(1)}{(1-\delta\gamma)(1-\delta(1-x_{n,S}\beta(1-\alpha)I^*))}$$

(OPT) becomes:

$$u_n'(\hat{x}_{n,S}) - \beta(1-\alpha)I^*\delta\left(V_{n,S} - V_{n,I}\right) = 0$$

Taking the derivative of the LHS of (OPT) with respect to α we have:

$$\beta I^*\delta\left(V_{n,S} - V_{n,I}\right) - \beta(1-\alpha)I^*\delta\frac{\partial(V_{n,S} - V_{n,I})}{\partial\alpha}$$

The first term is the marginal benefit to more *risk* as a result of testing while the second term is the marginal benefit to more *caution*. Note that:

$$\frac{\partial(V_{n,S} - V_{n,I})}{\partial\alpha} = \frac{\delta x_{n,s}\beta I^*(1-\delta\gamma)(V_{n,S} - V_{n,I}) + (1-\delta)(u_n(1) - u_n(0))}{(1-\delta\gamma)(1-\delta(1-x_{n,S}\beta(1-\alpha)I^*))}$$

which is positive for $V_{n,S} \geq V_{n,I}$. Putting the two effects together, the impact of α on the marginal return to activity is positive if:

$$V_{n,S} - V_{n,I} \geq (1-\alpha)\frac{\delta x_{n,s}\beta I^*(1-\delta\gamma)(V_{n,S} - V_{n,I}) + (1-\delta)(u_n(1) - u_n(0))}{(1-\delta\gamma)(1-\delta(1-x_{n,S}\beta(1-\alpha)I^*))}$$

Notice that as $\alpha \to 1$, this always holds. By contrast for $\alpha \to 0$, this becomes:

$$(1-\delta\gamma)(V_{n,S} - V_{n,I})_{\alpha=0} \geq u_n(1) - u_n(0)$$

which may not hold. Thus, while it is possible, for low α, that there may be an unambiguous comparative static that testing will reduce equilibrium infections, for high α, ambiguity remains. In this case, an increase in α (i.e., the effectiveness of testing and isolating) leads to a shift upwards in the BEH curve. In this model, therefore, as testing increases the relative safety of interactions this causes activity to rise by more than the effect driven by the decrease in the utility of the infected. Hence, the ambiguity remains for this comparative static.

4.4 Mandated masks

Encouraging the use of masks has been a strategy increasingly deployed and even mandated for dealing with COVID-19. In some medical circles there is debate regarding whether mandated masks would encourage less social distancing and potentially have a immiseration effect on infection rates ((Mantzari, Rubin and Marteau 2020)). As was the case with testing, the analytical approach here can be used to provide insight on that potential.

Suppose that, if all but recovered agents wear masks, the probability that the virus infects a susceptible person in an interaction with an infected one is $1-\alpha$; that is, a higher α means that a susceptible has more protection. Mask wearing is costly to individuals and, thus, it is assumed that all susceptible and infected agents, n, bear a cost, c_n, for each period they wear a mask. In this situation, the only difference between the impact of masks is this cost as well as the fact that infected people are not restricted in their activity and thus earn $u_n(1)$ while infected.

Thus, as was the case with testing, more effective masks (i.e., a higher α) leads to an increase in the returns to risky activity as well as a cautionary effect. The overall effect of masks is the impact of both. Note that, the impact of more effective masks on activity is positive if:

$$V_{n,S} - V_{n,I} \geq (1-\alpha)\frac{\delta x_{n,S}\beta I^*(1-\delta\gamma)(V_{n,S} - V_{n,I})}{(1-\delta\gamma)(1-\delta(1-x_{n,S}\beta(1-\alpha)I^*))} \implies 1 \geq \delta$$

where the last implication assumes that $V_{n,S} \geq V_{n,I}$. Thus, masks will always move the BEH curve to the right meaning that, given that they move the EPI curve to the left, there is no unambiguous comparative static result with respect to masks. Compared with testing and isolation, the returns to being infected are higher with mask wearing and so this reduces one driver of caution.[17]

5 The $\hat{\mathcal{R}} = 1$ Prediction

The analytical shortcut, whereby an equilibrium is analysed based on an assumption that S is fixed, gives rise to a prediction that $\hat{\mathcal{R}} = 1$. As noted earlier, this comports with the trends associated with the first few months of COVID-19 in a variety of countries that failed to suppress

[17](Gans 2021) and (Avery 2021) use a similar approach to the model here to examine the impact of vaccine policies (e.g., mandates) in a behavioural SIR model.

the pandemic. The question is: given that it is obtained using an analytical short-cut, how seriously should we take this prediction?

The potential error that arises from the short-cut can be seen by examining Figure 3. Note that rather than being constant, the share of susceptibles, S, will fall overtime. Indeed, if the level of infected persisted at I^*, S would fall by γI^* in each period. This means that the realised aggregate level of activity by susceptibles, X_S, would be expected to fall. This would not change the EPI curve as this change would be a movement along that curve. However, it would have an impact on the BEH curve. This is because the maximum value of X_S that can be generated by that relationship is S. Thus, a reduction in S may cause the feasibility constraint to bind. Without modelling how agents take into account the change in S in their own decisions – through expectations of a lower I in the immediate future – this curve, as derived, is only an approximation of what might occur.

That said, there is one special case for which a full equilibrium of the environment (sans the analytical shortcut) coincides with the limited equilibrium outcome examined thusfar. (Toxvaerd 2020) assumes that all agents are identical and that their activity choice $x \in 0, 1$. This gives rise to BEH as depicted in Figure 5. When $I(t) < I^*$, all agents choose $\hat{x} = 1$ and when $I(t) > I^*$, they choose $\hat{x} = 0$. He shows that I^* is independent of the share of susceptibles. The equilibrium arises when $I(t) = I^*$ and agents pursue a mixed strategy between $\{0, 1\}$. The total choosing $\hat{x} = 0$ averages $\frac{1}{S\mathcal{R}_0}$. Thus, $\hat{\mathcal{R}}_t = S\mathcal{R}_0 = 1$.

In Figure 5, it can be seen that as S falls, this reduces the maximum of BEH but otherwise leaves the line, and hence, equilibrium outcome in terms of infections and reproduction rate unchanged. This, of course, does not continue indefinitely. As (Toxvaerd 2020) shows eventually $S\mathcal{R}_0 < 1$ in which case, the equilibrium moves down the EPI line until the pandemic eventually ends. Thus, compared with the standard SIR model, in this model, the pandemic emerges and hits a ceiling of infecteds at $I(t) = I^*$ and stays that way until $S < \frac{1}{\mathcal{R}_0}$. Thus, the curve is not so much flattened as 'pancaked' at I^*.

This at least provides comfort that the $\hat{\mathcal{R}} = 1$ prediction is the outcome of a possible full equilibrium model. It also gives insight as to why the simulations of (Cochrane 2020) and (Keppo et al. 2020) were able to generate outcomes whereby a $\mathcal{R}_t = 1$ outcome was observed for considerable periods of time when calibrated with parameters based on COVID-19. Put simply, during the early months of the pandemic S was so large compared to I that it would not be expected to change very much meaning that $\mathcal{R}_t$ appeared to be relatively constant over time and close to 1 especially relative to SIR simulations that did not include a behavioural element.

6 Conclusions

Standard epidemiological models of pandemics often do not consider how susceptible, infected and recovered people will change their behaviour over the life cycle of the pandemic. Economists have made progress in building behavioural elements into these models but the

Figure 5: Binary Choice and Symmetric Agent Equilibrium

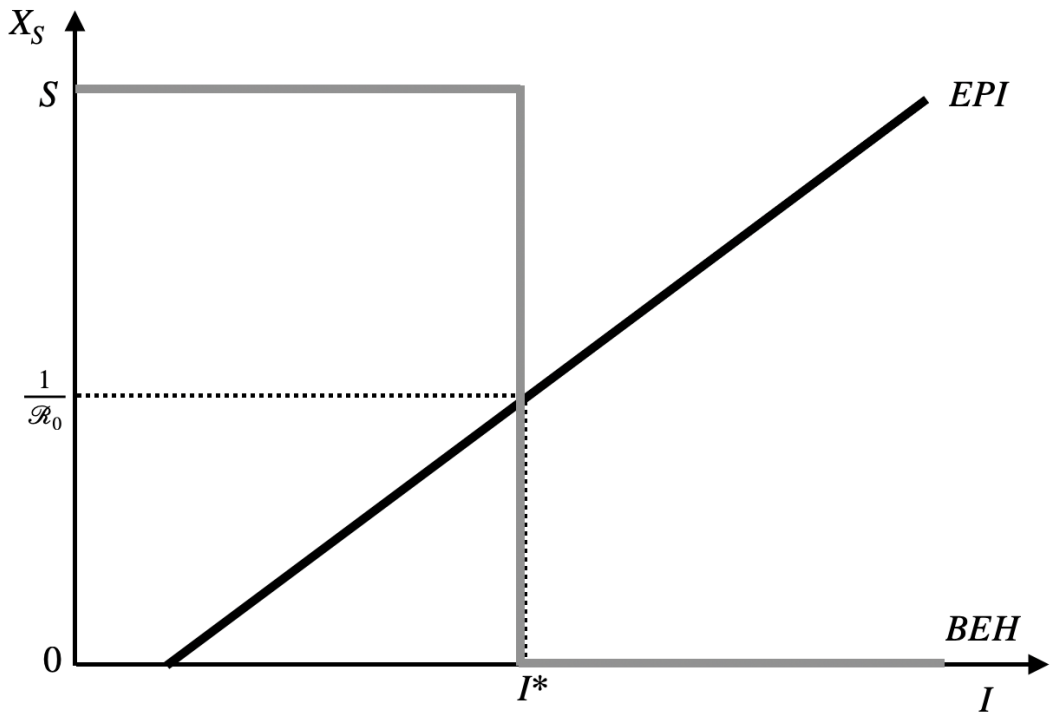

non-stationarity that is a key part of viral epidemics such as COVID-19 has prevented an easy characterisation of equilibrium paths of pandemics and the potential impact of interventions.

This paper argues that some analytical progress can be made on behavioural SIR models by taking inspiration for epidemiological models that do have stationary characteristics. In so doing, an equilibrium outcome is derived that allows intuitive comparative static outcomes on key variables such as infection rates and aggregate activity choices while at the same time generating a prediction that during much of a pandemic, without intervention, the effective reproduction number, $\mathcal{R}_t$ will tend towards 1. At this point, the infection rate is neither rising nor falling. This is consistent with the outcomes in many regions with respect to COVID-19. Nonetheless, the model here falls short of the usual requirements for a full equilibrium outcome. It does, however, have the benefit of being upfront about this limitation and what precisely we are getting in return in terms of tractibility and potential insight.

In doing this, this paper makes the case for treating $\hat{\mathcal{R}}_t = 1$ to be an expected outcome that can be used to evaluate, for both policy analysis and empirical predictions regarding pandemics. While being upfront regarding its 'cargo cultish' logic (i.e., based on observations of effective

reproduction numbers hovering around 1 rather than fully from primitive assumptions), as a shortcut it can provide some insight that might inform debates. For instance, (Budish 2020) has argued that $\mathcal{R}$ being just below 1 should be a constraint that is met by policies that impose lockdowns and other behaviour during pandemics. However, if the expectation is that, absent interventions, that goal would be mostly achieved anyway, that target is arguably of limited use compared, say, to a target of achieving $I = 0$ prior to what would otherwise be the natural course of pandemics. Moreover, with regard to lockdowns, the expectation that non-targeted lockdown activities may be adjusted to generate $\hat{\mathcal{R}}_t = 1$ requires us to not simply look at the epidemiological consequences of interventions (as (Acemoglu et al. 2022) do) but also to whether the non-targeted activities are such that they would be unable to adjust so that $\hat{\mathcal{R}}_t = 1$ was feasible. In other words, the criteria for lockdowns is not simply about spread but about the scope for behavioural adjustment.

Nonetheless, this analysis here remains purely normative. While it is tempting to conclude that if testing or mask use led to more infections this would be welfare-reducing, we must also remember the purpose of those interventions is precisely to allow activity to be safer and hence, allow for more of these at the margin. Thus, even though it is possible to draw some possible welfare conclusions from the fact that the BEH curve does not take into account external effects and so likely lies above a suitably derived social curve, the reality is more nuanced and requires an embrace of dynamic impacts. In particular, as (Rachel 2020) has shown the cumulative nature of pandemic impacts suggest that a focus on instantaneous external effects is unlikely to provide the correct insight into optimal policy-making.

References

Acemoglu, D., Chernozhukov, V., Werning, I. and Whinston, M. D. (2022), A multi-risk SIR model with optimally targeted lockdown, *AER: Insights*.

Acemoglu, D., Makhdoumi, A., Malekian, A. and Ozdaglar, A. (2020), Testing, Voluntary Social Distancing and the Spread of an Infection, *Technical report*, National Bureau of Economic Research.

Aguirregabiria, V., Gu, J., Luo, Y. and Mira, P. (2020), A Dynamic Structural Model of Virus Diffusion and Network Production: A First Report.

Allcott, H., Boxell, L., Conway, J. C., Ferguson, B. A., Gentzkow, M. and Goldman, B. (2020), What explains temporal and geographic variation in the early US coronavirus pandemic?, *Technical report*, National Bureau of Economic Research.

Atkeson, A. (2021), Behavior and the Dynamics of Epidemics, *Technical report*, National Bureau of Economic Research.

Atkeson, A. G., Kopecky, K. and Zha, T. (2021), Behavior and the Transmission of COVID-19, *AEA Papers and Proceedings*, Vol. 111, pp. 356–60.

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

Avery, C. (2021), A Simple Model of Social Distancing and Vaccination, *Technical report*, National Bureau of Economic Research.

Bethune, Z. and Korinek, A. (2020), COVID-19 Infection Externalities: Pursuing Herd Immunity or Containment?, *Technical report*, University of Virginia.

Bisin, A. and Moro, A. (2020), Learning Epidemiology by Doing: The Empirical Implications of a Spatial *SIR* Model with Behavioral Responses, *Available at SSRN 3625361*.

Boppart, T., Harmenberg, K., Hassler, J., Krusell, P. and Olsson, J. (2020), Integrated epi-econ assessment, *Technical report*, National Bureau of Economic Research.

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

Budish, E. B. (2020), R < 1 as an Economic Constraint: Can We 'Expand the Frontier' in the Fight Against Covid-19?, *University of Chicago, Becker Friedman Institute for Economics Working Paper*.

Chen, F. H. (2009), Modeling the effect of information quality on risk behavior change and the transmission of infectious diseases, *Mathematical biosciences* **217**(2), 125–133.

Chen, F. H. (2012), A mathematical analysis of public avoidance behavior during epidemics using game theory, *Journal of theoretical biology* **302**, 18–28.

Chen, F. H. and Toxvaerd, F. (2014), The economics of vaccination, *Journal of theoretical biology* **363**, 105–117.

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

Dasaratha, K. (2020), Virus dynamics with behavioral responses, *arXiv preprint arXiv:2004.14533*.

Deb, R., Pai, M., Vohra, A. and Vohra, R. (2020), Testing Alone Is Insufficient, *Available at SSRN 3593974*.

Di Guilmi, C., Galanis, G. and Baskozos, G. (2020), A Behavioural SIR Model and its Implications for Physical Distancing.

Eichenbaum, M. S., Rebelo, S. and Trabandt, M. (2020), The macroeconomics of epidemics, *Technical report*, National Bureau of Economic Research.

Eksin, C., Paarporn, K. and Weitz, J. S. (2019), Systematic biases in disease forecasting–the role of behavior change, *Epidemics* **27**, 96–105.

Ellison, G. (2020), Implications of Heterogeneous SIR Models for Analyses of COVID-19, *Technical report*, National Bureau of Economic Research.

Farboodi, M., Jarosch, G. and Shimer, R. (2021), Internal and external effects of social distancing in a pandemic, *Journal of Economic Theory* **196**, 105293.

Fenichel, E. P. (2013), Economic considerations for social distancing and behavioral based policies during an epidemic, *Journal of health economics* **32**(2), 440–451.

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

Fukuda, S., Kos, N. and Wolf, C. (2021), Epidemics with Behavior, *arXiv preprint arXiv:2103.00591*.

Gans, J. S. (2020a), *Economics in the Age of COVID-19*, MIT Press.

Gans, J. S. (2020b), *The Pandemic Information Gap: The Brutal Economics of COVID-19*, MIT Press.

Gans, J. S. (2021), Vaccine Hesitancy, Passports and the Demand for Vaccination, *Technical report*, National Bureau of Economic Research.

Geoffard, P.-Y. and Philipson, T. (1996), Rational epidemics and their public control, *International economic review* pp. 603–624.

Gersovitz, M. (2003), 25 Births, Recoveries, Vaccinations, and Externalities, *Economics for an imperfect world: Essays in honor of Joseph E. Stiglitz* p. 469.

Gersovitz, M. (2011), The economics of infection control, *Annu. Rev. Resour. Econ.* **3**(1), 277–296.

Gersovitz, M. and Hammer, J. S. (2004), The economical control of infectious diseases, *The Economic Journal* **114**(492), 1–27.

Goodkin-Gold, M., Kremer, M., Snyder, C. and Williams, H. (2020), Vaccination policy in the short and long run, mimeo., Harvard.

Goolsbee, A. and Syverson, C. (2021), Fear, lockdown, and diversion: Comparing drivers of pandemic economic decline 2020, *Journal of Public Economics* **193**, 104311.

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

Hellmann, T. F. and Thiele, V. (2020), A Theory of Voluntary Testing and Self-isolation in an Ongoing Pandemic, *Technical report*, National Bureau of Economic Research.

Jones, C. J., Philippon, T. and Venkateswaran, V. (2020), Optimal mitigation policies in a pandemic: Social distancing and working from home, *Technical report*, National Bureau of Economic Research.

Keppo, J., Quercioli, E., Kudlyak, M., Wilson, A. and Smith, L. (2020), The Behavioral SIR Model, with Applications to the Swine Flu and COVID-19 Pandemics, mimeo., Wisconsin.

Kermack, W. O. and McKendrick, A. G. (1927), A contribution to the mathematical theory of epidemics, *Proceedings of the royal society of london. Series A, Containing papers of a mathematical and physical character* **115**(772), 700–721.

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

Krueger, D., Uhlig, H. and Xie, T. (2020), Macroeconomic dynamics and reallocation in an epidemic, *Technical report*, National Bureau of Economic Research.

Makris, M. (2021), Covid and social distancing with a heterogenous population, *Economic Theory* pp. 1–50.

Mantzari, E., Rubin, G. J. and Marteau, T. M. (2020), Is risk compensation threatening public health in the covid-19 pandemic?, *BMJ*.

McAdams, D. (2020), Nash SIR: An Economic-Epidemiological Model of Strategic Behavior During a Viral Epidemic, *arXiv preprint arXiv:2006.10109*.

McAdams, D. (2021), The Blossoming of Economic Epidemiology, *Annual Review of Economics*.

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

Philipson, T. J. and Posner, R. A. (1993), *Private Choices and Public Health*, Cambridge, Mass.: Harvard University Press.

Rachel, L. (2020), An Analytical Model of Covid-19 Lockdowns, mimeo, London School of Economics.

Reluga, T. C. (2010), Game theory of social distancing in response to an epidemic, *PLoS Comput Biol* **6**(5), e1000793.

Rowthorn, B. R. and Toxvaerd, F. (2020), The optimal control of infectious diseases via prevention and treatment, *Technical report*.

Taylor, C. (2020), Information and Risky Behavior: Model and Policy Implications for COVID-19, mimeo, Stanford.

Toxvaerd, F. (2019), Rational disinhibition and externalities in prevention, *International Economic Review* **60**(4), 1737–1755.

Toxvaerd, F. (2020), Equilibrium social distancing, mimeo, University of Cambridge.