NBER WORKING PAPER SERIES

VACCINE HESITANCY, PASSPORTS AND THE DEMAND FOR VACCINATION

Joshua S. Gans

Working Paper 29075
http://www.nber.org/papers/w29075

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2021

Thanks to Ajay Agrawal, Avi Goldfarb and Richard Holden for helpful discussions. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

Vaccine Hesitancy, Passports and the Demand for Vaccination
Joshua S. Gans
NBER Working Paper No. 29075
July 2021
JEL No. I12,I18

ABSTRACT

Vaccine hesitancy is modelled as an endogenous decision within a behavioural SIR model with endogenous agent activity. It is shown that policy interventions that directly target costs associated with vaccine adoption may counter vaccine hesitancy while those that manipulate the utility of unvaccinated agents will either lead to the same or lower rates of vaccine adoption. This latter effect arises with vaccine passports whose effects are mitigated in equilibrium by reductions in viral/disease prevalence that themselves reduce the demand for vaccination.

Joshua S. Gans
Rotman School of Management
University of Toronto
105 St. George Street
Toronto ON M5S 3E6
CANADA
and NBER
joshua.gans@rotman.utoronto.ca

1 Introduction

Vaccine hesitancy, whereby individuals elect not to be vaccinated, has been a long-standing issue in public health. During the Covid-19 pandemic, such factors have meant that the adoption of vaccines has been short of what might generate lasting herd immunity in many populations. This has led to the creation of incentives to stimulate adoption, education campaigns to provide information on vaccine safety and calls for vaccine passports.

One factor that has been documented to counter vaccine hesitancy has been the prevalence of the relevant virus/disease itself. Oster (2018) found that an outbreak in a county in the year prior led to a 28 percent increase in childhood vaccinations. This raises an interesting question: *If vaccine adoption reduces prevalence, does that very adoption drive observed vaccine hesitancy?* And if there is an endogenous behavioural effect, does this change the impact of proposed policies to counter hesitancy?

This paper examines this issue. In so doing, vaccine adoption is endogenously driven by prevalence and vice versa. In addition, the underlying model is a behavioural epidemiological model where individuals can take costly social distancing actions to manage their own infection risk if they remain unvaccinated. The simple model presented here demonstrates that this can have a significant impact on the efficacy of various interventions to counter hesitancy. Blanket incentives/education for vaccination involve infra-marginal effects that are costly which limit their desirability as a means of encouraging vaccination. More critically, it is demonstrated that vaccine passports are ineffective (in this model, completely ineffective) at improving vaccination rates and may reduce them as the restrictions imposed by passports reduce prevalence. This implies that while vaccine passports (or credentials) may be useful in assuring others that someone has a lower risk of being infectious, they are unlikely to be useful as a punitive tool to counter vaccine hesitancy.

The literature on the economics of vaccines has proven more subtle than the textbook treatments of vaccination would suggest. These often state that vaccine adoption is, literally, a textbook example of a positive externality and so requires government intervention to encourage adoption and reduce free-riding. But as Francis (1997) demonstrated, when embodied within an epidemiological model of viral spread, those externalities do not necessarily manifest themselves in terms of socially suboptimal vaccinations. Indeed, he shows that with homogeneous agents and a perfect vaccination but for their personal cost of vaccination, the decentralised and socially optimal vaccination outcomes coincide.[1]

This paper builds on Francis (1997) and present a model with his underlying structure

[1]Chen and Toxvaerd (2014) demonstrate that when agents are heterogeneous and vaccination is imperfect, this equivalence not longer holds. See also Gersovitz (2003) and Toxvaerd and Rowthorn (2020).

of homogeneous agents but for vaccination costs and the introduction of a perfectly effective vaccine. But as policies such as vaccine passports involve costly restrictions on agent activity, a behavioural model is built where agents endogenously choose their activity/social distancing (using a similar structure to Toxvaerd (2020) but with discrete time).[2] Auld and Toxvaerd (2021) present preliminary evidence that vaccinations did impact on people's social activity in the Covid-19 pandemic. The focus here is not on the social optimality of any policies per se but their impact on total vaccination take-up.

The paper proceeds as follows. First, the standard susceptible-infected-recovered (or SIR) model is presented followed by its behavioural analog. Second, vaccines are introduced and the equilibrium with both endogenous activity and vaccine adoption is analysed. To resolve technical issues associated with non-stationarity, the analytical shortcut outlined in Gans (2020b) is utilised. Finally, various policies are explored and their impact on vaccination rates is characterised. A final section concludes with a list of the paper's obvious shortcomings in the hope of infecting someone to overcome them.

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 $\mathcal{R}_0 = \frac{\beta}{\gamma}$. $\mathcal{R}_0$

[2] Talamàs and Vohra (2020) also show that behavioural effects can impact on vaccination by changing the network of interactions – similar in spirit to Kremer (1996). However, their focus is on whether imperfect vaccines may create increased prevalence rather than on vaccine adoption per se which is costless in their model.

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.

Vaccination is a way of moving people from the susceptible to the recovered population segments. It is assumed here that vaccination is perfect. Of key interest in this model is the herd immunity threshold that is reached when $S(t)$ becomes low enough, through either acquired or vaccinated immunity, so that $\mathcal{R}_t < 1$. This threshold, for $1 - S(t)$ is $1 - \frac{1}{\mathcal{R}_0}$. A goal of many vaccination programs is to ensure that the total of the recovered and vaccinated shares exceeds this threshold allowing the pandemic to abate without any further interventions.[3]

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 of β on new infections. That variable is then assumed to be a decreasing function of $I(t)$; e.g., $x(I(t))$.[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; e.g., Philipson and Posner (1993), Geoffard and Philipson (1996), Kremer (1996) and empirically by Greenwood et al. (2019). The pioneering treatment that first introduced forward-looking, rational economic agents into epidemiological models was provided by Gersovitz and Hammer (2004). They explored 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. The literature is now extensive and has been reviewed by Philipson (2000) and Gersovitz (2011) and, more recently, Gans (2020b) and McAdams (2021). The models of Toxvaerd (2020)

[3]This is not to say that further vaccination beyond this threshold may be worthwhile as infections and health costs can continue thereafter. See Gans (2020c) for a discussion.

[4]See for example, Eksin et al. (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.

and Rachel (2020) examine microfounded models of endogenous social distancing comparing how decentralised outcomes compare with socially optimal outcomes.

In relation to vaccination, Rowthorn and Toxvaerd (2020) examines the appropriate mix of prevention and treatment while Goodkin-Gold et al. (2020) looks at vaccine pricing where epidemiological effects are anticipated and influenced. Makris and Toxvaerd (2020) examines how behaviour responds to the imminent arrival of a vaccine – showing that it tends to induce more caution. None of these behavioural models examine vaccine hesitancy that is the focus here.

3.2 Model Setup

For simplicity, it is assumed that all agents are the same in terms of their preference for activity and in terms of their costs of becoming infected. (Below, they will differ in terms of their preferences to being vaccinated).

Agents choose their level of activity, $x_i \in [0,1]$. This activity gives them a per period utility value of $u(x)$ where, for simplicity, this has a functional form of $u-(1-x)c$ (for $u \geq c$. 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 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, $x_R = 1$ and will earn an expected present discounted payoff of $\frac{u}{1-\delta}$. 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 $u-(1-x_I)c-L$ and their expected discounted payoff is:

$$V_I(t) = u-(1-x_I)c-L+\delta\Big(\gamma V_R + (1-\gamma)V_I(t+1)\Big)$$

where here $V_R = \rho\frac{u}{1-\delta}$. Note that, being self-interested, infected agents set $\hat{x}_I = 1$ in each period and, thus, their expected discounted payoff becomes:

$$V_I = \frac{u - L + \delta(1-\gamma)\rho\frac{u}{1-\delta}}{1-\delta\gamma}$$

3.4 Susceptible Agent Activity

For both the infected and recovered agents, 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 $u - (1 - x_S)c$ and their expected discounted payoff is:

$$V_S(t) = u - (1-x_S)c + \delta\Big(p(x_S, \hat{x}_I I(t))V_I + (1 - p(x_S, \hat{x}_I I(t)))V_S(t+1)\Big)$$

where $p(x_S, x_I I(t))$ is probability that an agent becomes infected at time t (the consequences of which are felt at time $t+1$).

The structure of $p(x_S, x_I I(t))$ 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. Thus, it is natural to posit that $p(x_S, x_I I(t)) = \beta x_S x_I I(t)$.

A susceptible individual will choose $x_S(t)$ to maximise $V_S(t)$ holding the state variables and their future path as given. Given this, note that:

$$\hat{x}_S = \begin{cases} 0 & c < \beta x_I I(t)\delta(V_S(t+1) - V_I) \\ x & c = \beta x_I I(t)\delta(V_S(t+1) - V_I) \\ 1 & c > \beta x_I I(t)\delta(V_S(t+1) - V_I) \end{cases} \tag{BEH}$$

where $x \in (0,1)$. Given this, the utility of a susceptible agent is:

$$\hat{V}_S = \begin{cases} \frac{u-c}{1-\delta} & c < \beta x_I I(t)\delta(V_S(t+1) - V_I) \\ \frac{u-c}{1-\delta} & c = \beta x_I I(t)\delta(V_S(t+1) - V_I) \\ \frac{u+\delta\beta x_I I(t)V_I}{1-\delta(1-\beta x_I I(t))} & c > \beta x_I I(t)\delta(V_S(t+1) - V_I) \end{cases} \tag{UTL-S}$$

To anticipate, in the equilibrium examined below, $\hat{x}_S = x$ and so agents will earn utility as if they were setting $x_S = 0$. This represents the lower-bound on the utility of a susceptible agent.

3.5 Equilibrium Analysis

Given this choice, the equilibrium outcome depends on the evolution of state variables taking into account the choices of all agents. 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 $X_S(t) \equiv x_S S(t)$. 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 $x_S = 1$ for all susceptible agents, 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.

This non-stationarity makes this whole thing a pain in the neck 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. However, he uses various approximations in order to analyse policies in the model. Gans (2020b) argues for an *analytical shortcut* inspired by SIS epidemiological models (that do have stationary equilibria). 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). The condition is a simple one: $S(t+1) = S(t) = S$ for all t.[5]

Note first that, from the law of motion in the SIR model and noting that $\hat{x}_I = 1$ we have:

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

With this shortcut, we can combine $\hat{x}_S$, setting this equal to $\hat{X}_S(I(t))$, and explore equilibria in which $I(t+1) = I(t)$ for all t. When this condition is satisfied then $\hat{x}_S(t+1) = \hat{x}_S(t)$ for all t which carries over to, $\hat{X}_S(I(t))$. Importantly, this means that:

[5]It is immediately is apparent that this condition violates the laws of motion of the SIR model whenever $\gamma > 0$. 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.

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

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.[6]

I^* can be found by solving BEH at the point where susceptible agents are indifferent between being infected or not. This gives:

$$I^* = \frac{c}{\beta\delta(V_S - V_I)} = \frac{(1-\delta)(1-\gamma\delta)c}{\beta\delta\Big((1-\gamma)\delta(1-\rho)u + (1-\delta)L - (1-\gamma\delta)c\Big)}$$

Given the analytical shortcut, this should be interpreted as an equilibrium prevalence that is mainly driven by exogenous model parameters than endogenous state variables during much of any pandemic.[7] To the extent that there are reductions in $S(t)$ overtime, these reductions result in an increase in $\hat{X}_S$ such that the probability of being infected remains constant overtime.

4 Vaccine Availability and Adoption

We now consider what happens when a vaccine becomes available. We will assume that the vaccine is perfect in that it allows a susceptible agent to be moved to the recovered category. Thus, an agent choosing to be vaccinated at t, will receive a benefit of $V_R - V_S(t)$. Note that $V_S(t)$ will change overtime and be impacted by the number of vaccinated people. In particular, the benefits from being vaccinated are typically higher when $I(t)$ is increased for t and beyond (Philipson (2000)).

To model vaccine hesitancy, agents must face an individual cost associated with being vaccinated. While the sources of such costs are nuanced and a mixture of real costs and perceived costs (MacDonald et al. (2015)), here we abstract from these sources by capturing them in a single parameter, θ. $\theta > 0$ is an on-going reduction in utility from being vaccinated.

[6] It can readily be seen that this equilibrium exists if $\mathcal{R}_0 > 1$. When $I(t) = 0$, all agents set $\hat{x}_S = 1$ 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 $I(t) = I^*$.

[7] Note that this is the same outcome as Rachel (2020)'s approximations.

It is assumed that θ is distributed amongst the population according to the distribution function, $F(\theta)$.

Each period after a vaccine is available, an agent chooses whether to vaccinate or not and, if not, want level of activity to choose. The following proposition characterises the equilibrium outcome:

Proposition 1 *If a vaccine becomes available at t, all agents for whom $\theta \leq c$ will be vaccinated and no additional vaccinations will be chosen after t. If $(1-F(c))S(t) < \frac{1}{\mathcal{R}_0}$, prevalence will decrease, otherwise it will remain constant at I^*.*

Proof. First note that, regardless of the number of susceptible agents, S in the population, in equilibrium, so long as $S\mathcal{R}_0 \geq 1$, the equilibrium prevalence will be I^* and the flow utility of a susceptible agent will be $u - c + \delta V_S$. By contrast, should an agent become vaccinated at t rather than $t-1$, their flow utility is $u - \theta + \delta V_R'$ where $V_R' = \frac{u-\theta}{1-\delta}$. Thus, it is easy to see that all agents for whom $\theta \leq c$ will be vaccinated and the total share of the population who become vaccinated at t will be $F(c)S(t)$ with $R(t) + F(c)S(t)$ being immune.

Second, if $(1 - F(c))S(t) < \frac{1}{\mathcal{R}_0}$, then even if all susceptible agents choose $\hat{x}_S = 1$, I^* will not be sustainable and prevalence will fall. This results in $V_S(t) > \hat{V}_S$ and, thus, no additional vaccinations will arise.

Third, if $(1 - F(c))S(t) \geq \frac{1}{\mathcal{R}_0}$, then the original equilibrium prevalence, I^*, will be maintained. This will also not induce any agents to be vaccinated as vaccinations will continue to be determined by whether c is greater than θ or not. Thus, no additional vaccinations will arise. ■

While prevalence, $I(t)$, impacts on $V_S(t)$ meaning that, other things being equal, an increase in $I(t)$ will induce some agents to vaccinate even if $\theta > c$, the equilibrium effects that arise from agents' endogenous activity choices immiserise this effect. Thus, all susceptible agents will be indifferent between being infected or susceptible. Moreover, the level of vaccinations in the population do not impact on prevalence so long as $\mathcal{R}_t \geq 1$. Thus, all agents for whom $c \geq \theta$ will choose to vaccinate immediately while others will never choose to vaccinate thereafter. Of course, if the availability of the vaccine means that $\mathcal{R} < 1$ and $I(t)$ will fall below I^*. However, this will mean that agents with $\theta > c$ will have no further inducement to be vaccinated.

This result relies critically on the fact that when $\mathcal{R}_t \geq 1$, I^* is a constant which, in turn, is derived from the analytical shortcut. However, as Rachel (2020) shows (see also McAdams (2021)), in behavioural models, $I(t)$, will be likely to fall over time as agents anticipate a point where $\hat{x}_S = 1$ and the net harm from becoming infected is increasing in t. Given this, it remains the case that all those who want to be vaccinated, will be vaccinated at time t.

5 Policies to Counter Hesitancy

5.1 Lowering vaccine cost

The first set of policies targets the cost θ associated with being vaccinated. While this may include persuasion and educational efforts, it is also the case that a direct monetary subsidy, τ, could also increase vaccination rates.[8]

While a welfare analysis involves subtle trade-offs, if we take a view, common amongst public health officials, that vaccination rates should be the minimum amount to cause $\mathcal{R}_t$ to fall below 1 (see also Budish (2020)), then we would set τ so that:

$$(1 - F(c + \tau))S(t) < \frac{1}{\mathcal{R}_0}$$

That is, τ is set to achieve herd immunity.

Note that, under our assumptions, in the absence of an intervention, agents who were previously infected and have recovered from the disease would not choose to vaccinated as they already do not socially distance, and so do not incur c, and face a strictly positive cost θ from being vaccinated. However, for those whom $\tau > \theta$, they will choose to be vaccinated under the intervention. Given this, the total cost of the subsidy program would be $\tau\big(F(c+\tau)S(t) + F(\tau)R(t)\big)$ in order to achieve a marginal effect on vaccine hesitancy of $(F(c+\tau) - F(c))S(t)$.

5.2 Vaccine passports

Another set of policies targets the utility of susceptible agents; namely, reducing that utility. The primary example of this is the vaccine passport. A passport policy restricts that activity, x, of *all* agents who do not show proof of vaccination. This is done by restricting activity of certain types or perhaps restricting other aspects of activity (such as mask-wearing requirements).

Here a simple view of a passport policy is taken and it is assumed that the policy caps the activity of all unvaccinated agents at $\bar{x}$. The idea here could be that activity along the $[0, 1]$ dimension is ranked from those that are most essential (close to 0) to those that are most optional (close to 1). Unvaccinated agents are prohibited from the most optimal activities where $x > \bar{x}$.[9] A benefit of this is that infected agents are kept away from those activities as

[8] For instance Chevalier et al. (2021) show that locating distribution centres closer to lower income neighbourhoods increases vaccination rates. Brehm et al. (2021) provide an analysis of lotteries in encouraging vaccinations.

[9] There are other ways of modelling passports. For instance, distinct activities could be explicitly mod-

are those who are still susceptible. In this way, one side-effect of this policy is to potentially reduce disease prevalence.

Note, first, that one impact of this policy is to limit the activities of recovered and infected agents. In the absence of a restriction, they would set $x = 1$. For the recovered agent, this is a pure welfare loss but it now means that those for whom $(1 - \bar{x})c \geq \theta$ will be vaccinated. Even though their status is temporary, infected agents for whom $(1 - \bar{x})c \geq \theta$ will also be vaccinated. However, as a matter of practice, this is the same as a recovered agent being vaccinated as vaccine dose regimens likely last more than the infectious period for most diseases. For practical purposes, however, we can treat infected people as unvaccinated and, thus, by implication, their activity is restricted to $\bar{x}$.

The target of the passport policy are the susceptible agents – both to counter hesitancy and protect them. The protection element is subtle. Given the restrictions on activity (and assuming they bind on the susceptible agents), the probability that a susceptible agent becomes infected is now $\beta\bar{x}^2 I(t)$ as both infected agents and susceptible agents are potentially restricted in their activity. However, while this is possible in terms of exposure probabilities, the fact that infected agents and susceptible agents are restricted to the same activities means that the probability that a given agent a susceptible agent encounters is infected is equal to the portion of infected people in the population unadjusted for activity level. Thus, following the spirit of analyses such as Kremer (1996) and Talamàs and Vohra (2020), it is reasonable to suppose that the probability of infection remains at most $\beta\bar{x}I(t)$.

Suppose, for the moment, that a susceptible agent who becomes infected and then recovers, will prefer not to be vaccinated if they had chosen not to be vaccinated while susceptible. (This will be checked in equilibrium). In this case, for unvaccinated agents, we have:

$$V_R = \frac{u - (1 - \bar{x})c}{1 - \delta}$$

$$V_I = \frac{u - L - (1 - \bar{x})c + \delta(1 - \gamma)\rho\frac{u-(1-\bar{x})c}{1-\delta}}{1 - \delta\gamma}$$

Note that, consistent with the assumption, $(1 - \bar{x})c < \theta$ for these agents. The lower $\bar{x}$ is, the lower are the utilities from being infected and then recovered. This will impact on the incentives of a susceptible agent in their choice of activity as it removes one of the benefits from being infected – not having to manage the risk of infection.

To examine this choice, note that the marginal condition $c = \beta I(t)\delta\big(V_S(t+1) - V_I\big)$ can

elled. However, the approach here surfaces the first-order effects and, thus, keeps the analysis simple.

be solved to yield the prevalence if susceptible agents are not constrained by $\bar{x}$:

$$I^* = \frac{(1-\delta)(1-\gamma\delta)c}{\beta\delta\Big(\delta(1-\rho)(1-\gamma)(u-c) - \big(1-\delta+\delta(1-\gamma)\rho\big)\bar{x}c + (1-\delta)L\Big)}$$

Importantly, I^* is increasing in $\bar{x}$. This is because, when x_S is unconstrained, then $V_S = \frac{u-c}{1-\delta}$ while V_I is decreasing in $\bar{x}$. Thus, reducing $\bar{x}$ makes $V_S - V_I$ larger and so susceptible agents are more likely to be cautious in equilibrium.

However, in terms of the choice of whether to vaccinate or not, anticipation of this equilibrium outcome means that the choice is identical to the case where there is no vaccine passport. It is easy to see that susceptible agents will only prefer to vaccinate if $c \geq \theta$. Thus, while reducing $\bar{x}$ reduces prevalence of the virus, it does not change the equilibrium amount of vaccination by susceptible agents. As for recovered agents, some will choose not to be vaccinated if $(1-\hat{x})c < \theta$. Thus, our earlier supposition that susceptible agents who choose to not be vaccinated will not subsequently become vaccinated holds.

Thusfar, the analysis has assumed that $\hat{x}_S \leq \bar{x}$. Recall that $\hat{x}_S = \frac{\gamma}{\beta S(t)}$. This condition is more likely to be satisfied if $\mathcal{R}_0$ and/or $S(t)$ are high. Vaccination, however, reduces $S(t)$ and as a lower $\bar{x}$ decreases I^*, we need to check whether this is possible in equilibrium.

To examine this, suppose now that $\hat{x}_S > \bar{x}$ and so the vaccine passport constrains all agents who are not vaccinated. Note that this implies that $S(t)\bar{x} < \frac{1}{\mathcal{R}_0}$ and so $I(t)$ will be falling over time. Given this constraint, note that:

$$V_S = \frac{u - (1-\bar{x})c + \delta\beta\bar{x}I(t)V_I}{1-\delta(1-\beta\bar{x}I(t))}$$

In particular, as $c > \beta I(t)\delta(V_S - V_I)$, then this V_S exceeds $\frac{u-c}{1-\delta}$ that arises when activity is unconstrained. Because activity involves an externality on all susceptible agents, constraining activity causes the utility of a susceptible agent to rise. Thus, as V_R is unchanged, this implies that the incentive to vaccinate, $V_R - V_S$, will fall as $\bar{x}$ is reduced. This means that the total number of vaccinated susceptibles will be lower than $F(c)S(t)$.

The above analysis can be summarised in the following proposition.

Proposition 2 *A vaccine passport policy with restriction $\bar{x}$ on unvaccinated agents will result in either the same or fewer vaccinated susceptibles compared with the case where there is no vaccine passport.*

In summary, either the vaccine passport restrictions bind on susceptible agents or they do not. If they do not bind, those restrictions reduce the utility of infected/recovered agents

which makes susceptible agents more cautious in equilibrium. The equilibrium prevalence adjusts to this making susceptible agents indifferent in terms of activity levels chosen. Thus, their utility remains as if they chose to mitigate all activity. This leaves the gain from being vaccinated the same as when there were no vaccine passports. If the restrictions do bind on vaccinated agents, prevalence is reduced and this increases the utility of susceptible agents as they receive more by meeting the constraint than choosing a zero level of activity. This means that the gain from vaccination is lower than when there were no vaccine passports. Hence, passports, at best, do not counter vaccine hesitancy and, at worse, increase it. Moreover, as discussed earlier, passports lead to unnecessary vaccinations from recovered agents or to unnecessary restrictions on their activity. Thus, they are likely to be strictly suboptimal for a social welfare perspective.

This analysis focuses purely on the motivation to use vaccine passports as a means of overcoming hesitancy. When a vaccine is imperfect (as all vaccines are) and some people cannot be vaccinated because the medical costs will always be too high, then a passport can be useful in assessing the riskiness of interactions between people and protecting unvaccinated people from becoming infected. In that sense, like testing, it is a means of overcoming the pandemic information problem (Gans (2020a)). Thus, a vaccine passport is not a set of rights per se but, instead, a means of efficiently communicating risk. For this reason, many countries in managing Covid-19 are adopting credentials that signal low risk for a person if they have been vaccinated or if they have received a recent negative test. In certain environments, such as higher risk health or aged care facilities, vaccines can become mandatory to protect some who are not vaccinated or who face higher risks with imperfect vaccines. However, this is not a tool of overcoming hesitancy per se but overall protection of the vulnerable.

6 Conclusions

The model here exposes some first-order interactions between vaccine hesitancy, social distancing and various policy interventions. However, like the work of Francis (1997) that is based on, it is simplified and the results may not be robust to the introduction of augmentations to the epidemic model, variable vaccine costs, imperfect vaccines, imperfect information, spatial heterogeneity, heterogeneity in susceptibility, behavioural variability and stochastic effects. These are all potentially fruitful directions for future research.

References

Auld, C. and Toxvaerd, F. (2021). The great covid-19 vaccine rollout: Behavioral and policy responses.

Brehm, M. E., Brehm, P. A., and Saavedra, M. (2021). The ohio vaccine lottery and starting vaccination rates. Technical report, Technical report, 2021. 345.

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*, (2020-31).

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

Chevalier, J. A., Schwartz, J. L., Su, Y., and Williams, K. R. (2021). Distributional impacts of retail vaccine availability. 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.

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

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

Gans, J. S. (2020b). The economic consequences of r= 1: Towards a workable behavioural epidemiological model of pandemics. Working Paper, NBER.

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

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

Gersovitz, M. (2003). 25 births, recoveries, vaccinations, and externalities. *Economics for an imperfect world: Essays in honor of Joseph E. Stiglitz*, page 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. M., and Williams, H. L. (2020). Optimal vaccine subsidies for endemic and epidemic diseases. Technical report, National Bureau of Economic Research.

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

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

MacDonald, N. E. et al. (2015). Vaccine hesitancy: Definition, scope and determinants. *Vaccine*, 33(34):4161–4164.

Makris, M. and Toxvaerd, F. (2020). Great expectations: Social distancing in anticipation of pharmaceutical innovations. mimeo.

McAdams, D. (2021). The blossoming of economic epidemiology. *Annual Review of Economics*, 13.

Oster, E. (2018). Does disease cause vaccination? disease outbreaks and vaccination response. *Journal of health economics*, 57:90–101.

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.

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

Talamàs, E. and Vohra, R. (2020). Free and perfectly safe but only partially effective vaccines can harm everyone. *Games and economic behavior*, 122:277–289.

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

Toxvaerd, F. and Rowthorn, R. (2020). On the management of population immunity. mimeo.