



A Simple Model of Social Distancing and Vaccination

Citation

Avery, Christopher. "A Simple Model of Social Distancing and Vaccination." HKS Faculty Research Working Paper Series RWP21-030, November 2021.

Published Version

<https://www.hks.harvard.edu/publications/simple-model-social-distancing-and-vaccination>

Permanent link

<https://nrs.harvard.edu/URN-3:HUL.INSTREPOS:37370346>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)



HARVARD Kennedy School
JOHN F. KENNEDY SCHOOL OF GOVERNMENT

A Simple Model of Social Distancing and Vaccination

Faculty Research Working Paper Series

Christopher Avery

Harvard Kennedy School

November 2021

RWP21-030

Visit the **HKS Faculty Research Working Paper Series** at:

https://www.hks.harvard.edu/research-insights/publications?f%5B0%5D=publication_types%3A121

The views expressed in the **HKS Faculty Research Working Paper Series** are those of the author(s) and do not necessarily reflect those of the John F. Kennedy School of Government or of Harvard University. Faculty Research Working Papers have not undergone formal review and approval. Such papers are included in this series to elicit feedback and to encourage debate on important public policy challenges. Copyright belongs to the author(s). Papers may be downloaded for personal use only.

A Simple Model of Social Distancing and Vaccination
Christopher Avery
NBER Working Paper No. 29463
November 2021
JEL No. D0,I10,I18

ABSTRACT

This paper analyzes a simple model of infectious disease where the incentives for individuals to reduce risks through endogenous social distancing take straightforward cost-benefit form. Since disease is transmitted through social interactions, the threat of spread of infection poses a collective action problem. Policy interventions such as lockdowns, testing, and mask-wearing serve, in part, as substitutes for social distancing. Provision of a vaccination is the only intervention that unambiguously reduces both the peak infection level and the herd immunity level of infection. Adoption of vaccination remains limited in a decentralized equilibrium, with resulting reproductive rate of disease $R_t > 1$ at the conclusion of vaccination. Vaccine mandates yield increases in vaccination rates and corresponding reductions in future infection rates but do not increase expected payoffs to individuals.

Christopher Avery
Harvard Kennedy School of Government
79 JFK Street
Cambridge, MA 02138
and NBER
christopher_avery@hks.harvard.edu

For acknowledgments and financial disclosure information, see:
<https://www.nber.org/papers/w29463>

1 Introduction

A recent New York Times article observed that self-interested behavior tends to counterbalance the effects of vaccination against infectious disease. “The administration’s emphasis on vaccines has undermined the importance of building other precautions into people’s lives in ways that are comfortable and sustainable”, referring to precautions such as “wearing a mask, or avoiding indoor dining and crowded bars.”¹ Early evidence (Auld and Toxvaerd (2021)) supports the contention of that article that increases in vaccination tend to be followed by reduction in protective behaviors. The fact that the introduction of a vaccine erodes incentives for these behaviors is ingrained in economic theory but still seems to be novel in some current policy discussions.

This paper develops a simple framework for analyzing the effects of policy interventions in response to an epidemic of an infectious disease and highlighting the economic intuition that underlies the predicted effects of those interventions. The analysis draws insights from several different strains of recent literature: (1) models of the trajectory of an epidemic with endogenous social distancing (Reluga (2010); Engle et al., (2021); Makris and Toxvaerd (2020); Toxvaerd (2020); Weitz et al., (2020); McAdams et al., (2021); McAdams (2021) provides an incisive summary of this literature); (2) models designed for calibration to the observed spread of COVID-19 (Atkeson (2021a, 2021b); Atkeson et al (2020b); Droste and Stock (2021)); (3) models studying optimal regulation in response to an epidemic (Acemoglu et al (2020); Baqaee et al (2020); Budish (2020); Eichenbaum et al., (2020); Farboodi et al., (2020); Rachel (2020); Alvarez et al., (2021)); (4) models studying the incentives for vaccination (Geoffard and Philipson (1997); Galvani and Reluga (2011); Galeotti and Rogers (2013); Chen and Toxvaerd (2014); Toxvaerd and Rowthorn (2020); Goodkin-Gold et al. (2021)). Appendix Table 1 summarizes the features included in papers with theoretical models of strategic social distancing and/or vaccination in response to an epidemic. Despite the explosion of recent papers related to COVID-19, there appear to be comparatively few studies that consider the interaction of vaccination and other policy interventions with endogenous social distancing.

The paper relies on a barebones discrete time framework for expositional purposes while utilizing an associated continuous time model for more formal analysis. The text focuses on the case of myopic social distancing behavior, whereby individuals choose current actions to minimize immediate expected costs rather than attempting to map out and respond optimally to the projected future course of infections; this assumption encompasses social distancing choices that depend only

¹<https://www.nytimes.com/2021/08/18/health/covid-cdc-boosters-elderly.html>

on the current infection rate. The baseline model yields individual-level incentives that can be summarized in simple cost-benefit terms, highlighting the economic factors underlying equilibrium with endogenous social distancing and vaccination decisions.

The assumption of myopic behavior has the twin virtues of plausibility, since the waxing and waning of COVID infection rates in 2020 and 2021 have consistently confounded expert predictions, and tractability, since a fully rational model is necessarily non-stationary as the proportion of previously uninfected people necessarily declines over the course of an epidemic.² The paper considers two primary alternatives to myopic behavior: *forward-looking optimization* where social distancing choices are based on fully rational optimization of present and future discounted payoffs and *pandemic fatigue* as suggested by Atkeson (2021a, 2021b) and Droste and Stock (2021). By comparison to myopic behavior, forward-looking optimization yields greater inclination for social distancing over time, since protective behavior has more benefits when the future is less risky, whereas pandemic fatigue induces the opposite behavior with less social distancing over time at a given infection rate. The assumption of myopic social distancing facilitates formal analysis and proofs; most of the qualitative results hold for all three descriptions of individual behavior.

The paper proceeds as follows. Section 2 presents the baseline model of infectious disease and endogenous social distancing. Section 3 provides equilibrium and comparative static analysis and considers the effect of individual policy interventions on the trajectory of infections in equilibrium. Section 4 extends the model to the case where a perfect vaccine (that provides full immunity from infection) is available and considers the effect of mandates designed to promote vaccination rates. Section 5 extends the model to allow for forward-looking social distancing choices as well as pandemic fatigue. Section 6 sketches an extension of the model with perfect vaccination to an imperfect vaccine that provides partial but incomplete immunity from infection. Section 7 concludes.

²Gans (2020) discusses this point in detail; Gans (2021) refers to nonstationarities in a fully rational model of infectious disease as a “pain in the neck”. McAdams, Song, and Zou (2021) show that steady state analysis is possible if births exactly offset deaths from infection.

2 The Model

Consider a discrete time model of infectious disease for a population of measure 1 where sick people are contagious for a single period and re-infection is not possible. The disease is characterized by exogenous parameters $T > 1$ for the baseline transmission rate, I_0 for the initial (small) positive proportion of individuals who are sick, δ for the per-period discount rate, and D for the cost of contracting the disease. This paper falls into the subcategory of game theoretic SIR (Susceptible-Infectious-Recovered) models, because each person is in one of these three categories at each moment in time and because social distancing and other decisions are made by individuals in a Nash equilibrium framework. Each person is assumed risk neutral with a separable utility function with total cost equal to the sum of the discounted costs (in utility units) of disease, financial outlays, and reduced activities from social distancing.

In the base case, the ordinary level of social activity per person is normalized to 1 and the level of social activity for infected people is given by an exogenous constant $A_I \leq 1$, where activity level $\alpha < 1$ corresponds to proportion α of that ordinary activity level.³ Section 3.4 relaxes the assumption that A_I is a known constant by allowing A_I to be determined endogenously in the case of asymptomatic infection. The time-varying parameters in the baseline model, indexed by t , are S_t for the proportion of susceptible (never infected) people, I_t for the current infection rate and A_{S_t} for the (average) activity level chosen endogenously by those who are susceptible.

In each period t , each susceptible person j chooses activity level A_{jt} , once again with activity level normalized so that $A_{jt} = 1$ in the absence of threat of infection. The probability that a person who is susceptible in period t is infected in that period is proportional to that person's number of interactions with those who are currently infected:

$$P(\text{person } j \text{ is infected in period } t) = TA_I A_{jt} I_t,^4$$

With no social distancing, the probability of infection for Person j in period t is $TA_I I_t$. Since each infectious person is assumed to be contagious for exactly one period, TA_I is equivalent to R_0 , the baseline reproduction rate per infection, when everyone is susceptible and there is no social distancing. The text assumes $TA_I > 1$ so that the disease tends to spread under baseline

³McAdams (2021) assumes $A_I = 1$ because infected people have no incentive to reduce activity to avoid illness. Altruism and the possibility of asymptomatic infection each suggest $A_I < 1$. Toxvaerd (2021) studies the effects of altruism and reduced social interactions when asymptomatic infections are possible.

⁴This formulation assumes random interactions between people. Acemoglu et al., (2020) allows for different contact rates across subgroups.

conditions. Aggregating across susceptible individuals, the period-by-period transition equation for the infection rate is proportional to A_{St} , the average activity rate for susceptible individuals in period t :

$$I_{t+1} = TA_I A_{St} I_t S_t,$$

In a homogeneous population, as assumed for most of this paper, all susceptible individuals choose the same activity level and so $A_{St} = A_{jt}$ for any susceptible person in period t .

The cost to person j for employing social distancing by forgoing proportion x of activities to reduce activity to $1 - x$ in a given period is $C(x)$, which represents the summed value of those forgone activities. $C(x)$ is assumed to be continuous with $C'(0) = 0$, $C'(x) > 0$ for $x > 0$ and $C''(x) > 0$, where these properties follow from time-use optimization by Person j according to the following logic. First, $C'(0) = 0$ follows from optimization under default conditions since then Person j must be indifferent between more and fewer activities at the default optimum $x_j = 0$ (corresponding to activity level $A_{jt} = 1$) with no risk of disease. Second, since $A_{jt} = 1$ with no threat of this particular disease, each inframarginal activity up to $A_{jt} = 1$ yields positive utility, thereby suggesting $C'(x) > 0$ for $x > 0$. Finally, optimization by Person j in response to the risk of infection suggests reduction of activities in priority order, starting with the least valuable activities, indicating $C''(x) > 0$, so that the earlier forgone activities are less costly than later forgone activities.⁵

With myopic maximization, each susceptible person j seeks to maximize expected payoff for a single period

$$\pi(x_{jt}, I_t) = -DP(\text{person } j \text{ is infected in period } t) - C(x_{jt}) = -TA_I A_{jt} I_t D - C(x_{jt}),$$

where $A_{jt} = 1 - x_{jt}$.⁶ Since I_t is the only time-varying parameter in the single-period utility function for susceptible person j , the optimal (myopic) one-period level of social distancing $x_{jt}^*(I_t)$ is a function of I_t alone.

⁵The case of constant marginal costs for social distancing $C(x) = Kx$, where K is a positive constant, greatly simplifies equilibrium analysis and has been studied by Gans (2020, 2021), Rachel (2020) and Toxvaerd (2019). We discuss the implications of this cost structure in greater detail in Appendix B.

⁶Cochrane (2021) illustrated by a series of examples that myopic social distancing based on observation of the current number of deaths rather than infections can produce a trajectory of infections with multiple peaks because social distancing is not synchronized with peak infection rates. Atkeson (2021a, 2021b), Atkeson et al., (2020b), Droste and Stock (2021), and Weitz et al., (2020) elaborate on this observation.

2.1 Continuous Time Version of the Model

The continuous time version of the model allows for instantaneous changes in all variables including the level of social distancing chosen by each person, with time-varying values written as functions of time t so that $S(t)$ and $I(t)$ represent the proportion of susceptible and infected people at time t . The time paths of these variables are characterized by their derivatives, which are drawn directly from the discrete time transition equations above:

$$S'(t) = -TA_I A_{St} I(t) S(t) = -TA_I (1 - x(t)) I(t) S(t).$$

$$I'(t) = TA_I (1 - x(t)) I(t) S(t) - I(t).$$

We can think of the continuous time model as the limit of the discrete time model where each period of the original discrete time model is divided into N subperiods as N tends to infinity. Assuming that cost of social distancing and the probability of infection transmission are proportional to the length of time for each subperiod, individuals face the same myopic incentives for the discrete and continuous time models and thus choose the same level of social distancing for a given infection rate in each case. One difference between the discrete time and continuous time model is that the length of infection is deterministic in the discrete time model and stochastic with mean 1 in the continuous time model.

A primary advantage of the continuous time formulation is that it yields an exact moment that the infection rate reaches a given value. By contrast, a discrete time model yields coarser sets of realized values for current and cumulative infection rates, which in turn can lead to an alternating series of infection rates above and below that threshold for some number of periods of time. For this reason, Proposition 1(d) and all of the formal properties of the model with vaccination rely on the continuous time version of the model.

3 The Trajectory of Infections with Endogenous Social Distancing

Herd immunity is reached when the susceptible population declines to the point where a new infection produces an average of one further infection in the absence of social distancing. Given this definition, herd immunity occurs when the remaining proportion of susceptible people is the inverse of the baseline transmission rate TA_I meaning $1 - H = \frac{1}{TA_I}$. Thus, the **herd immunity level** of cumulative infections is the proportion $H = \frac{TA_I - 1}{TA_I}$. Once herd immunity is reached, the

level of disease must gradually die out because any additional infection reduces the reproduction rate of the disease to less than 1, even without allowing for the possibility of social distancing.

A **full-population steady state** infection level \bar{I} produces the same probability of infection for susceptible people as the current rate of infection, accounting for social distancing: $TA_I A_{St} \bar{I} = I_t$. In practice, as emphasized by Gans (2021), \bar{I} is only an *approximate* steady state. Since the proportion of susceptible people in period t is always less than 1, if $I_t = \bar{I}$, $I_{t+1} = \bar{I} S_t < \bar{I}$. Proposition 1 records several standard properties of the time path of infection.

Proposition 1 (a) *There is a unique (myopic) optimal activity level A_{St} for susceptible people for each infection rate I_t with A_{St} strictly declining in I_t .*

(b) *The cumulative infection rate is eventually at least equal to the herd immunity level H .*

(c) *There is a unique full-population steady state infection level \bar{I} .*

(d) *The infection rate increases in each period until reaching a peak level less than \bar{I} .*

Property (a) follows from the fact that the probability of infection is linear in one's own activity rate, which means in turn that the marginal value of reduced risk per unit reduction in activity rate $TA_I I_t D$ is constant for any activity level for person j at time t . That is, Person j maximizes utility by conducting a cost-benefit assessment with constant marginal benefit and increasing marginal cost of forgone activity. The optimizing plan of activities corresponds to the first-order condition for social distancing at level x : $C'(x) = TA_I I_t D$.

Property (b) follows from the definition of herd immunity. Prior to reaching herd immunity, the susceptible population is large enough that the reproductive rate of new infections is greater than 1 without social distancing. If the infection rate ever becomes small when the cumulative rate of infections is less than H , susceptible people will choose close-to-normal activity levels and so the epidemic cannot die out at that point.

Property (d) follows from Property (c) and the fact that the optimal activity level for a susceptible person is declining in the current infection rate I_t . By assumption, I_0 is small and positive, so there is initially very little social distancing and thus exponential growth of infections. Since re-infections are not possible, S_t declines systematically over time. Each increase in I_t from period to period promotes social distancing, so A_{St} is declining at the start of an epidemic and the infection rate increases until the product $S_t A_{St}$ falls to less than 1, at which point the infection rate has peaked and begins to decline. The realized peak infection rate occurs when $S_t A_I = \frac{1}{A_I}$, which occurs at an infection rate $I_t < \bar{I}$.

3.1 Flattening the Curve

The concept of “flattening the curve” is inherent in the trajectory of infection with endogenous social distancing. Figure 1 depicts the relationship between the infection rate and the probability of infection for susceptible people in the current period after accounting for social distancing. At low levels of infection, the probability curve is steep with slope near $R_0 = TA_I$, reflecting exponential growth of disease when there is little immediate incentive to invest in social distancing. Increases in the infection rate promote social distancing, reducing the slope of the probability curve (and of the observed infection rate over time). In Figure 1, Point *A* represents the full-population steady state \bar{I} where the current infection rate is exactly equal to the probability of infection for current susceptible people, while Point *B* represents the observed peak infection rate.⁷ Consistent with the discussion above, the peak infection rate must be lower than this approximate steady state.




Figure 1. *Flattening the Curve*

Figure 1 also highlights two oft-emphasized, but essential observations. First, flattening the curve is inherent in laissez-faire equilibrium because initially exponential growth of a disease promotes social distancing, though there may well be policy motives to flatten the curve even more than in the ordinary social distancing equilibrium (e.g. to diminish the strain on hospitals). Second, flattening the curve has only limited effect on cumulative infections in the long-term. That is, endogenous social distancing spreads out the infection curve over time but does not change the

⁷Figure 1 uses the parameters $T = 1.5, A_I = 1, D = 8$ with $C(x) = \frac{x^2}{4}$. The peak infection rate at Point *B* is computed for initial rate of infection $I_0 = 0.02\%$ and each period divided into 5 subperiods.

herd immunity level, precisely because there is no incentive for social distancing when an epidemic dies out and the infection rate tends to zero.

3.2 Comparative Statics of the Steady State Infection Level

Comparative static analysis for the effects of exogenous parameters on the trajectory of the infection rate provides the foundation for policy applications because most interventions influence the course of an epidemic by changing the values of these parameters.

Comparative Statics of H

The herd immunity threshold H is increasing in T and A_I because $H = 1 - \frac{1}{TA_I}$. More precisely, H increases whenever the effective transmission rate TA_I increases. Since herd immunity comes into play in the end stage of an epidemic, comparative static analysis for H is predicated on a infection rate close to 0, at which point there is very little incentive for anyone to practice social distancing. For that reason, a policy intervention that results in a reduction of TA_I can only affect H if it maintains that effect on TA_I as disease dies out and the infection rate I_t is close to 0.

Comparative Statics of \bar{I} given changes in D

A change in D has no effect on T or A_I , so does not change the activity level for susceptible people $A_S^* = \frac{1}{TA_I}$ associated with the steady state. An increase in D increases the marginal value of social distancing, since D is the cost of an actual infection. For this reason, an increase in D induces an increase in A_S which in turn means that the activity rate corresponding to the population level steady state is chosen at a lower infection rate than before. That is, an increase in D reduces \bar{I} .

Comparative Statics of \bar{I} given changes in T or A_I

The parameters T and A_I appear as the product form in both the steady state identity $A_S = \frac{1}{TA_I}$ and the first-order condition for social distancing $C'(x) = TA_I ID$. That is, the value of the product TA_I influences the steady state infection level but, holding TA_I constant, the separate values of T and A_I do not influence I^* . For this reason, it makes sense to conduct comparative statics analysis as a function of this product TA_I of those two parameters.

Proposition 2 *An increase in the effective transmission rate TA_I can produce an increase or decrease in \bar{I} . If the cost function takes polynomial form $C(x) = \gamma x^n$ where γ is a positive constant and $n > 1$ (so that $C(x)$ is strictly convex), then \bar{I} is decreasing in TA_I for TA_I sufficiently large.*

Comparative statics for \bar{I} are determined by the effect of changes of exogenous parameters T and A_I on the identity $\bar{I} = A_S^{-1}(\frac{1}{TA_I})$. Representing A_S as a function of $\bar{I} = \frac{1}{TA_I}$ in this identity

serves to emphasize the fact that the value of \bar{I} adjusts to maintain the identity in response to a change in TA_I . For example, an increase in TA_I yields a reduction in the activity level A_S associated with \bar{I} . Yet, since an increase in TA_I increases the marginal value of social distancing, it is not obvious whether $A_S(I_t) = \frac{1}{TA_I}$ holds at the same value, a higher value, or a lower value of I_t when the product TA_I increases. Further, $\frac{1}{TA_I}$ is decreasing and strictly convex in TA_I and becomes quite flat as TA_I increases. Unless $C(x)$ is extremely convex, an increase in TA_I when TA_I is large produces a larger effect on the activity level A_S than on $\frac{1}{TA_I}$ and so will reduce the peak infection rate.

Figure 2a compares the equilibrium path of infection rates over time for $T = 1.5, 2$ and 4 with $A_I = 1$ with quadratic cost of social distancing $C(x) = \frac{\gamma}{2}x^2$ and $A_I = 1$ so that the initial reproductive rate R_0 is exactly equal to T .⁸ An increase from $TA_I = 1.5$ to $TA_I = 2$ shifts the infection rate curve to the left and also results in a higher peak. A further increase from $T = 2$ to $T = 4$ shifts the infection rate curve farther to the left but reduces the peak infection level. As the figure suggests, with quadratic costs of social distancing, each increase in the initial reproductive rate of disease increases the initial slope and lengthens the right tail of the infection rate curve, but does not necessarily increase the peak infection rate.




Figure 2a. *Comparative Statics of the Peak Infection Rate*

⁸Figure 2 uses the same parameters as Figure 1 apart from changes in the values of TA_I for two of the three curves. See the Appendix for calculations for the general case of quadratic costs. With quadratic costs, $TA_I = 2$ yields the highest value for the full-population steady state \bar{I}

Figure 2b graphs equilibrium peak infection rates as a function of the transmission rate T in this example with quadratic cost of social distancing. As illustrated in Figure 2a, the peak infection rate first increases and then decreases in T . At lower values of T , the peak infection rate is clearly below than the full-population steady state level, while at higher values of T , the peak infection rate approaches the full-population steady state level.

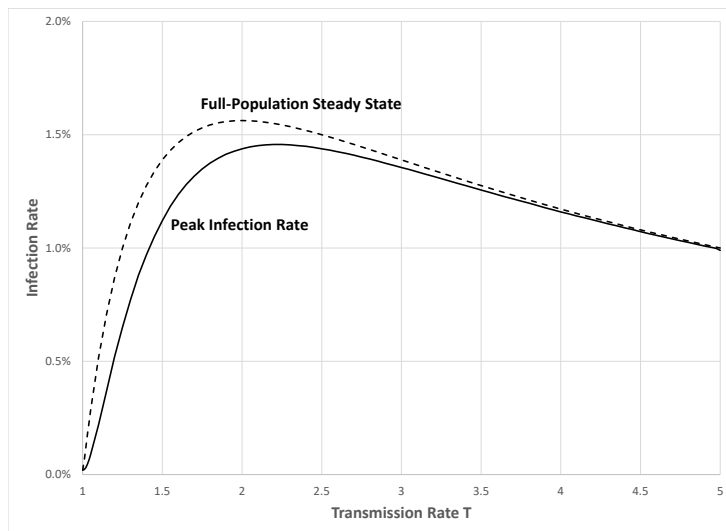


Figure 2b. *Comparative Statics of the Peak Infection Rate*

3.3 Implications for Policy Interventions

We classify policy interventions in two ways: (1) whether they are permanent or temporary; (2) their effects on the exogenous parameters of the infectious disease framework. Table 1 classifies interventions with regard to their effects on parameters and on their predicted effects on herd immunity H and steady state infection rate \bar{I} . Vaccination is the only intervention that unambiguously reduces both H and \bar{I} , and is also the only intervention that reduces the herd immunity level without costly effort when the infection rate has declined and is close to zero.

Intervention	Duration	Effect on Parameters	Effect on H	Effect on \bar{I}	Reference
Lockdown	Temporary	Restricts A_I, A_S	None	Possible Reduction	Section 3.3.1
Tests	Possibly Permanent	Reduction in A_I	Reduces (*)	Ambiguous	Section 3.3.2
Masks	Possibly Permanent	Reduction in T	Reduces (*)	Ambiguous	Section 3.3.3
Vaccination	Permanent	Reduction in S	Reduces	Reduces	Section 4.1

Table 1. Summary of Policy Interventions

(*) Tests and Masks only reduce H if continued even when $I(t)$ approaches 0.

3.3.1 Lockdowns

A lockdown is a government-imposed reduction in the set of possible activities. A full lockdown order confines individuals to their homes and could be viewed as reducing activity to 0, or perhaps close to 0 because of continuing interactions between members of the same household. As practiced in 2020 and 2021, lockdowns in most parts of the world were generally partial, in part because people in designated “essential” jobs were allowed to continue to work. Approximately 40% of jobs in the U.S. count as “essential” based on the definition of the Department of Homeland Security (McCormack et al., (2020); see also Blau et al., (2021)).

By nature, a lockdown is temporary and cannot be expected to affect the herd immunity threshold. Although an explicit goal of a lockdown is to flatten the curve, there are several reasons why it may not be effective. If we consider a partial lockdown as an upperbound \bar{A} on individual activity levels, the restriction may not be restrictive enough to bind: if the optimal level of activity chosen by susceptible individuals in the absence of a lockdown is less than \bar{A} , then the lockdown actually has no effect. Even if \bar{A} does restrict activity, it may not do so in an economically efficient manner – it may be implemented in a way that rules out high-priority activities for an individual but allowing that person to continue with lower-priority activities. That is, if $\bar{A} = 0.7$, individuals might choose to forgo a different 30% of activities than the ones that are precluded by a lockdown.

A lockdown does not change the full population steady state \bar{I} , but will likely both delay the timing and reduce the magnitude of the realized peak infection rate. Accounting for the number of susceptible individuals remaining at time t , the actual peak of infection is reached when $TA_I A_S S_t = 1$. Thus, for example, a lockdown may result in a peak at a time with relatively low value of S_t , since some people will get sick during the partial lockdown, which would reduce the peak rate of infection.⁹

3.3.2 Testing

Widespread testing has the potential advantage of alerting people with asymptomatic infections that they are sick and should obey quarantine rules to protect others. The effects of testing listed in Table 1 are drawn from the simplifying assumption that there are no false test results so that a testing regime can be summarized by a reduction in the value of A_I , the average activity level

⁹In this case, the realized peak infection rate occurs when $A_S = \frac{1}{TA_I S_t}$, so any reduction in A_S means that the identify holds at a higher value A_S and lower level of infection than without the lockdown.

for people when they are infectious.¹⁰ Several considerations influence the role of testing in the context of the model: (1) the proportion of people who are tested and the frequency with which they are tested; (2) the proportion of false negatives and positives at a given time (which depend both on the general accuracy of the test and the current infection rate since, for example, there can't be any false negatives if no one is sick); (3) the degree of restriction on activities for those who have tested positive. A further question is whether the government will continue to carry out an aggressive testing regime at times when the infection rate is relatively low. Testing can only reduce the herd immunity level if it is maintained during the period of time when the infection is receding, since otherwise TA_I returns to its baseline value when the infection rate tends to zero and thus H is maintained at its original value.

3.3.3 Masks

Mask wearing reduces the rate of transmission per interaction between a susceptible person and an infected person, especially for indoor interactions. When Person j and Person k interact, a choice by Person j to wear a mask likely has asymmetric effects and may well be more beneficial to Person k than to Person j . The results listed in Table 1 rely on the simplifying assumption that a mask-intervention reduces the value of T to a new average level $T_m < T$ for all people, corresponding, for example, to a government regulation observed by a constant proportion of people across time.

Under these assumptions, testing and mask-wearing have analogous effects since each reduces the effective transmission rate TA_I to susceptible people without social distancing. There are considerable subtleties to both testing and mask wearing beyond the results summarized in Table 1. The effect of false positives and false negatives produced by an ambitious testing regime could be considerable. At relatively high rates of disease, false positives could have an unintended positive consequence – helping to solve the collective action problem of reducing activity towards the socially optimal rate given the positive externalities created by social distancing. At a low rate of disease, false positives dominate false negative tests, likely resulting in substantive costs of forgone activities by people who are not sick.

Similarly, there are a number of unmodeled incentives weighing for and against the choice to wear a mask. Mask wearing by others reduces the incentive to wear a mask to protect oneself from infection, but also may result in newfound social pressures to conform and wear a mask as well. Thus, one could consider the level of mask-wearing in a given period as the equilibrium of a binary

¹⁰Atkeson et al., (2020a) provides a more sophisticated treatment of the economic value of testing.

choice population game.¹¹ Since it is comparatively easy to put on or take off a mask, this binary choice game is repeated over time, with likely positive correlation in each person’s choice across time periods – this description gives the mask-wearing game some resemblance to the model of imperfect vaccination that is discussed in Section 6.

3.4 Asymptomatic Infection

With purely asymptomatic infection, infected and susceptible people can be expected to choose the same level of activity. For simplicity, assume that everyone who was not previously sick assumes that they are susceptible at time t (even though some are currently infected) so that the first-order condition for social distancing remains the same as above: $C'(x) = TA_ID$. From the perspective of an individual person, A_I is a fixed constant, but in fact, it will be determined endogenously by this condition and the further restriction that $A_I = A_S$, meaning $A_I = 1 - x$. We can write the steady state condition as $A_I A_S = \frac{1}{T}$ or $A_S = \frac{1}{\sqrt{T}}$ since $A_I = A_S$ given these conditions. The comparison between steady states with A_I as a fixed value and $A_I = A_S$ turns on the ordinal ranking of (1) $\frac{1}{A_I T}$ when A_I is a constant and (2) $\frac{1}{\sqrt{T}}$ when $A_I = A_T$.

In the case of quadratic cost of social distancing $C(x) = \frac{\gamma}{2}x^2$, we can still solve for \bar{I} in closed form, though this seems to be more difficult for the more general case of polynomial costs and asymptomatic infection. With quadratic costs and $A_I = A_S$, the first-order condition for social distancing x is $x = \frac{TA_ID}{\gamma}$ so $A_S = 1 - \frac{TA_ID}{\gamma}$. Substituting the condition $A_I = A_S$ gives $A_S = 1 - \frac{TA_S ID}{\gamma}$ or equivalently $A_S(\gamma + TID) = \gamma$ with solution $A_S = \frac{\gamma}{\gamma + TID}$. This corresponds to steady state $\bar{I} = \frac{\gamma}{TD}(\sqrt{T} - 1)$, which is increasing if $T < 4$ and decreasing if $T > 4$.

4 Equilibrium Adoption of Perfect Vaccination

Now consider an extension of the model to allow for endogenous vaccination, focusing on the case where a perfect vaccination is available from the start of the epidemic. Vaccination costs $C_V \geq 0$ (the same cost for everyone). and a one-time dose is sufficient to provide complete protection from infection once and for all for the recipient. In equilibrium, each person chooses whether and when to get vaccinated to maximize expected payoff, i.e. minimizing the expected cost of

¹¹McAdams (2020) describes how social distancing and other preventative actions can be either strategic complements or substitutes in the context of an epidemic - see also Engle et al., (2021).

vaccination combined with the expected cost of possible infection.¹² All propositions in the text rely on subgame perfect equilibrium, where each person maximizes expected payoff given the dynamic strategies of others since vaccination by one person provides positive externalities by reducing the future probabilities of infection to others. Using the continuous time version of the model makes it possible to pinpoint the exact moment that the infection rate reaches a particular threshold.

Proposition 3 *There is a unique equilibrium of the continuous time vaccination game. In equilibrium, all vaccination takes place at the first moment t^* where the infection rate surpasses a constant threshold I^* .*

The logic of Proposition 3 follows from a heuristic cost-benefit analysis as described below for a single period of time in a discrete time setting with no social distancing. A default plan of vaccinating right away has immediate cost C_V . The nominal benefit of waiting one period is a reduction in cost from C_V to δC_V . The cost of waiting one period is an increase in this cost from δC_V to D if infected in the current period. Denoting the probability that a susceptible person is infected in the current period by P_t , the costs $P_t(D - \delta C_V)$ outweigh the benefits $(1 - \delta)C_V$ of waiting one period for vaccination if $P_t(D - \delta C_V) \geq (1 - \delta)C_V$ which takes the form $P_t \geq P^*$. Without social distancing, P^* is the known constant $\frac{(1-\delta)C_V}{D-\delta C_V}$. Since the probability of infection in period t is proportional to I_t , the infection rate in period t , the threshold probability $P^* = \frac{(1-\delta)C_V}{D-\delta C_V}$ corresponds to a threshold value for the infection rate I^* . With endogenous social distancing, the comparison is complicated by the inclusion of the cost of social distancing, but the cost-benefit result is of the same form – it is preferable to vaccinate now than to wait and do so at the next opportunity if the infection rate is above a threshold I^* .

In equilibrium, vaccination results in an immediate change in the slope but maintains the hump-shaped pattern and single peak of the infection rate curve. The cost-benefit comparison indicates that it can only be individually optimal to vaccinate at the first time t^* when $I(t^*) = I^*$, as otherwise it would be possible to increase expected payoff by delaying vaccination (if $I_t < I^*$) or accelerating the timing of vaccination (if $I_t > I^*$). Further, each increase in vaccination rate at time t^* reduces the future spread of disease and thus reduces incentives for others to vaccinate.

Given this logic, there is a unique equilibrium where either (1) no one vaccinates because vaccination at time t^* brings lower expected payoff than remaining unvaccinated forever when no

¹²We assume economically rational vaccination decisions by individuals even though that is not consistent with myopic social distancing; Section 5.2 discusses an extension of the model to a fully rational version of the model with forward-looking social distancing decisions and economically rational vaccination decisions.

one is expected to vaccinate; (2) vaccination proceeds at time t^* until the point where the expected payoff from vaccination is exactly equal to the expected payoff from remaining unvaccinated. Note that with perfect vaccination, there cannot be an equilibrium where everyone chooses to vaccinate at time t^* because the disease would die out quickly after that, leaving an incentive for the marginal person to remain unvaccinated.

Figure 3 depicts the Nash equilibrium trajectory in a discrete time numerical example with no social distancing and I^* approximately equal to 0.1%.¹³ Initially, infections proceed at exponential rate with no vaccination. When the infection rate crosses the critical threshold, a bit more than a third of the population chooses to vaccinate.

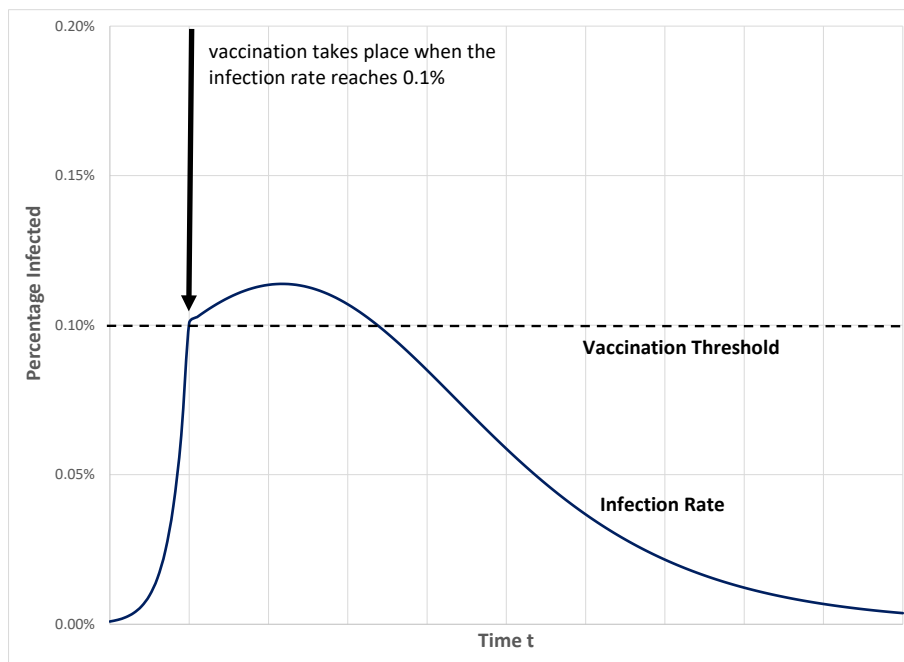


Figure 3. *Equilibrium Infection Trajectory with Perfect Vaccination*

Figure 3 illustrates a key property of equilibrium with a perfect vaccination, which is that the infection rate is initially increasing beyond the point of vaccination at time t^* . There is an interesting juxtaposition between the optimal vaccination strategy and cost-benefit analysis. Since the infection rate typically increases to a peak level and then declines, vaccination is only cost effective for Person j on a moment-by-moment basis from t^* until the next moment t^{**} when $I(t^{**}) = I^*$. Though vaccination is not *individually* cost-effective on a moment-by-moment basis from the perspective of an individual person in the end stages of an epidemic when $I(t) < I^*$, it is

¹³The parameters for this example are $T = 1.6, A_I = 1, \delta = .95, C_V = 1, D = 3.$ and $I_0 = 0.00092\%$.

not possible to choose to be vaccinated during peak periods and then return to being unvaccinated during fallow periods of an epidemic. Instead, equilibrium decision making involves a tradeoff between per-period expected gains from vaccination just after t^* when $I(t) > I^*$ and expected losses for $t > t^{**}$ and $I(t) < I^*$. If vaccination at time t^* proceeds to the point where $I(t) < I^*$ for all $t > t^*$, then vaccination is not cost effective for any period of time, which motivates the following Corollary.

Corollary 1 *Equilibrium adoption of vaccination in the continuous time vaccination game yields reproductive rate $R_{t^*} > 1$ at the time of vaccination.*

4.1 Policy Effects of Perfect Vaccination

Perfect vaccination reduces the herd immunity level on a one-for-one basis as a direct effect of reducing the overall susceptible population. If proportion V of the population receives a perfect vaccination, herd immunity requires only proportion $H - V$ of the population to have been infected.¹⁴ Similarly, perfect vaccination reduces the full-population steady state infection level \bar{I} according to the condition $TA_I A_S = \frac{1}{1-V}$ – corresponding to a higher activity rate for remaining susceptible (i.e. unvaccinated) people and thus a lower infection rate than without vaccination.¹⁵

As suggested by the New York Times article, vaccination and social distancing strategies are substitutes: the reduction of risk from either action reduces the value of the other. In equilibrium, fewer people choose to vaccinate when social distancing is possible in future periods than when it is not, precisely because future social distancing by others reduces both the immediate and cumulative risk of being unvaccinated. At the same time, these two methods of reducing risk are supplementary, with fewer cumulative infections in equilibrium when a person can choose between them than when only one is available.

4.2 Vaccine Mandates

We consider two versions of vaccine mandates: (1) a work-based requirement imposed by employers and (2) an activity-based requirement imposed by local governments and businesses. A number of

¹⁴The herd immunity threshold would decline instead to $H(1-V)$ if unvaccinated people maintained overall activity rates and interacted only with each other.

¹⁵If the baseline value for \bar{I} is higher than the threshold for vaccination, then no one will choose vaccination; in that case, vaccination has no effect on either H or \bar{I} .

employers have adopted policies with deadlines for employees to be vaccinated.¹⁶ President Biden has enacted regulation requiring federal contractors and companies with more than 100 employees to require vaccination for their workers. This paper models a *work-based vaccine mandate* as a per-period fine c_F (in utility terms) for anyone who is unvaccinated.

In August 2021, New York City imposed a new rule requiring proof of vaccination for indoor dining at restaurants.¹⁷ France adopted a more expansive requirement for a “pass sanitaire”, requiring government certification that one is vaccinated for access to cultural venues (as of July 21, 2021), restaurants, cafes, and bars (as of August 9, 2021) with initial approval to impose this rule through November 15, 2021. This paper models an *activity-based vaccine mandate* as a constraint limiting activities by unvaccinated people to $A_M < 1$, imposing minimum cost of social distancing $C(1 - A_M)$ per unit-time since it is still possible for an unvaccinated person to choose an even greater level of social distancing than is indicated by the new regulation.¹⁸

The analysis below assumes that the details and timing of the mandate are anticipated in advance and that it takes effect after time t^* , the moment when vaccination takes place in equilibrium without the mandate.¹⁹

Proposition 4 *An anticipated work-based vaccine mandate in the continuous time vaccination game that begins at time $t_M > t^*$ increases vaccination rates and reduces the cumulative infection rate, but has no effect on the expected utility for any person.*

A vaccine mandate changes incentives and equilibrium actions by making it unappealing to remain unvaccinated beyond time t_M . A work-based vaccine requirement provides unambiguous incentives for vaccination because it imposes direct costs on those who are unvaccinated, but does not alter incentives for the timing of a perfect vaccination.²⁰ If a work-based vaccine is anticipated,

¹⁶<https://www.nbcnews.com/business/business-news/here-are-companies-mandating-vaccines-all-or-some-employees-n1275808>

¹⁷<https://www.cnn.com/2021/08/15/new-york-city-vaccine-mandate-presents-new-challenges-for-restaurants.html>

¹⁸This framework assumes that the mandate pertains to proportion $1 - M$ of activities that are of lowest value to each person.

¹⁹ If sufficiently punitive, an unanticipated mandate imposed after some people have been vaccinated at t^* induces additional vaccination, thereby reducing subsequent infection rates but also reducing expected utility for those who are unvaccinated at the time of the mandate.

²⁰There is one important exception: if it is never cost-effective to vaccinate in the absence of a mandate, the mandate may yield a new equilibrium with vaccination at time t_M , in which case the mandate results in an increase in the vaccination rate but a reduction in expected utility for each person in equilibrium.

it induces an increase in vaccination rates at time t^* and thus lower cumulative infection rates thereafter than in the equilibrium of Proposition 3. However, since all vaccination continues to take place at time t^* , a work-based mandate does not change the infection rates prior to t^* or the expected payoffs for vaccination at time t^* . Further, given the equilibrium indifference condition for vaccination at time t^* , expected payoffs for those who remain unvaccinated are also unaffected by the work-based mandate.²¹

It may seem natural to relax a work-based vaccine mandate once the infection rate reaches a sufficiently low level, both since there is diminished future risk of infection and because it is implausible that anyone else will choose to be vaccinated from then on. Yet, any relaxation of a vaccine mandate undermines its purpose.

Corollary 2 *Ending a work-based mandate at a predetermined time or infection level reduces vaccination rates, increases cumulative infections, and has no effect on the realized expected utility for either those who choose to vaccinate at time t^* or those who remain unvaccinated.*

An activity-based vaccine requirement has an additional effect because it constrains activity in some periods subsequent to vaccination and thereby changes the future trajectory of infection conditional on a given level of vaccination. Since there are positive externalities associated with social distancing, mild restrictions on activities (i.e. A_M close to 1) may yield Pareto-gains while very strong restrictions on activities (i.e. A_M close to 0) yield Pareto-losses for the unvaccinated relative to the vaccine equilibrium with no mandates. Thus, if A_M is sufficiently small, an activity-based mandate is analogous to a work-based mandate.

Corollary 3 *An activity-based vaccine mandate that begins at time $t_M > t^*$ has the same qualitative effects as a work-based vaccine mandate if A_M is sufficiently small.*

4.2.1 Comment on the Expected Utility Effect of Vaccine Mandates

Although a vaccine mandate does not change expected utility for anyone in equilibrium, it shifts the burden for the unvaccinated from costs imposed by disease to financial losses and losses from forgone activities. Since the vaccine mandate imposes new forms of costs on the unvaccinated, their discounted expected losses from infection and thus their discounted probability of infection must also fall. In a more expansive model (such as Baqaee et al., (2020) and Alvarez et al., (2021)),

²¹There is an exception here as well. If the penalty for a work-based mandate is sufficiently large, then everyone will choose vaccination at time t^* to avoid the penalty.

these reductions in infection rates might yield population-level externalities - perhaps by allowing for more aggressive “re-opening” practices - that don’t alter the incentives for social distancing and vaccination that are present in the current model. In that case, a vaccine mandate would produce Pareto gains with universal gains in expected utility.

4.3 Heterogeneous Costs of Vaccination

There is a growing body of anecdotal evidence that some people are more averse to vaccination than others. This observation suggests an adjustment to the model with heterogeneous costs of vaccination across people according to continuous cumulative distribution function $F_{CV}(x)$. This heterogeneity of costs spreads out the timing of vaccinations but produces similar qualitative results to those for the model with homogeneous vaccination costs.²²

Proposition 5 *With heterogeneous cost of vaccination, there is a unique equilibrium with gradual vaccination during time period $[\underline{t}^*, \bar{t}^*]$ in order of ascending vaccination cost and $R_t > 1$ at the completion of vaccination. A vaccine mandate that begins at $t_M > \bar{t}^*$ increases vaccination rates and reduces future infections but reduces realized expected utility for all types with vaccination cost too high to vaccinate in the equilibrium without a vaccine mandate.*

5 Alternative Models of Social Distancing

This section considers extensions with alternative descriptions of individual behavior. Forward-looking social distancing accounts for future risk, which is ignored by myopic social distancing, and thus involves less protective behavior at any infection level than myopic behavior. Pandemic fatigue also involves less protective behavior than myopic social distancing, but only in the later stages of an epidemic. If the same infection rate occurs twice in the course of an epidemic, forward-looking social distancing predicts greater endogenous social distancing at the later time because there is less future risk for a given infection rate when the infection rate is declining than when it is increasing, but pandemic fatigue yields the opposite prediction.

The choice among these three models has little effect on herd immunity because they all yield similar behavior with very little social distancing near the point of herd immunity when the infection rate is low and $R_t \rightarrow 0$. In all three cases, social distancing remains quite responsive to changes

²²See the discussion of linear utility in Appendix Section E for a comparison to Gans (2021), which considers a vaccination mandate with heterogeneous cost of vaccination and no social distancing.

in the infection rate, which limits the possibility of large and lasting differences in equilibrium trajectories for the infection rate across these models.

5.1 Pandemic Fatigue

Assume that pandemic fatigue causes a change in the cost function from $C(x)$ to $C_{PF}(x)$ from time t_{PF} onwards, where $C'_{PF}(x) > C'(x)$ for each x so that the marginal cost of social distancing increases after pandemic fatigue sets in. (Assume also that $C_{PF}(x)$ retains the essential properties of $C(x)$, i.e. it strictly increasing and strictly convex.) Since pandemic fatigue has no effect on the cost function prior to time t_{PF} , the realized infection rate is the same with or without pandemic fatigue until some people choose to vaccinate at t^* or pandemic fatigue sets in at t_{PF} .

Proposition 6 *With perfect vaccination and $t_{PF} > t^*$ in the continuous time vaccination game, if pandemic fatigue is anticipated, it results in an increase in vaccination rates and has no effect on expected utility in equilibrium.*

With perfect vaccination, pandemic fatigue is superficially regrettable, but actually reduces equilibrium infection rates because it induces an increase in vaccination that more than offset the reductions in social distancing corresponding to pandemic fatigue. Since the timing of vaccination (and myopic social distancing decisions prior to that) are unaffected by pandemic fatigue, the indifference condition for perfect vaccination at time t^* ensures that equilibrium expected payoffs are the same as in Proposition 3. The anticipation of pandemic fatigue in the future makes vaccination more attractive and so equilibrium vaccination rates must increase in response. In this sense, pandemic fatigue resembles a work-based vaccine mandate as each has the primary effect of reducing expected utility to those who remain unvaccinated after t^* , thereby promoting vaccination. Of course, this conclusion relies critically on the assumption that future pandemic fatigue is anticipated accurately at the time of vaccination.²³ As in the case of vaccine mandates, if there are unmodeled societal benefits to a reduction in peak infection rate and in the rate of increase of infections prior to that peak, pandemic fatigue may actually result in improved outcomes since it shifts some of the costs incurred in equilibrium from infections to more costly reductions in activities after t_{PF} .

²³By contrast, unanticipated pandemic fatigue would induce additional vaccinations at time t_{PF} , thereby reducing but not offsetting the effect of reduced social distancing on expected utility.

5.2 Forward-Looking / Economically Rational Social Distancing

With forward-looking social distancing, each susceptible person chooses a level of social distancing at each moment in time to minimize the expected discounted sum of future payoffs. Denoting V_t as the expected sum of discounted future costs at time t (in time- t utility units), the objective function at time t for susceptible person j when vaccination is not possible is

$$-P(x_t, I_t)D - C(x_t) + \delta(1 - P(x_t, I_t))V_{t+1} = -TA_I(1 - x_t)I_tD + C(x_t) + \delta[1 - TA_I(1 - x_t)I_t]V_{t+1}.$$

The first-order condition for the optimal level of social distancing x_t at time t is then

$$C'(x_t) = TA_I I_T (D - \delta V_{t+1}).$$

This first-order condition is distinct from the first-order condition for myopic maximization because of the inclusion of the term δV_{t+1} . That is, a forward-looking susceptible person chooses a level of social distancing based on the understanding that avoiding infection now saves D now but does not rule out the possibility of future infection and costs. Thus, forward-looking optimization results in a lower level of social distancing than does myopic optimization for any given infection rate (assuming the same social distancing cost function $C(x)$ in both cases).

The inclusion of the term V_{t+1} in the period t objective function brings a recursive element into the model, since expected payoffs at time t now depend on future actions at time $t + 1$ and beyond. Nevertheless, most of the results with myopic optimization carry over to the model with forward-looking behavior. A key first observation is that the gain from avoiding immediate infection is at least $(1 - \delta)D$,²⁴ which means that the first-order condition for x_t corresponds to positive marginal cost and thus $x_t^* > 0$ for all t . That is, forward-looking equilibrium behavior always includes some social distancing.

Figure 4 repeats Figure 2 with forward-looking rather than myopic social distancing.²⁵ There are higher infection rates for each value of T with forward-looking than with myopic social distancing but the general relationships of the curves are much the same as in Figure 2. As in Proposition 2, an increase in the effective transmission rate from TA_I has ambiguous effects on \bar{I} with forward-looking social distancing.

²⁴Since one possible strategy is to set $x_{jt} = 0$ in all periods, incurring no cost of social distancing, V_t must be at least $-D$.

²⁵The equilibrium trajectory rates in this figure are based on the results of iterative approximation of the value function V_{t+1} .




Figure 4. *Peak Infection Rates with Forward-Looking Social Distancing*

5.2.1 Perfect Vaccination with Forward-Looking Social Distancing

Forward-looking social distancing complicates analysis of the optimal timing of vaccination. Vaccination at time t reduces future infection rates, which in turn, alters the continuation payoff $V(t)$ for remaining unvaccinated, and thus alters social distancing choices and infection rates prior to time t . Despite this circularity, a one-step comparison of payoffs for vaccination in periods t and $t + 1$ yields a similar threshold rule for vaccination as the comparison described in Section 4 with myopic maximization.

In anticipation of vaccination next period, a forward-looking susceptible person at time t who plans to vaccinate in period $t + 1$ chooses x_t to minimize the discounted expected cost $P_t D + C(x_t) + (1 - P_t)\delta C_V$ in time- t units, where P_t is the probability of infection at time t after accounting for social distancing x_t . Perfect vaccination at time t yields cost C_V in time t units, so vaccination at time t is preferable to waiting until time $t + 1$ if $C_V \leq P_t D + C(x_t) + (1 - P_t)\delta C_V$; neither of these expected payoffs includes the continuation value V_{t+1} , which greatly simplifies analysis. The minimum expected cost for vaccinating next period, $P_t D + C(x_t) + (1 - P_t)\delta C_V$ is increasing in the infection rate I_t by revealed preference,²⁶ so $C_V \leq P_t D + C(x_t) + (1 - P_t)\delta C_V$ is equivalent to a threshold requirement $I_t \geq I_{FL}^*$. Thus, incentives for vaccination with forward-looking behavior match the cost-benefit structure from perfect vaccination with myopic social distancing; once again,

²⁶If $I_1 < I_2$, then choosing the level of social distancing x_2 that is optimal for infection rate I_2 (given the plan to vaccinate in the next period) at infection rate I_1 yields lower discounted expected cost than the minimum discounted expected cost for infection rate I_2 .

vaccination is cost effective over periods of time when the infection rate surpasses a fixed threshold.

There is a new technical challenge for analysis of forward-looking social distancing as the properties and even the existence of equilibrium are no longer obvious in this case. The complicating factor is the connection between the continuation value and current level of social distancing, as higher expected cost for $V(t+1)$ induces lower social distancing x_t , which in turn suggests higher expected cost next period. For expositional purposes, Propositions 7 and 8 assume that equilibrium with forward-looking social distancing exists for any initial level of disease and that it yields a single-peaked infection rate curve.²⁷

Proposition 7 *In any equilibrium with forward-looking social distancing, all vaccination in the continuous time game takes place at the first time t_{FL} where the infection rate surpasses threshold I_{FL}^* . In equilibrium, the infection rate continues to rise after vaccination takes place.*

Proposition 7 verifies that the key properties of the vaccination equilibrium of Proposition 3 are maintained with forward-looking social distancing: all vaccination takes place at a single moment in time and is not sufficient to cause an immediate decline in infection rate. As in Figure 3, the peak infection rate occurs after vaccination is complete. One difference between the results of Propositions 3 and 7 is that vaccination changes the initial trajectory of the infection rate with forward-looking behavior but has no effect on the infection rate prior to vaccination with myopic social distancing. For this reason, it is possible to identify the timing of vaccination with myopic social distancing from the equilibrium outcome when vaccination is not possible: $I(t)$ reaches I^* at the same time whether vaccination is possible or not. With forward-looking decisions, however, vaccination reduces future risks and increases incentives for prior social distancing (as emphasized by Makris and Toxvaerd (2021)), so the infection rate reaches threshold I_{FL}^* at a later time in equilibrium with vaccination than in equilibrium without vaccination. This difference between the models is primarily technical in nature – it increases the difficulty of computation of equilibrium strategies with forward-looking behavior without yielding especially substantive changes in observed vaccination behavior.

As in Proposition 4, a vaccine mandate with forward-looking social distancing promotes incentives for vaccination by penalizing the unvaccinated. With forward-looking distancing, the greater

²⁷Iterative approximation of the continuation value function $V(t)$ appears to work well for the case of quadratic costs $C(x) = \frac{\gamma x^2}{2}$ with resulting single-peaked trajectories for equilibrium infection rates as illustrated in Figure 4. Appendix B of Farboodi et al., (2020) provides details about the (minor) difficulties with the related procedure used to estimate equilibrium outcomes in that paper.

incentives for vaccination both alter the timing of vaccination and the payoffs for both vaccinated and unvaccinated. In particular, a work-based mandate delays the spread of infection and the timing of vaccination because the threshold infection rate I_{FL}^* is only achieved at a time t_{FLV} after t_{FL}^* when vaccination takes place with forward-looking behavior and no vaccine mandate.

Proposition 8 *With forward-looking social distancing and a work-based vaccine mandate, all vaccination takes place at a single time t_{FLV} where $t_{FLV} > t_{FL}^*$.*

Since the work-based mandate delays the timing of vaccination, it reduces equilibrium infection rates in the sense of delaying the time at which the infection rate reaches the threshold I_{FL}^* .

6 Imperfect Vaccination

Imperfect vaccination that provides less than complete immunity from infection is quite topical given the observed phenomenon of “breakthrough infections”. The phenomenon of imperfect vaccination opens new depths to equilibrium analysis because the activity levels of unvaccinated and vaccinated people each influences the transmission rate of infections. If vaccinated people choose a single level of social distancing x_{Vt} and non-vaccinated people choose a separate common level of social distancing x_{Nt} , then these choices (likely) serve as strategic substitutes in a collective action game with the twin goals of reducing the risk of infection at the societal level while minimizing one’s own cost of doing so. Needless to say, this is a much more complicated game than any of the versions of the model with perfect vaccination.

This section sketches of a model with imperfect vaccination and uses a numerical example to provide some suggestive results. Assume that that imperfect vaccination is available to everyone at cost C_V (homogeneous cost) and that it reduces the probability of infection in period t from $TA_I A_{jt} I_t$ to $\alpha TA_I A_{jt} I_t$ where α is a known constant less than 1 and $TA_I \alpha < 1$ so that a sufficient level of vaccination yields herd immunity.²⁸ Then myopic social distancing by all susceptible people yields a lower level of social distancing x_{Vt} for vaccinated people than for non-vaccinated people V_{Nt} as a function of the infection rate at time t .

With imperfect vaccination and no social distancing, the transition equation for infection rates at period $t+1$ includes the measures of both susceptible non-vaccinated S_{Nt} and susceptible vaccinated

²⁸While it might be natural for vaccinated and unvaccinated people to choose different levels of activity when infected, particularly in the case of asymptomatic infection, this section maintains the assumption that A_I is an exogenous fixed value and that it is the same for both unvaccinated and vaccinated people who are infected.

people S_{Vt} as well as the current infection rate.

$$I_{t+1} = TA_I I_t S_{Nt} + \alpha TA_I I_t S_{Vt}.$$

This equation indicates that herd immunity can no longer be summarized by the cumulative infection rate because each vaccination reduces the reproductive rate R_t of disease.

Imperfect vaccination also complicates equilibrium analysis because the expected payoff for vaccination includes the possibility of a future infection. Since future infection rates vary with current vaccination rates, there are positive externalities of vaccination for those who are vaccinated and not just for those who are unvaccinated (as is the case with perfect vaccination), which creates possibilities for multiple equilibria.

With imperfect vaccination and no social distancing, vaccinating in period t yields a higher expected payoff than waiting to vaccinate in period $t + 1$ if not infected in period t if $\delta^t C_V + \delta^t \alpha P_t D + \delta^{t+1} V_{t+1} \geq \delta^t P_t D + \delta^{t+1} (1 - P_t) C_V + \delta^{t+1} (1 - P_t) V_{t+1}$, where P_t is the probability of infection in period $TA_I I_t$. We can write this equation to identify a threshold probability for the benefits of immediate vaccination to outweigh the costs of waiting one period before vaccinating.

$$P_t \geq \frac{C_V(1-\delta)}{(1-\alpha)(D-\delta V_{t+1})-\delta C_V}.$$

Here V_{t+1} varies with time and need not be constant for a given infection rate: an increase in V_{t+1} reflects an increase in future expected cost, thereby reducing the value of the denominator and increasing the threshold probability value above. This result reflects the economic intuition that it is less desirable to choose imperfect vaccination when facing greater risk of infection in the future.

As in Proposition 3, vaccination can only take place in equilibrium when the cost-benefit condition holds with equality: $P_t = \frac{C_V(1-\delta)}{(1-\alpha)(D-\delta V_{t+1})-\delta C_V}$, or equivalently, $I_t = \frac{C_V(1-\delta)}{TA_I(1-\alpha)(D-\delta V_{t+1})-\delta C_V}$. If δ is close to 1, the continuation value V_{t+1} for those who are vaccinated changes much more slowly over time than does I_t , which guarantees that the cost-benefit condition holds with equality at most twice, once when the infection rate is increasing and once when it is decreasing. Under these conditions, equilibrium with imperfect vaccination takes the same general form as equilibrium with perfect vaccination, with all vaccination taking place at a single point in time.

Figure 5 compares the equilibrium outcomes with perfect and imperfect vaccination for the parameter values from Figure 3 and an imperfect vaccine that reduces transmission rates by 75%. With imperfect vaccination, the threshold infection rate for vaccination is initially increasing and subsequently is decreasing and is higher than the threshold with perfect vaccination. Thus, vaccination takes place at a later time with imperfect vaccination and the peak equilibrium infection

rate is higher with imperfect than with perfect vaccination. In this example, $R_t > 1$ results at the time of vaccination in equilibrium, since if $R_t < 1$, vaccination would not be cost effective over any period of time, following the same logic as in Corollary 1.²⁹ Vaccination is still quite effective at limiting the epidemic. With no possibility of vaccination, the peak infection rate is 9.5% and cumulative infections to 68.7% by comparison to peak infection rates of 1.14% and 1.72% and cumulative infections of 41.2% and 53.4% with perfect and imperfect vaccination respectively.




Figure 5. *Equilibrium Infection Rates with Perfect vs. Imperfect Vaccination*

Allowing for endogenous social distancing along with imperfect vaccination further complicates the threshold condition for vaccination now to be more cost effective than vaccination in one period. In equilibrium, vaccination still takes place when the cost-benefit condition holds with equality but there may be more than one time when that condition holds with equality. Though beyond the scope of this current paper, we believe that further investigation of the model with imperfect vaccination may be of interest, especially in extensions that allow for different patterns of social contact between vaccinated and unvaccinated people who remain susceptible and for different levels of social activity between vaccinated and unvaccinated people when infected.

²⁹One new feature of the equilibrium with imperfect vaccination for this example is that vaccination takes place at an earlier time in equilibrium than would be optimal if no one ever vaccinated; this feature results from the positive externalities of imperfect vaccination, which causes the infection rate to reach the cost-benefit threshold at an earlier time when some people vaccinate than when no one chooses to vaccinate.

7 Discussion and Conclusion

The paper develops a model with the minimal set of assumptions to study precautionary behavior in response to the spread of infectious disease. and identifies cost-benefit considerations that are central to individual decision-making in this context. The results emphasize the ways that an epidemic poses society with a collective action problem, so that some policy interventions crowd out existing incentives for private individuals to mitigate risk through social distancing. Vaccination is the only approach that reliably reduces the herd immunity level of cumulative disease and yet endogenous adoption of vaccination yields a reproductive rate of disease $R_t > 1$ at the time of vaccination; this finding reflects the insight of Geoffard and Philipson (see also Goodkin-Gold et al., (2021)) that increases in vaccination reduce the marginal value of vaccination to the point where private incentives for vaccination are not sufficient for disease eradication. Given this background, the interaction of economic forces in general equilibrium can yield seemingly counterintuitive results – for example, that the phenomenon of pandemic fatigue can lead to a Pareto improvement in equilibrium outcome because it promotes incentives for vaccination.

The minimalist design of the baseline model should facilitate further study of extensions that add detail to the interventions mentioned in the text or incorporate new features that have been prominent in practice. For instance, it seems natural to consider heterogeneous populations with different levels of risk and social connections, seasonality in the infectiousness of a disease, and the development of variants of an existing disease during the course of an epidemic. As indicated in Appendix Table 1, many of these ideas are already present in the literature but relatively few of them have been studied in combination with endogenous vaccination. It would also be natural to consider extensions with combinations of symptomatic and asymptomatic infection, and strategic interactions between unvaccinated and imperfectly vaccinated people where the social distancing decisions of each person have spillover effects for the immediate level of risk of infection to others.

References

- [1] Acemoglu, Daron, Victor Chernozhukov, Iván Werning and Michael D. Whinston (2020), “Optimal Targeted Lockdowns in a Multi-Group SIR Model,” NBER Working Paper 27102.
- [2] Alvarez, Fernando, David Argente, and Francisco Lippi (2021), “A Simple Planning Problem for COVID-19 Lock-down, Testing, and Tracing,” *American Economic Review: Insights*, 3:3, 367-382.
- [3] Atkeson, Andrew (2021a), “Behavior and Dynamic of Epidemics”, *Brookings Papers on Economic Activity*, Spring 2021.
- [4] Atkeson, Andrew (2021b), “A Parsimonious Behavioral SEIR Model of the 2020 COVID Epidemic in the United States and the United Kingdom”, UCLA Working Paper, 2021.
- [5] Atkeson, Andrew, Michael Droste, Michael Mina, and James H. Stock (2020a), “Economic Benefits of Covid-19 Screening Tests:”, NBER Working Paper 28031.
- [6] Atkeson, Andrew, Karen Kopecky, and Tao Zha (2020b), “Behavior and the Transmission of COVID-19”, UCLA Working Paper, 2020.
- [7] Auld, Christopher and Flavio Toxvaerd (2021), “ The Great COVID-19 Vaccine Rollout: Behavioral and Policy Response”, Discussion Paper, Centre for Economic Policy Research.
- [8] Baqaee, David, Emmanuel Farhi, Michael Mina, and James H. Stock (2020), “Policies for a Second Wave”, *Brookings Papers on Economic Activity*, Spring 2020, p. 385-431.
- [9] Blau, Francine D., Josefine Koebe, and Pamela A. Meyerhofer (2021). “Who are the essential and frontline workers?,” *Business Economics* 56(3), p. 168-178, 2021.
- [10] Budish, Eric (2020), “Maximize Utility subject to $R \leq 1$: A Simple Price-Theory Approach to Covid-19 Lockdown and Reopening Policy,” NBER Working Paper 28093
- [11] Chen, Frederick and Flavio Toxvaerd (2014), “The Economics of Vaccination,” *Journal of Theoretical Biology*, 363, p. 105-117.
- [12] Cochrane, John H. (2020), “An SIR Model with Behavior,” blog post (May 4, 2020); <https://johnhcochrane.blogspot.com/2020/05/an-sir-model-with-behavior.html>

- [13] Droste, Michael and James H. Stock (2021), “Adapting to the COVID-19 Pandemic,” *AEA Papers and Proceedings*, 111, 2021.
- [14] Eichenbaum Martin S., Sergio Rebelo, and Mathias Trebandt (2020), “The Macroeconomics of Epidemics,” NBER Working Paper 26882.
- [15] Ellison, Glenn, (2020), “Implications of Heterogeneous SIR Models for Analyses of COVID-19,” NBER Working Paper 27373.
- [16] Engle, Samuel, Jussi Keppo, Marianna Kudlyak, Elena Quercioli, Lones Smith, and Andrea Wilson (2021), “The Behavioral SIR Model with Applications to the Swine Flu and COVID-19 Pandemics,” Wisconsin University Working Paper, 2021.
- [17] Farboodi, Maryam, Gregor Jarosch, and Robert Shimer (2020), “Internal and External Effects of Social Distancing in a Pandemic”, NBER Working Paper 27059.
- [18] Galeotti, Andrea and Brian W. Rogers (2013) “Strategic Immunization and Group Structure” *American Economic Journal: Microeconomics*, 5(2), 1-32.
- [19] Gans, Joshua, S. (2021), ‘Vaccine Hesitancy, Passports, and the Demand for Vaccination,’ NBER Working Paper 29075.
- [20] Gans, Joshua S. (2020) “The Economic Consequences of $R = 1$: Towards a Workable Behavioural Epidemiological Model of Pandemics,” NBER Working Paper 27632.
- [21] Geoffard, Pierre-Yves and Tomas Philipson (1997) “Disease Eradication: Public vs. Private Vaccination,” *American Economic Review* 87, 222-230.
- [22] Goodkin-Gold, Matthew, Michael Kremer, Christopher M. Snyder, and Heidi Williams (2021), “Optimal Vaccine Subsidies for Endemic and Epidemic Diseases,” NBER Working Paper Diseases 28085.
- [23] Makris, Miltiadis (2020), “Covid and Social Distancing with a Heterogenous Population,” working paper, University of Kent.
- [24] Makris, Miltiadis and Flavio Toxvaerd (2020), “Great Expectations: Social Distancing in Anticipation of Pharmaceutical Innovations”, Cambridge-INET Working Paper Series No: 2046.

- [25] McAdams, David (2021), “The Blossoming of Economic Epidemiology,” *Annual Review of Economics*, forthcoming.
- [26] McAdams, David (2020), “Nash SIR: An Economic-Epidemiological Model of Strategic Behavior During a Viral Epidemic”, *Covid Economics*, Issue 16, p. 115-134.
- [27] McAdams, David, Yangbo Song, and Dihan Zou (2021), “Equilibrium Social Activity During an Epidemic,” Duke University Working Paper.
- [28] McCormack, Grace, Christopher Avery, Ariella Kahn-Lang Spitzer, and Amitabh Chandra, “Economic Vulnerability of Households With Essential Workers,” *JAMA*. 324(4):388-390, 2020.
- [29] Rachel, Lukasz, “An Analytical Model of Covid-19 Lockdowns”, London School of Economics Working Paper, 2020.
- [30] Reluga, Timothy C. (2010), “Game Theory of Social Distancing in Reponse to an Epidemic”, *PLOS Computational Biology*, 6(5).
- [31] Reluga, Timothy C. and Alison P. Galvani (2011), “A General Approach for Population Games with Application to Vaccination”, *Mathematical Biosciences* 230(2), p. 67-78.
- [32] Toxvaerd, Flavio (2020a), “Equilibrium Social Distancing”, Cambridge-INET Working Paper Series No. 2008.
- [33] Toxvaerd, Flavio and Robert Rowthorn (2020b), “On the Management of Population Immunity,” Cambridge-INET Working Paper Series No. 2020.
- [34] Toxvaerd, Flavio (2021), “Contacts, Altruism, and Competing Externalities,” Cambridge-INET Working Paper Series No. 2135.
- [35] Weitz, Joshua S., Sang Woo Park, Ceyhun Eksin, and Jonathan Dushoff, “Awareness-driven behavior changes can shift the shape of epidemics away from peaks and toward plateaus, shoulders, and oscillations,” *Proceedings of the National Academy of Sciences* 117 (51) 32764-32777.

A Classification of Related Papers

Reference	Social Distancing	Vaccination	Calibration	Interventions	Other
This Paper	M	I, P		V_M	
Acemoglu (2020)	M			L, T	G
Atkeson (2021a,b)	M_D		Yes		PF, S
Atkeson (2020a)	M			T	G
Atkeson (2020b)	M_D		Yes		PF
Baqae (2020)	M		Yes,	Various	
Chen (2014)		I, P			
Droste (2021)	M_D		Yes		S
Eichenbaum (2020)	F		Yes	L	
Engle (2021)	M		Yes		
Farboodi (2020)	F		Yes	L	
Galeotti (2013)		P		V_S	G
Gans (2021)	F_L	P		V_M	
Geoffard (1997)		P			
Goodkin-Gold (2021)		I, P	Yes	V_S	P
Makris (2020)	F		Yes	L	G
Makris (2020b)	F	P		S	
McAdams (2021)	F_L				
Rachel (2020)	F_L			L	
Reluga (2010)	F				
Toxvaerd (2020a)	M_L				
Toxvaerd (2020b)		I			T
Weitz (2020)	M_D				PF

Table 2. Summary of Papers with Strategic Social Distancing and Vaccination

Social Distancing Codes: F = Forward Looking; F_L = Forward Looking with Linear Preferences; M = Myopic; M_D = Myopic based on Deaths.

Vaccination Codes: I = Imperfect; P = Perfect.

Intervention Codes: L = Lockdown; Te = Tests; V_M = Vaccine Mandate; V_S = Vaccine Subsidy.

Other Codes: G = Groups; P = Producers; PF = Fatigue; S = Seasonality; T = Treatment.

B Proofs

Proof of Proposition 1

(a) Person j 's objective function for period t is $\pi(x_{jt}, I_t) = -TA_I A_{St} I_t D - C(x_{jt})$, which gives first-order condition $C'(x_{jt}) = TA_I I_t D$. Here, the marginal value of social distancing (per unit of activity) is constant and equal to $TA_I I_t D$, since activity increases both social interactions and the probability of contracting an infection at a constant rate. Since C is strictly convex with $C'(0) = 0$, this first-order condition either holds for a unique optimal level of social distancing $x^*(I_t)$ for each I_t or else $C'(x) < TA_I I_t D$ for all $x < 1$ in which case there is a boundary solution with $x^*(I_t) = 1$. Since the marginal value of social distancing is increasing in I_t , $x^*(I_t)$ is strictly increasing and A_{St} is strictly decreasing in I_t , once again because C is strictly convex.

(b) Define G_t as the cumulative proportion of the population that has been infected by period t . By the Monotone Convergence Theorem, G_t must converge to a limit since it is increasing and bounded on the range $[0, 1]$. Suppose that this limit is some value $L < H$. By definition $(1 - H)TA_I = 1$, so $(1 - L)TA_I > 1$. Working backwards from the optimal social distancing / activity level, first choose a constant κ that is sufficiently small that $(1 - L)TA_I(1 - \kappa) > 1$. Person j chooses a social distancing level less than κ and thus an activity level greater than $1 - \kappa$ whenever $C'(\kappa) > TA_I D(1 - \kappa)$ or equivalently, $I_t < \frac{C'(\kappa)}{TA_I D}$. Next choose $\epsilon = \frac{C'(\kappa)}{TA_I D}$. Since $G_t = \sum_{n=0}^t I_n$ and G takes a limiting value of L , there is some period N such that $H - S_t < 1 - \epsilon$ for each $t > N$, implying $I_t < \epsilon$ for each $t > N$ as well. Then the reproductive rate of disease $R_t = S_t TA_I A_{jt} > (1 - L)TA_I(1 - \kappa) > 1$ for each $t > T$ (with $G_t \leq L$) which contradicts the assumption that G_t converges to L .

(c) The steady state equation (not adjusted for the size of the susceptible population) is $TA_I A_{St} I_t = I_t$ or equivalently $TA_I A_{St} = 1$. The steady state corresponds to the infection level such that each susceptible individual chooses activity level $\frac{1}{TA_I}$, meaning $x = 1 - \frac{1}{TA_I} = \frac{TA_I - 1}{TA_I}$. Since x is determined by the (single-period) first-order condition $C'(x) = TA_I I_t D$, the steady state is characterized by the equation $C'(\frac{TA_I - 1}{TA_I}) = TA_I I_t D$. If $C'(\frac{TA_I - 1}{TA_I}) < TA_I D$, then there is some value $I < 1$ such that $C'(\frac{TA_I - 1}{TA_I}) = TA_I I_t D$ holds with equality. Then, solving for \bar{I} given $I_t = \bar{I}$ yields the identity $\bar{I} = \frac{C^{-1}(\frac{TA_I - 1}{TA_I})}{TA_I D}$.

(d) In the continuous time version of the model, $I(t)$ is initially increasing and so $S(t)$ is decreasing at rate at least $-I(0)$ while $I(t)$ is increasing. Since $I(t)$ must decrease once herd immunity is reached at $S(t) = H$, there must be some time \bar{t} such that $I'(\bar{t}) = 0$. $I(t)$ cannot

remain constant beyond \bar{t} at $I(\bar{t})$ because $S(t)$ and thus $I'(t)$ would decrease if $I(t)$ stayed at the peak value past \bar{t} . That is, $I(t)$ is declining just after time \bar{t} .

Next observe that $I(t)$ achieves each value between $I(0)$ and $I(\bar{t})$ twice, once when $I(t)$ is increasing before time \bar{t} and once when $I(t)$ is decreasing after \bar{t} . If not, then there must be some value $I_K < I(\bar{t})$ such that $I(t_1) = I(t_2) = I_K$, $t_2 > t_1 > \bar{t}$, and $I''(t_1) < 0$ and $I''(t_2) > 0$. But then since $S(t_2) < S(t_1)$ by definition of $S(t)$, $I'(t_2)$ must be less than $I'(t_1)$ and so cannot be greater than 0.

Proof of Proposition 2: Section C provides computations that encompass this result.

Proof of Proposition 3: Vaccinating at time t (conditional on no prior infection) yields a certain loss of $\delta^t C_V$. Vaccinating at time $t + 1$ (again conditional on no infection prior to time t) yields either a loss of $\delta^t D$ for contracting an infection at time t or a loss of $\delta^{t+1} C_V$ for not contracting an infection at time t and then choosing to be vaccinated in the next period. Denoting the probability of infection at time t as P_t , it is preferable to vaccinate immediately in period t than to wait one period if

$$\delta^t C_V \leq \delta^t P_t D + \delta^{t+1} (1 - P_t) C_V \text{ or } P_t \geq \frac{C_V(1 - \delta)}{D - \delta C_V}.$$

Since $P_t = T A_I I_t$ with no social distancing, this condition can be reorganized as a restriction on I_t , namely $I_t \geq \frac{C_V(1 - \delta)}{T A_I (D - \delta C_V)}$ or $I_t \geq I^*$ where $I^* = \frac{C_V(1 - \delta)}{T A_I (D - \delta C_V)}$.

With social distancing, the cost-benefit inequality for a perfect vaccine in period t to yield higher expected payoff than waiting until period $t + 1$ to vaccinate if

$$C_V \leq P_t D + (1 - P_t) \delta + C(x_t).$$

In this equation, P_t is the probability of infection after accounting for social distancing and $C(x_t)$ is the cost of social distancing given myopic choice of x_t in period t . Reorganizing this inequality with all terms that vary with the infection rate on the left-hand side gives $P_t(D - \delta C_V) + C(x_t) \geq C_V(1 - \delta)$. We know that $C(x_t)$ is increasing in I_t but it is possible, as suggested by Figure 1, that P_t is not increasing for all ranges of I_t after accounting for the effect of social distancing. At infection rates such that P_t is increasing in I_t , both terms on the left-hand side of the inequality are increasing in P_t . At infection rates such that P_t is decreasing in I_t , we can rewrite the left-hand term as $P_t D + C(x_t) - P_t \delta C_V$, where by assumption $-P_t \delta C_V$ is increasing in t . Endogenous social distancing is chosen to minimize immediate expected cost to an individual $P_t D + C(x_t)$ where (by revealed preference), this optimized value must also be increasing in I_t . Thus, whether P_t is

increasing or decreasing in I_t , the left-hand side of the inequality $P_t(S - \delta C_V) + C(x_t)$ is increasing in I_t and so $P_t(D - \delta C_V) + C(x_t) \geq C_V(1 - \delta)$ holds for $I_t \geq I^*$ where I^* is defined implicitly so that $P_t(D - \delta C_V) + C(x_t) = C_V(1 - \delta)$.

This comparison indicates that vaccination only takes place in equilibrium in the first of some number of periods with $I_t \geq I^*$. With no social distancing, the reproductive rate of infection at time t is $R_t = TA_I S_t$, where S_t is the proportion of the population that is susceptible at time t . Since S_t declines over time as previously susceptible people choose to be vaccinated or become infected, R_t declines over time as well. That is, the infection is increasing for some period of time when $R_t > 1$ and then subsequently decreasing after that since $R_t < 1$ implies $R_n < 1$ for all $n > t$. As a result, there can be at most one period t^* where $I_{t-1} < I^*$ and $I_t \geq I^*$, i.e. where the infection rate “crosses the threshold” for possible vaccination “from below” that threshold.

In equilibrium, vaccination occurs only at time t^* and the remaining question is about the amount of vaccination (if any) at that time.

Lemma 1 *Each increase in vaccination at time t^* reduces cumulative infection at each subsequent time t .*

Vaccination at time t^* provides a one-time reduction in the susceptible population but does not change the dynamics of infection, with an initial period of increase followed by subsequent decrease in the infection rate. Each increase in vaccination reduces the immediate size of the susceptible population and so reduces the realized infection rate just after time t . Consider two levels of vaccination V_1 and $V_2 > V_1$ at time t^* . Denote the current infection rates for the two regimes as $I_{V_1}(t)$ and $I_{V_2}(t)$, the infection rates as $I_{V_1}(t)$ and $I_{V_2}(t)$, the cumulative rates of infection as $G_{V_1}(t)$ and $G_{V_2}(t)$, and the remaining susceptible populations as $S_{V_1}(t)$ and $S_{V_2}(t)$ respectively. By construction, $G_{V_1}(t^*) = G_{V_2}(t^*)$ and $I_{V_1}(t^*) = I_{V_2}(t^*)$ while by definition $V_1 + S_{V_1}(t) + G_{V_1}(t) = V_2 + S_{V_2}(t) + G_{V_2}(t) = 1$ for $t \geq t^*$.

Claim: The peak infection level $\overline{x_{V_1}}$ for regime V_1 is at least as large as the peak infection level $\overline{x_{V_2}}$ for regime V_2 .

If $I'_{V_2}(t) \leq 0$ for t then the peak infection rate for Regime V_2 occurred previously and the result is immediate.³⁰

³⁰Here $I'_{V_2}(t^*)$ refers to the right-hand derivative of the infection rate at time t^* , taking account of vaccination at that time. A subsequent result, Corollary 1 demonstrates that vaccination takes place in equilibrium prior to the peak of the infection rate, so this case does not occur in equilibrium.

Next consider the remaining case where $I'_{V_2}(t^*) > 0$. By assumption, vaccination leaves a larger susceptible population in Regime V_1 than in Regime V_2 and so $I'_{V_1}(t^*) > I'_{V_2}(t^*)$. Then there must be a range of infection levels $y \in (I_{V_2}(t^*), \bar{y}]$ such that for each y : (1) infection level y is first reached at an earlier time in Regime V_1 than in Regime V_2 ; (2) the susceptible population when infection level y is greater in Regime V_1 than in Regime V_2 . Denote the times when infection rate y is first reached in the two regimes as t_{1y} and t_{2y} respectively. Property (2) implies that $I'_{V_1}(t_{1y}) > I'_{V_2}(t_{2y})$ for each $y \in (I_{V_2}(t^*), \bar{y}]$. By construction, the infection rate curve is steeper for Regime V_1 than for Regime V_2 on this range of infection rates, and so $I_{V_2}(t + t_{2y} - t_{1y}) > I_{V_1}(t)$ for $t^* < t < t_{1y}$. That is, comparing the infection rates over equal length time periods from t^* to t_{1y} for Regime 1 and from $t^* + t_{2y} - t_{1y}$ to t_{2y} for Regime 2, the infection rate is lower at each time t in Regime 1 than at the corresponding time $t + t_{2y} - t_{1y}$ in Regime 2. Therefore, there must be more recoveries from infection in Regime 2 during time period $[t^* + t_{2y} - t_{1y}, t_{2y}]$ than in Regime 1 during time period $[t^*, t_{1y}]$ and by extension, more recoveries from infection from Regime 2 during time period $[t^*, t_{2y}]$ than in Regime 1 during time period $[t^*, t_{1y}]$. By construction, $I_{V_1}(t^*) = I_{V_2}(t^*)$ and $I_{V_1}(t_{1y}) = I_{V_2}(t_{2y})$, so since there are more recoveries from infection in Regime 2 during $[t^*, t_{2y}]$ than in Regime 1 $[t^*, t_{1y}]$, there also must be a greater number of cumulative infections in Regime 2 than in Regime 1 during these corresponding time periods for the two regimes. That is, the difference in susceptible populations starts at $V_2 - V_1$ at $I_{V_1}(t^*) = I_{V_2}(t^*)$ and is diverging rather than converging at higher infection rates. Thus, the slope of the infection rate curve is larger in Regime V_1 than in Regime V_2 for each new infection rate while the infection rate is increasing in each regime. As a result, the infection rate must still be increasing in Regime V_1 when it peaks in Regime V_2 .

End Proof of Claim

Since $I_{V_1}(t) > I_{V_2}(t)$ and $G_{V_1}(t) > G_{V_2}(t)$ for t just greater than t^* , $G_{V_1}(t) > G_{V_2}(t)$ is only possible if the infection rate curves cross at some time t^* and I_1 then lies above I_2 for some period of time $[t^*, t^{**}]$. By the logic from the proof of the claim above, I_2 can only cross and then lie above I_1 to the right of the peak infection point for I_1 . Suppose that there exists time t^* such that $I_1(t^*) = I_2(t^*)$ and $I_2'(t^*) > I_1'(t^*)$ so that infection rate I_2 lies above I_1 just after time t^* . Then $S_2(t^*) > S_1(t^*)$, which means that the cumulative infection rates differ by more than $V_1 - V_2$ at that point. For each subsequent time t such that $I_2(t) > I_1(t)$, there is a previous (most recent) time $t^-(I_2(t))$ such that $I_1(t^-(I_2(t))) = I_2(t)$. If $S_2(t^-(I_2(t))) > S_1(t)$ for each such t , then the cumulative infection rate for Regime 2 never catches up to the cumulative infection rate for Regime

1. If, instead, there is some t such that $S_2(t^-(I_2(t))) = S_1(t)$, then (as above) since the susceptible populations and infection rates are equal across the two regimes at infection rate $I_2(t)$, the future course of infection in the two regimes must be identical from then on (from different starting times), and the difference in cumulative infection rates at these pairs of times will be $V_2 - V_1$ and once again, the cumulative infection rate for Regime 2 never catches up to the cumulative infection rate for Regime 1.

End Proof of Lemma

Given the Lemma, each increase in vaccination rate at time t^* reduces cumulative infection rates in all subsequent periods. Since vaccination has no effects on prior infection rates, each increase in vaccination rate at time t^* thus reduces the expected gain from vaccination at t^* . Thus, there is either a boundary solution where no one chooses to vaccinate at time t^* (or at any other time) or a mixed strategy equilibrium where susceptible individuals are indifferent between vaccination at time t^* and never vaccinating.

Proof of Corollary 1

If $I'(t^*) < 0$, then by Proposition 1, the infection rate is decreasing at each $t > t^*$ and $I(t) < I^*$ for $t > t^*$. Then it is not cost effective to be vaccinated over any time interval, which contradicts the presumption that a positive measure choose vaccination at time t^* .

Proof of Proposition 4

Assume throughout that a positive measure of people choose vaccination at time t^* in equilibrium with no vaccine mandate. (Footnote 15 describes the result for the case where vaccination is never cost effective without the mandate.) The introduction of a penalty for unvaccinated people from time t_M on yields a new threshold infection rate I_M for vaccination to be cost-effective over short-term periods from time t_M onwards, but does not change the relative payoffs for vaccination at time t^* (or any other time $t < t_M$) by comparison to vaccination at time t_M : the expected payoff for vaccination at time $t \leq t_M$ depends only on the number of people who choose to be vaccinated prior to time t with or without the imposition of the mandate. Thus, given that all vaccination took place at t^* without the mandate, that must continue to be true with the mandate.

In equilibrium without a work-based mandate, susceptible people are indifferent between vaccinating and not vaccinating at t^* , which requires a balance of expected gains from vaccination for $[t^*, t_M)$, when infection rates are relatively high and expected losses from vaccination at t_M and beyond, when infection rates are relatively low. If a positive measure of people vaccinate at time t_M , the expected payoff for vaccination relative to non-vaccination from t_M onward must be even

higher than it was with no mandate. If instead, no susceptible people choose to vaccinate at time t_M , the penalty for unvaccinated people still increases the expected payoff from vaccination from time t_M onward than it was with no mandate. In either case, more people choose to vaccinate in equilibrium at time t^* than would vaccinate without the work-based mandate.

Proof of Proposition 5

In any equilibrium where a positive measure of people choose to vaccinate, it must be those with the lowest costs of vaccination who choose to do so. To see this, note that if a person with cost C_1 chooses to vaccinate while a person with cost $C_2 < C_1$ chooses not to do so, then either the person with cost C_2 would increase expected payoff by vaccinating at the same time as the person with cost C_1 or the person with cost C_1 would increase expected payoff by choosing not to vaccinate. By the logic of Proposition 3, a person with given cost should only vaccinate when the infection rate reaches the threshold corresponding to that cost. Given these observations, equilibrium requires a set of lowest-cost types to vaccinate in turn when the infection rate reaches the relevant threshold and then vaccination when the marginal-cost type is indifferent between vaccinating and not vaccinating. Each type in ascending order of cost achieves lower expected gain from vaccination in equilibrium (because types with higher costs of vaccination receive lower expected gains from vaccination for each time period after vaccination and because they are not vaccinated during periods of time when lower-type costs increase their total gains from vaccination). If the lowest-cost type prefers to vaccinate at the appropriate time, there must be a higher-cost type that is exactly indifferent between vaccinating and not vaccinated when the infection rate reaches the relevant threshold (since if all previous types plan to vaccinate, the highest-cost type would gain by not vaccinating). By construction, there is a unique equilibrium and it takes the anticipated form.

The effect of a work-based mandate with heterogeneous costs is similar to the effect of that mandate with homogeneous costs of vaccination. One again, this mandate increases the value of vaccination beyond t_M without affecting social distancing choices or future infection rates conditional on vaccination by a given set of lowest-cost types. Thus, a mandate with heterogeneous costs of vaccination induces a new set of next-highest cost types to choose vaccination. In a new equilibrium, some of those types may vaccinate when infection reaches the relevant threshold prior to the time that the mandate becomes active and others may choose to vaccinate precisely at t_M . The equilibrium indifference condition only applies to the highest-cost type choosing to vaccinate as all lower cost types strictly prefer to follow their equilibrium strategies for vaccination than to remain

unvaccinated forever; by similar logic to the proof of Proposition, this type achieves lower expected utility in the mandate equilibrium than in the equilibrium without a mandate, so the equilibrium indifference condition indicates that the highest-cost types who never vaccinate in equilibrium with or without a mandate achieve strictly lower expected utility with the mask mandate than without it. Similarly, all other types who are induced to vaccinate by the mask mandate achieve lower expected utility with the mandate than without it.

Proof of Proposition 6 The result follows almost directly from the following Lemma.

Lemma 2 *If there is no possibility of vaccination, then pandemic fatigue results in an increase in cumulative infections for each time $t > t_{PF}$ by comparison to cumulative infections in equilibrium with the same cost function $C(x)$ for social distancing at each time t .*

First observe that by comparison of the first-order conditions in the two cases, there is less social distancing in equilibrium at each infection level with $C_{PF}(x)$ than with $C(x)$.

Step 1: If t_{PF} occurs at a point when the infection rate is rising, then pandemic fatigue results in an increase in the peak infection rate. Define $I_{PF}(t)$ as the equilibrium infection rate in the case of pandemic fatigue. By assumption $I(t) = I_{PF}(t)$ for $t \leq t_{PF}$ since the cost function is the same in both cases before the onset of pandemic fatigue and then $I'_{PF}(t_{PF}) > I'(t_{PF})$ since there is less social distancing at a given infection rate with pandemic fatigue than without it. By a similar argument to that above, if $I_{PF}(t_1) = I(t_2) > I(t_{PF})$, there are greater cumulative infections at time t_2 without pandemic fatigue than there are at time t_1 with pandemic fatigue. Thus, there are more susceptible people remaining and less social distancing with pandemic fatigue than without it, so $I_{PF}(t_1) > I(t_2)$, implying a greater peak infection rate with pandemic fatigue than without it.

Step 2: For each value of remaining susceptible people that occurs after the peak level of infection without pandemic fatigue, the corresponding infection level is higher with pandemic fatigue than without it.

If t_{PF} occurs when the infection rate is rising, then since the peak infection level of infection is higher with pandemic fatigue than without it, I_{PF} equals I^* , the peak value of $I(t)$, twice – once before and once after the peak infection rate with pandemic fatigue. Denote \bar{t} as the time when the peak infection rate is reached without pandemic fatigue and t_{PF}^{**} as the second time that rate is reached with pandemic fatigue. By assumption, $I_{PF}(t_{PF}^{**}) = I^*$ and $I'(t^*) = 0 > I'_{PF}(t^{**})$ since t_{PF}^* is the second time that this infection rate is reached with pandemic fatigue, at which point, the infection rate is declining. Since there is less social distancing at any infection rate with pandemic

fatigue than without it, so $I'(t^*) > I'_{PF}(t^{**})$ is only possible if there are more susceptible people at time \bar{t} without pandemic fatigue than at t_{PF}^{**} with pandemic fatigue. Further, for each $t > \bar{t}$, there is also a higher infection rate corresponding to that level of cumulative infection with pandemic fatigue than without it. If not, there must be some pair of times t_A and t_B such that there are an equal number of susceptible people at t_A without pandemic fatigue and at t_B with pandemic fatigue and also an equal infection rates without pandemic fatigue at t_A and with pandemic fatigue at t_B . But then there is more social distancing without pandemic fatigue at the given infection rate, the infection rates and susceptible populations must diverge once again in the two cases.

If t_{PF} occurs when the infection rate is falling, then once again $I(t) = I_{PF}(t)$ for $t \leq t_{PF}$ and $I'_{PF}(t_{PF}) > I'(t_{PF})$. By similar logic to the case where $t_{PF} < t^*$, for any measure of susceptible people reached without pandemic fatigue, the infection rate must be higher when that measure of susceptible people is reached with pandemic fatigue.

Step 3: The cumulative infection rate is higher with pandemic fatigue than without it at each time t .

At time t_{PF} the infection rate curves diverge with higher infection rate with pandemic fatigue than without it. Cumulative infections are clearly higher with pandemic fatigue than without it throughout any time interval $[t_{PF}, t_H]$ such that $I_{PF}(t) \geq I(t)$ given $t_{PF} \leq t \leq t_H$. If there is some later time t such that $I(t) > I_{PF}(t)$, then by the result of Step 2, the number of remaining susceptible people must be lower with pandemic fatigue than without it.

End Proof of Lemma 2

Lemma 2 indicates that pandemic fatigue reduces utility for unvaccinated people in equilibrium from time t_{PF} on. Since $t_M > t^*$ by assumption, $I(t)$ is the same with and without pandemic fatigue until the moment of vaccination when $I(t) = I^*$ at t^* . For a given level of vaccination, the future expected payoff for unvaccinated people is less with pandemic fatigue than without it.³¹ In equilibrium from Proposition 3, the marginal person is indifferent between vaccination at t^* and remaining unvaccinated forever, so with this level of vaccination and pandemic fatigue anticipated at future time t_M , this same marginal person strictly prefers vaccination. Therefore, there must be more vaccination at time t^* with pandemic fatigue than without it.

³¹By the same logic of Lemma 2, pandemic fatigue reduces utility for unvaccinated people by comparison to the result without pandemic fatigue assuming a fixed amount of vaccination at time t^* with and without pandemic fatigue.

Proof of Proposition 7

Let \underline{t}_{FL} denote the first moment when the infection rate reaches threshold I_{FL}^* in equilibrium when vaccination is not possible. Let $V^*(t)$ denote the measure of remaining susceptible people who choose to vaccinate at time t such that $I(t) = I_{FL}^*$ where $V^*(t)$ is continuous and well-defined on an interval $[\underline{t}_{FL}, \overline{t}_{FL}]$ where $V^*(\underline{t}_{FL}) = 0$ and $V^*(\overline{t}_{FL})$ corresponds to full vaccination at time \overline{t}_{FL} (i.e. $V^*(\overline{t}_{FL})$ is a fixed value chosen in advance to equal the measure of all susceptible people who remain at \overline{t}_{FL} given anticipation of that level of vaccination at that time). Thus, we have constructed a range of potential equilibrium times with exactly the amount of vaccination at given time t so that the infection rate reaches the threshold I_{FL}^* at that time. There must be a profitable deviation to vaccination at \underline{t}_{FL} from a proposed equilibrium with vaccination $V^*(\underline{t}_{FL})$ since by assumption there is no equilibrium without vaccination. If all vaccination takes place at time \overline{t}_{FL} and all remaining susceptible people vaccinate at that time, then the infection rate must be lower than I_{FL}^* for all $t > \overline{t}_{FL}$ and so it is preferable not to vaccinate than to vaccinate at time \overline{t}_{FL} when everyone else follows these strategies. The payoffs from vaccination at given time and remaining unvaccinated are continuous in time t for this range of proposed equilibria. Since it is preferable to vaccinate in the proposed equilibrium with vaccination at \underline{t}_{FL} and it is preferable not to vaccinate in the proposed equilibrium with vaccination at \overline{t}_{FL} , there must be some intermediate time t_{FL}^* such that the expected payoffs from vaccination and remaining unvaccinated are equal at time t_{FL}^* , so this is an equilibrium.

Proof of Proposition 8

Following the same construction as for Proposition 7, maintain the same range of proposed equilibria, while applying a fixed reduction in the expected payoff for not vaccinating. Vaccination remains preferable to remaining unvaccinated in the proposed equilibrium with vaccination at time \underline{t}_{FL} . If it remains preferable to remain unvaccinated in the proposed equilibrium with vaccination at time \overline{t}_{FL} , then there is still a time where the payoffs from both strategies are equal, but this time must be later than t_{FL}^* since by construction with the reduction in payoffs from the mandate, vaccination is preferable to remaining unvaccinated in all proposed equilibria with vaccination at time t_{FL}^* or before. Otherwise, if it is now preferable for everyone to vaccinate than remain unvaccinated in the proposed equilibrium with vaccination at last possible time \overline{t}_{FL} , then there is an equilibrium where everyone waits to vaccinate at time \overline{t}_{FL} and all remaining people vaccinate at that time. In each of these cases vaccination is delayed by the mandate to a time subsequent to the time of vaccination in equilibrium with forward-looking behavior and no mandate.

C Comparative Statics of \bar{I} with a Polynomial Cost Functions

With quadratic cost function $C(x) = \gamma x^2/2$, the first-order condition for social distancing is $C'(x) = I_t T A_I D$ with single period solution for susceptible person j satisfying $\gamma x_{jt} = I_t T A_I D$, meaning $x_j = \frac{I_t T A_I D}{\gamma}$ and associated activity level $A_{jt} = 1 - x_{jt}^* = 1 - \frac{I T A_I D}{\gamma}$. Substituting this value into the steady state equation gives $1 - \frac{I T A_I D}{\gamma} = \frac{1}{T A_I}$, with associated solution (after cross-multiplying by $\gamma T A_I$ or $\bar{I} = \frac{\gamma(A_I T - 1)}{A_I^2 T^2 D}$.

We can differentiate this closed form solution to determine the comparative static effect of any parameter on \bar{I} . Since D appears only in the denominator, \bar{I} is decreasing in D . Next, $dI^*/d(TA_I) = \frac{\gamma}{S} d(\frac{y-1}{y^2}/dy)$ using the change of variables $y = T A_I$ to simplify notation. Using the quotient rule, $d(\frac{y-1}{y^2}/dy) = \frac{y^2 - 2y^2 + 2y}{y^4}$, where the sign of this derivative matches the sign of the numerator, which is positive if $2y - y^2 > 0$ or $y > 2$. That is, the steady state \bar{I} is increasing in $T A_I$ if $T A_I < 2$ and is decreasing in $T A_I$ if $T A_I > 2$.

The same result holds for the more general case of polynomial costs with $C(x) = \gamma x^{m+1}/(m+1)$ so that $C'(x) = x^m$. in this case, the first-order condition for social distancing is $x^m = I T A_I D$ with solution $x = (I T A_I D)^{1/m}$. As for the case of quadratic costs, substituting this value into the steady state identify $A_S = \frac{1}{T A_I}$ gives $1 - (\frac{I T A_I D}{\gamma})^{1/m} = \frac{1}{T A_I}$ or $(\frac{T A_I I D}{\gamma})^{1/m} = \frac{T A_I - 1}{T A_I}$. Raising both sides to the n th power and solving for I gives $\frac{T A_I I D}{\gamma} = (\frac{T A_I - 1}{T A_I})^n$, which is equivalent to

$$I^* = \frac{\gamma(T A_I - 1)^m}{D(T A_I)^{m+1}}.$$

The formula for the steady state I^* takes the form $f(y) = K \frac{(y-1)^m}{y^{m+1}}$, where $y = T A_I$ and $K = \frac{\gamma}{D}$. Again using the quotient rule for differentiation, $\frac{df}{dy} = \frac{n(y-1)^m - 1y^{m+1} - (n+1)y^n(y-1)^m}{y^{2m+2}}$, which has the same sign as $ny - (n+1)(y-1) = m+1 - y$. That is, I^* is increasing in $y = T A_I$ for $T A_I < n$ and is decreasing in I^* for $T A_I > m$.

D Relationship between Discrete and Continuous Time Models

Suppose that we divide each discrete time period of unit length into N equal subperiods of length $\frac{1}{N}$ each. The adjusted discount factor δ_N per subperiod satisfies $\delta_N^N = \delta$ so $\delta_N = \delta^{\frac{1}{N}}$. Assuming infection probabilities and cost are proportional to period length, the cost of social distancing level x_n in subperiod n is $\frac{1}{N}C(x_n)$ and the probability of infection in subperiod N is $\frac{1}{N}TA_I(1 - x_n)I_n$, where I_n is the infection rate at the start of subperiod n . With myopic maximization, the objective function for Person j in subperiod n is $-TA_I(1 - x_n)I_nS - \frac{1}{N}C(x_n)$, which yields the identical first-order condition to the objective function for the same level of infection in the original discrete time model with periods of length 1. To maintain the property that the recoveries are proportional to the current infection rate, we relax the assumption of a deterministic length of time per infection and instead assume that each infected person recovers in the next subperiod with probability $\frac{1}{N}$ so that the average length of infection is still one (original) period. In addition to this change from deterministic to stochastic length infections, the continuous time model allows for the possibility of compounding in infections, similar to the phenomenon of compound interest, with the result that the infection rate increases at a faster rate per period without social distancing in the continuous time model than in the discrete time model even though $R_0 = TA_I$ is the same in both cases.

Threshold for Perfect Vaccination:

The following calculation identifies the threshold probability P^* and associated infection threshold I^{**} for perfect vaccination in the continuous time model without social distancing. Simplifying the one period discrete time inequality $P_t(D - \delta C_V) \geq (1 - \delta)C_V$ gives $P_t \geq \frac{C_V(1-\delta)}{D-\delta C_V}$, so $P^* = \frac{C_V(1-\delta)}{D-\delta C_V}$ and $I^* = \frac{C_V(1-\delta)}{TA_I D - \delta C_V}$. The continuous time thresholds correspond to the limiting values for P^* and I^* with each original length period divided into N subperiods: $P^{**} = \lim_{N \rightarrow \infty} \frac{NC_V(1-\delta^{1/N})}{D-C\delta^{1/N}}$ with solution $P^{**} = \frac{C_V \ln(\delta)}{D-C_V}$ so that $I^{**} = \frac{C_V \ln(\delta)}{TA_I(D-C_V)}$.

E Social Distancing with Linear Costs

With linear cost of social distancing, $C(x) = Kx$, where K is a positive constant while the benefit of social distancing through reduced risk of infection is $B(x) = TA_I I_t D x$. This linear cost structure also results when each person has a binary choice between (1) ordinary activity / no social distancing; (2) specific pre-specified reduction in activity with social distancing. With binary choice of social distancing actions, a mixed strategy yields proportional reductions in benefits and costs and thus linear marginal benefits and linear marginal costs of social distancing.

A key property of this case is that the marginal benefit of social distancing $TA_I I_t D$ and the marginal cost K of social distancing do not depend on x . In addition, marginal cost does not vary with time while I_t is the only time-varying parameter in the marginal benefit function. Starting from a low level of initial infection, there is no social distancing in equilibrium until the infection rate grows to exactly the point where marginal benefit equals marginal cost of social distancing: $I_t = I_{MBC} = \frac{K}{TA_I D}$, then remains at that plateau for some time. While $I_t = I_{MBC}$, social distancing by individuals x_{jt} adjusts from moment to moment to maintain the infection rate at this same level in order to maintain the equality of marginal benefit and marginal cost. Since the susceptible population declines over time, social distancing must increase steadily while $I_t = I_{MBC}$ to maintain I_t at this value. The infection rate remains at I_{MBC} until it cannot be maintained at this level even without social distancing (i.e. herd immunity is reached).

The comparative statics of policy interventions that influence parameters T , A_I , and D are driven by the effects of these parameters on the transition function of disease without social distancing and on I_{MBC} . For example, an increase in T necessarily causes a reduction in I_{MBC} , so that an increase in transmissibility always yields a lower peak infection rate.

Incentives for vaccination follow the logic of the analysis of Proposition 3 with linear costs of social distancing. Since vaccination is cost-effective during periods of time when the infection rate is above a critical threshold, with homogeneous cost of vaccination, in equilibrium all vaccination takes place at a single time. Further, the level of vaccination is determined by equilibrium indifference - vaccination proceeds until everyone is indifferent between vaccinating at that time and remaining unvaccinated forever.³² A vaccine mandate has the same qualitative effect with linear costs of social distancing as in the baseline model with convex and increasing costs of social distancing. For instance, a work-based mandate imposes costs on unvaccinated people at time t_M and thereafter.

³²Since vaccination does not affect any of the parameters in the marginal benefit function, it does not change the peak level of infection but does change the cumulative infection level required for herd immunity.

As in Proposition 4, a mandate yields an increase in vaccination rates, along with reduction in cumulative infection rates, but no change in expected utility since it does not change expected utility for those who vaccinate in equilibrium without a mandate.³³

Similarly, with heterogeneous costs of vaccination and linear costs for social distancing in a continuous model, vaccination proceeds with lowest-cost types choosing to vaccinate in sequence, with each type vaccinating when the infection rate reaches the cost-effectiveness threshold for that type. Once again, vaccination stops when the marginal person (i.e. lowest cost-type that is not yet vaccinated) is indifferent between vaccinating and remaining unvaccinated. Since a work-based mandate nominally reduces future payoffs for unvaccinated people beyond time t_M , it induces additional people - those with next lowest vaccination costs - to choose vaccination. That is, the effect of a vaccine mandate with heterogeneous costs of vaccination and linear costs of social distancing mirrors the qualitative effect of the mandate observed in Proposition 5. Gans (2021) finds a different result for an activity-based mandate under these assumptions (linear costs of social distancing and heterogeneous costs of vaccination) primarily because that paper focuses on the (temporary) steady state where the infection rate is at peak level for some period of time.

³³An increase in vaccination does not reduce the peak infection rate associated with linear costs but does reduce the length of time (possibly to zero) that the infection rate is maintained at that peak level for a shorter time.