Steady-State Social Distancing and Vaccination

Christopher Avery, Frederick Chen, and David McAdams

On-Line Appendix

A Quadratic Costs

Suppose that the flow cost of social distancing x is quadratic, $c(x) = \frac{\alpha x^2}{2}$, so that $c'(x) = \alpha x$. First, we verify that the steady-state continuation value for infected agents C_I and the individually-optimal social-distancing level x^* for susceptible agents are each linear functions of the steady-state continuation value for susceptible agents C_S, namely,

$$C_I = a_1 + b_1 C_S \tag{A.1}$$

$$x^* = a_2 + b_2 C_S. \tag{A.2}$$

(To simplify exposition, we suppress the dependence of C_I, C_S, and x^* on the steady-state infection prevalence, which we denote here simply as I.) Equation (A.1) follows directly from equation (7) in the main text, which states that $C_I = \frac{d+\gamma e^{-rt} C_S}{\gamma + r}$; so,

$$a_1 = \frac{d}{\gamma + r} \tag{A.3}$$

$$b_1 = \frac{\gamma e^{-rt}}{\gamma + r}. \tag{A.4}$$

By the first-order condition (10), $c'(x^*) = \alpha x^* = \beta I(C_I - C_S)$; so, $x^* = \frac{\beta I}{\alpha}(C_I - C_S)$. This verifies equation (A.2) with

$$a_2 = \frac{\beta I}{\alpha} a_1 \tag{A.5}$$

$$b_2 = \frac{-(1 - b_1)\beta I}{\alpha}. \tag{A.6}$$

Next, we verify that the steady-state value of C_S for any given infection prevalence I can be determined by solving a quadratic equation. By equation (8), $C_S = \frac{c(x) + \beta(1-x) I C_I}{\beta(1-x)I + r}$. Cross-multiplying and substituting $C_I = a_1 + b_1 C_S$ and $x = a_2 + b_2 C_S$ gives

$$C_S[\beta(1 - a_2 - b_2 C_S)I + r] = \frac{\alpha}{2}(a_2 + b_2 C_S)^2 + \beta I(1 - a_2 - b_2 C_S)(a_1 + b_1 C_S). \tag{A.7}$$

This yields a quadratic equation of form $XC_S^2 + YC_S + Z = 0$ where

$$X = \beta I b_2 + \frac{\alpha}{2} b_2^2 - \beta I b_1 b_2 \tag{A.8}$$

$$Y = \alpha a_2 b_2 + \beta I(1 - a_2) b_1 - \beta I a_1 b_2 - r - \beta I(1 - a_2) \tag{A.9}$$

$$Z = \frac{\alpha}{2} a_2^2 + \beta I(1 - a_2) a_1. \tag{A.10}$$

Finally, we provide details on how to compute SS_V, the "supply of vaccination required for equilibrium" shown in Figures 2-3. In a steady-state equilibrium with infection rate I, there is a constant flow of agents from the infected to the recovered state at rate γI. Since recovered agents have temporary immunity for t_R units of time, the proportion of temporarily-immune agents is $\gamma I t_R$ in this steady-state equilibrium. After accounting for infected and temporarily immune agents, and assuming that a stationary proportion SS_V of the population is vaccinated, the proportion of susceptible agents at each moment in the steady-state equilibrium is $S = 1 - I(1 + \gamma t_R) - SS_V$. Thus, the flow of agents from the susceptible to the infected state is $S\beta I(1 - x^*)$. In a steady-state equilibrium, the flow in and out of the infected state must be the same, so

$$\gamma I = (1 - I(1 + \gamma t_R) - SS_V)\beta I(1 - x^*). \tag{A.11}$$

Solving (A.11) allows us to compute SS_V.

B Details of Example 1 and Related Figures

We use the quadratic cost framework for the examples and figures in the paper.

B.1 Choice of Parameters:

We set the recovery parameter $\gamma = 1$ so that one unit of time is equal to the average length of contagiousness for an infected person. As explained in Footnote 7 in the main text, we chose $\beta = 3$ and $t_R = 20$ to be broadly consistent with recent estimates for SARS-CoV2 in 2020.[i] We set the remaining parameters $d = 1$, $\alpha = 0.1$ (so that $c(x) = 0.05x^2$), $t_V = t_R = 20$ and chose $r = -\ln(0.95)$ to be consistent with a discount factor $\delta = 0.95$ per unit time. We used this set of parameters for the computations in Figures 1a and 1b.

We augmented those choices of parameters with additional assumptions about the distribution of vaccination costs for Figure 2. Specifically, we used a trial-and-error process with the aid of a spread-

[i] Estimating the underlying biological parameters of the new variants that emerged later during the pandemic is much more conceptually challenging. For example, Omicron's *effective* reproduction number (i.e., its actual success at causing new infections per infected host) has been estimated as being several times higher than Delta's after it first emerged (Liu and Rocklöv (2022)). However, it is unclear how much of Omicron's relative success at that time was due to its ability to re-infect hosts who had already recovered from another SARS-CoV-2 strain versus its inherent transmissibility.

sheet to choose the upper limit $\overline{c}_V = .6468$ of a uniform distribution $U(0, \overline{c}_V)$ of vaccination costs to yield (approximate) steady-state infection rate 1% with endogenous social distancing and endogenous vaccination. We used a similar trial-and-error process to choose vaccine subsidy .1917 to yield (approximate) steady-state infection rate 0.5% with endogenous social distancing and endogenous vaccination given that subsidy.

We maintained these assumptions and $U(0, .6468)$ distribution of vaccination costs for the baseline case in Figure 3a. For the social-distancing subsidy in Figure 3a, we halved the original value of α to 0.05 so that the cost of social-distancing level x with the subsidy is $.025x^2$.

For Figure 3b, we maintained the choices $\beta = 3, t_R = t_V = 20, d = 1, r = -\ln(0.95)$, and $c(x) = .05x^2$ from the baseline case. In this case, we used a trial-and-error process to identify bounds for a tighter uniform distribution of vaccination costs to yield (approximate) steady-state infection rate 1% (as in the baseline case for Figures 2 and 3a) with endogenous social distancing and endogenous vaccination as well as (approximate) steady-state infection rate 1.1% with a social-distancing subsidy.

B.2 Computations for the Examples and Figures

Equation (A.11) identifies the required supply of vaccination for a steady-state equilibrium at each infection level.

Given the assumption that vaccination provides immunity for t_v units of time, the lifetime cost of vaccination at discount factor .95 per unit time is $\frac{c_V}{1-.95^{t_v}}$. Comparing the lifetime cost of vaccination to the expected lifetime cost $C_S^*(I)$ in the steady state with infection prevalence I when (initially) susceptible and unvaccinated gives the threshold cost for vaccination.

$$c_V \leq \frac{C_S^*(I)}{1-.95^{t_V}}. \tag{B.1}$$

Thus, the "demand for vaccination" in the steady state with infection prevalence I with uniform distribution of vaccination costs $U(0, \overline{c}_V)$ is

$$D_V(I) = \frac{C_S^*(I)(1-.95^{t_V})}{\overline{c}_V}. \tag{B.2}$$

With vaccine subsidy s_V per vaccination, the lifetime cost of vaccination adjusts to $\frac{c_V - s_V}{1-.95^{t_v}}$. In addition, the distribution of vaccination costs net of the subsidy is $U(-s_V, \overline{c}_V - s_V)$. Accounting for these changes, the "demand for vaccination" with subsidy s_V in the steady state with infection prevalence I is

$$D_V(I) = \frac{C_S^*(I)(1-.95^{t_V}) + s_V}{\overline{c}_V}. \tag{B.3}$$

B.3 Details of the Spreadsheet containing Computations

The computations for Figures 1a, 2, 3a, and 3b are provided in the Excel spreadsheet "Computations for Figures 1a, 2, 3a, and 3b." There is one separate worksheet of data and computations for each of Figure 1, Figure 2a, and Figure 3a and two separate worksheets for Figure 3b. Each of these worksheets follows the same format with one row per infection rate.

The entries in these worksheets are as follows. Column A contains the fixed parameter values. Column D lists the conjectured steady-state infection prevalence. Column E contains the associated proportion of recently recovered and currently immunized agents. The right-most columns find the roots of the quadratic equation given the parameters, as defined by (A.8), (A.9), and (A.10). Column F identifies the relevant root of the quadratic equation as the steady-state-cost value for susceptible agents C_S. Columns G and H contain the associated values of costs C_R and C_I. Given these values, Column I contains the optimal steady-state level of social distancing for a susceptible agent, which follows from the fixed parameter values and the cost values in Columns F through H. Column J contains the flow rate of new infections. Column K contains the flow rate of recoveries, where these values correspond to the relevant portions of equations (1) and (2) from the main text.

The computations in Columns A through K assume that no one is vaccinated, so the susceptible proportion is simply 1 minus the proportion who are either infected or immune. Column L uses the following logic to identify the vaccination proportion that would equate the flow rate of new infections and recoveries at the given infection prevalence: The difference in the flow rates of infection absent vaccination is "Column J - Column K", so dividing by "Column J" gives the proportion of new infections that would have to be avoided by vaccination to equate these flow rates. Multiplying this fraction ("Column J" - "Column K")/"Column J") by the susceptible proportion (i.e. 1 - Column D - Column E) gives the vaccination rate required for a steady-state equilibrium with this infection prevalence.

It is also possible to calculate the vaccination rate required for a steady-state equilibrium (as a function of underlying parameters of the model) from the steady-state condition. In particular,"Proportion Susceptible" * "Flow Rate of New Infections" = "Proportion of Recoveries", or

$$(1 - V - I - t_R)N = \gamma I, \tag{B.4}$$

where V is the proportion who are vaccinated, I is the proportion infected, $\gamma t_R I$ is the proportion immune, N is the flow rate of new infections (i.e. $N = \beta I(1 - V - I - \gamma t_R I)$, and γI is the proportion of infections per unit time in a steady state. Solving (B.4) for V gives

$$V = 1 - I - \gamma t_R I - \frac{\gamma}{\beta(1 - x)}. \tag{B.5}$$

Column M of the worksheet for Figure 1a, Column N of the worksheets for Figures 3a and 3b,

and Column O of the worksheet for Figure 2 uses (B.5) to compute the level of vaccination required for a steady-state equilibrium with infection prevalence given by the value in Column D of that row. These computations match the values in Column L, verifying that these two methods for computing the required vaccination level are equivalent.

Column M of the worksheets for Figures 2, 3a, and 3b uses (B.3) to compute the demand for vaccination (i.e. the proportion of susceptible agents who would adopt vaccination) at the given stationary infection prevalence. Column N of the worksheet for Figure 2 recomputes the demand for vaccination given a subsidy for vaccination.

The relevant equilibrium outcomes are highlighted in the top rows of each table. In the table labeled "Figure 1a", Row 1 highlights the (approximate) equality of the flow rate of new infections and recoveries at stationary infection rate 2.226% without vaccination, while Row 2 highlights the fact that a 36.69% vaccination rate is required for an equilibrium with stationary infection rate 1% given endogenous vaccination.

In the table labeled "Figure 2", Row 1 highlights the (approximate) equality of supply and demand for vaccination at stationary infection rate 1% given the parameters listed above, while Row 2 highlights the approximate equality of supply and demand for vaccination at stationary infection rate 0.5% with the vaccination subsidy identified above.

In the table labeled "Figure 3a", Row 2 highlights the (approximate) equality of supply and demand for vaccination at stationary infection rate 0.825% given the original parameters *except* for a change in the value of α to .05 after accounting for a social-distancing subsidy.

In the table labeled "Figure 3b, No Subsidy", Row 1 highlights the (approximate) equality of supply and demand for vaccination at stationary infection rate 1% given the original parameters *except* for a change in the lower and upper limits of the uniform distribution of vaccination costs.

In the table labeled "Figure 3b, Subsidy", Row 1 highlights the (approximate) equality of supply and demand for vaccination at stationary infection rate 1.1% given the original parameters *except* for a change in the lower and upper limits of the uniform distribution of vaccination costs and a reduction in the value of α to .05397.

B.4 Details of the Spreadsheets for Producing Figures 1a, 2, 3a, and 3b

The Excel spreadsheet "Figures" uses results from the "Technical Appendix" spreadsheet to produce Figures 1, 2, 3a, and 3b. This spreadsheet includes one table labeled "Data" with the relevant values from the other spreadsheet and separate worksheets for each of the figures themselves.

Columns C and D of the worksheet "Figures / Data" contain data from Columns J and K of the worksheet "Technical Appendix / Figure 1".

Columns G through I of the worksheet "Figures / Data" contain data from Columns L through N of the worksheet "Technical Appendix / Figure 2".

Columns L and M in the worksheet "Figures / Data" repeat the data from Columns G and H in that same worksheet. Columns N and O in the worksheet "Figures / Data" contain data from Columns L and M of the worksheet "Technical Appendix / Figure 3a."

Columns R and S in the worksheet "Figures / Data" contain data from Columns L and M ofthe worksheet "Technical Appendix / Figure 3b No Subsidy." Columns T and U in the worksheet "Graphs / Data" contain data from Columns L and M of the worksheet "Technical Appendix / Figure 3b Subsidy."

B.5 Details of the Spreadsheet for Figure 1b

The Excel spreadsheet "Computations for Figure 1b" generates the results and produces the graph for Figure 1b.

We used a straightforward algorithm to estimate convergence of an epidemic to steady-state equilibrium. Our computations are shown in a series of worksheets titled "Sheet 1" through "Sheet 8". We consider a 150-period model where each period is divided into 50 equal length segments of time with period-length-adjusted parameters and discounting structure derived from the parameters of Example 1.

In each iteration of the algorithm, we assume an initial infection rate of .01% and use a conjectured time series for C_S and C_I. Our analysis for each iteration proceeds in two steps. In the first step for a given iteration, we compute the infection trajectory for the current iteration with the following procedure. Starting at time $t = 0$, we compute the individually-optimal level of social distancing for susceptible people at time t given the current infection rate and the anticipated values for C_S and C_I at time $t + .02$, then use that level of social distancing to compute the resulting infection rate at time $t + .02$. Iterating this process forward from time $t = 0$ to time $t = 150$, we trace the infection trajectory induced by the conjectured time-series values for C_I and C_S.

In the second step of each iteration, we work backwards from period 150 to compute the realized time-series values for C_S and C_I given the infection trajectory and social-distancing pattern that was identified in the first step of analysis. We start in period 150 by assuming that the steady-state values for C_S and C_I will be realized in period 150.02. Then at each time t, we compute realized expected costs for each state from the level of social distancing at time t and the **computed expected costs** for time $t + .02$. Iterating this processbackwards from time $t = 150$ to time $t = 0$, we identify the full time series of induced expected costs for this iteration of the algorithm.

The only difference between one iteration of the algorithm to the next is the choice of the conjectured time series for C_S and C_I. In the very first iteration, we assume constant values for C_S and C_I equal to

0 at each point in time from $t = 0$ to $t = 150$. Then for iteration $n + 1$, we use the realized time-series values for C_S and C_I in iteration n as the approximations for the conjectured time series values for C_S and C_I in iteration $n + 1$.

We present each iteration of these computations as a separate sheet in the spreadsheet. The algorithm converged, with time series of infection prevalence as illustrated in Figure 1b.

C Derivation of the Bellman Equation

Here we show how the continuous-time dynamic programming problem ((5) in the paper) is derived from a discrete-time problem. At any time t, we have

$$C_I(t) = d\triangle + \gamma\triangle(\frac{1}{1+r\triangle})C_R(t+\triangle) + (1-\gamma\triangle)(\frac{1}{1+r\triangle})C_I(t+\triangle), \tag{C.1}$$

where $\triangle > 0$ is a small time interval. Since susceptible agents choose to minimize the expected present value of costs incurred over the rest of their lifetime, where these costs come from the cost of social distancing and the cost of being infected, we have

$$C_S(t) = \min_{x\in[0,1]}\{c(x)\triangle + \beta(1-x)I\triangle(\frac{1}{1+r\triangle})C_I(t+\triangle) + (1-\beta(1-x)I\triangle)(\frac{1}{1+r\triangle})C_S(t+\triangle)\}. \tag{C.2}$$

In (C.2), $c(x)\triangle$ is the cost of social distancing and $\beta(1-x)I\triangle$ is the probability of acquiring an infection over the time period from t to $t+\triangle$.

From (C.1), we get

$$C_I(t) - C_I(t+\triangle) = d\triangle + \gamma\triangle(\frac{1}{1+r\triangle})C_R(t+\triangle) - (\gamma+r)\triangle(\frac{1}{1+r\triangle})C_I(t+\triangle).$$

Divide by $\triangle$ and take limits as $\triangle \to 0$. This gives us

$$-C_I'(t) = d + \gamma C_R(t) - (\gamma+r)C_I(t), \tag{C.3}$$

where $C_R(t) = e^{-rt_R}C_S(t+t_R)$.

Similarly, from (C.2), we get

$$C_S(t) - C_S(t+\triangle) = \min_{x\in[0,1]}\{c(x)\triangle + \beta(1-x)I\triangle(\frac{1}{1+r\triangle})C_I(t+\triangle) - (\beta(1-x)I+r)\triangle(\frac{1}{1+r\triangle})C_S(t+\triangle)\}.$$

Dividing by $\triangle$ and taking limits as $\triangle \to 0$ yields the continuous time Bellman equation (Hamilton-Bellman-Jacobian equation):

$$-C_S'(t) = \min_{x\in[0,1]}\{c(x) + \beta(1-x)IC_I(t) - (\beta(1-x)I+r)C_S(t)\}. \tag{C.4}$$

In a steady state, $C_I'(t) = C_S'(t) = 0$. Let $C_S^*(I)$, $C_I^*(I)$, and $C_R^*(I)$, respectively, denote the steady

state expected lifetime cost for susceptible, infected, and newly recovered agents. We have

$$C_R^*(I) = e^{-rt_R} C_S^*(I).$$

From (C.3), we get

$$0 = d + \gamma C_R^*(I) - (\gamma + r) C_I^*(I),$$

which yields

$$C_I^*(I) = \frac{d}{\gamma + r} + \frac{\gamma}{\gamma + r} C_R^*(I).$$

Similarly, (C.4) gives, in a steady state,

$$0 = \min_{x \in [0,1]} \{c(x) + \beta(1-x) I C_I^*(I) - (\beta(1-x)I + r) C_S^*(I)\}.$$

From this, we obtain

$$C_S^*(I) = \min_{x \in [0,1]} \left\{ \frac{c(x) + \beta(1-x) I C_I^*(I)}{\beta(1-x)I + r} \right\}.$$

D Details of Example 2

We chose $\beta = 3$, $\gamma = 1$, $r = -\ln(0.95)$, $d = 1$, $t_R = 100$, $c(x) = 0.05x^2$, and $V = 0$. The solution to the individual optimization problem is a solution $x(I)$ to a quadratic equation. The steady-state condition is

$$\beta(1-x)(1 - I - \gamma t_R I) = \gamma. \tag{D.1}$$

Solving for x gives

$$x = \frac{303I - 2}{3(101I - 1)}. \tag{D.2}$$

The steady-state equilibrium solves $x(I)$ and (D.2) simultaneously, which pinpoints the steady-state infection level $I = 0.00610$ and associated level of social distancing $x = 0.131$.

By contrast, the objective function for the social planner is

$$I(x)\frac{d}{r} + [1 - (1 + \gamma t_R) I(x)] \frac{0.05x^2}{r}. \tag{D.3}$$

Solving (D.1) for I gives

$$I(x) = \frac{2 - 3x}{303(1 - x)}. \tag{D.4}$$

Combining (D.3) and (D.4), the planner's objective function (in expected cost) is minimized at $x = 0.104$. That is, there is **too much** social distancing in the steady-state equilibrium relative to the social

optimum.

E A Property of the Function $C_S(I,x)$

Let $x^*(I)$ denote the individually-optimal social-distancing level for susceptible agents given stationary infection prevalence I. Here we show that $\frac{\partial C_S(I,x^*(I))}{\partial I} > 0$ whenever $x^*(I) < 1$.[ii] From (11) in the paper, we have

$$C_S(I,x) = \frac{c(x) + \beta(1-x) I\left(\frac{d}{\gamma+r}\right)}{\beta(1-x) Ik + r}, \tag{E.1}$$

where $k = 1 - \frac{\gamma e^{-rt_I}}{\gamma+r}$. Differentiating (E.1),

$$\frac{\partial C_S(I,x)}{\partial x} = \frac{1}{\beta(1-x) Ik + r}\left[c'(x) - \frac{\beta I\left(\frac{rd}{\gamma+r} - c(x) k\right)}{\beta(1-x) Ik + r}\right]. \tag{E.2}$$

By the first-order condition for $x^*(I)$ to be a minimum, $\frac{\partial C_S(I,x^*(I))}{\partial x} = 0$. Since $c'(x^*(I)) > 0$, this requires by (E.2) that

$$\frac{rd}{\gamma+r} > c(x^*(I)) k. \tag{E.3}$$

From (E.1) we now obtain as desired that

$$\frac{\partial C_S(I,x)}{\partial I} = \frac{\beta(1-x)}{(\beta(1-x) Ik + r)^2}\left[\frac{rd}{\gamma+r} - c(x) k\right] > 0$$

at $x = x^*(I)$.

[ii] For sufficiently large I, susceptible agents find it optimal to completely isolate themselves, choosing $x^*(I) = 1$ and having expected lifetime costs $C_S(I,1) = \frac{c(1)}{r}$. However, no steady-state equilibrium can have such a high infection level, since full isolation by susceptible agents causes the flow of new infections to fall to zero, a contradiction.