Federal Reserve Bank of Cleveland Working Paper Series

Optimal Epidemic Control in Equilibrium with Imperfect Testing and Enforcement

Thomas Phelan and Alexis Akira Toda

Working Paper No. 21-15

August 2021

Suggested citation: Phelan, Thomas, and Alexis Akira Toda. 2021. "Optimal Epidemic Control in Equilibrium with Imperfect Testing and Enforcement." Federal Reserve Bank of Cleveland, Working Paper No. 21-15. https://doi.org/10.26509/frbc-wp-202115.

Federal Reserve Bank of Cleveland Working Paper Series
ISSN: 2573-7953

Optimal Epidemic Control in Equilibrium with Imperfect Testing and Enforcement*

Thomas Phelan† Alexis Akira Toda‡

August 3, 2021

Abstract

We analyze equilibrium behavior and optimal policy within a Susceptible-Infected-Recovered epidemic model augmented with potentially undiagnosed agents who infer their health status and a social planner with imperfect enforcement of social distancing. We define and prove the existence of a perfect Bayesian Markov competitive equilibrium and contrast it with the efficient allocation subject to the same informational constraints. We identify two externalities, static (individual actions affect current risk of infection) and dynamic (individual actions affect future disease prevalence), and study how they are affected by limitations on testing and enforcement. We prove that a planner with imperfect enforcement will always wish to curtail activity, but that its incentives vanish as testing becomes perfect. When a vaccine arrives far into the future, the planner with perfect enforcement may encourage activity before herd immunity. We find that lockdown policies have modest welfare gains, whereas quarantine policies are effective even with imperfect testing.

Keywords: efficiency, externalities, lockdown, perfect Bayesian equilibrium, quarantine.

JEL codes: C73, D50, D62, I12.

1 Introduction

Soon after evidence of the first community spread of the coronavirus 2019 (COVID-19) outside of China was reported in Italy in late February 2020, European countries promptly introduced drastic mitigation measures ("lockdown"), such as the closure of schools, restaurants, and other businesses. Many states and provinces in

*The views stated herein are those of the authors and are not necessarily those of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System.

†Federal Reserve Bank of Cleveland. Email: tom.phelan@clev.frb.org.

‡Department of Economics, University of California San Diego. Email: atoda@ucsd.edu.

the United States and Canada as well as countries around the world had followed suit by mid-March. While implementing policies to slow the spread of an infectious disease appears to be an obvious course of action for a prudent government, such mitigation policies have evidently not been without costs. Curtailing economic activities can cause unemployment, bankruptcy, and reduced access to education. Further, engaging in non-pharmaceutical interventions (better known as "social distancing") reduces new infections but also delays achieving herd immunity, and so may prolong the epidemic in the absence of a vaccine. It is therefore possible that lockdown policies can slow the progress of the epidemic but do little to alter its ultimate toll.

Although it appears we are now in the later stages of the pandemic, with several vaccines developed and administered, the immense disruption to economic and social activity wrought by the virus and the possibility of future pandemics (either due to variants or new viruses) motivates a theoretical analysis of an epidemic model suitably augmented with realistic features to capture policy-relevant tradeoffs. In this paper we build upon the standard Kermack and McKendrick (1927) Susceptible-Infected-Recovered (SIR) model and add two important features. First, the agents in our model are forward-looking and endogenously respond to the epidemic, but they must continuously update their beliefs about their own health status because they may lack symptoms or testing may not be available. Second, we study the extent to which prescriptions for policy depend upon the ability of the government to enforce their recommended actions over the long term. The first feature is motivated by the fact that the infection fatality rate (IFR, fraction of deaths among *all* cases) of COVID-19 inferred from seroprevalence studies is an order of magnitude smaller than the case fatality rate (CFR, fraction of deaths among *confirmed* cases), which suggests substantial underreporting.[1] The second is motivated by the fact that some governments were both slow to impose social distancing measures and may lack the ability to enforce such measures over the long run, possibly due to opposition from constituents. To the best of our knowledge, no existing work considers these two features and studies the role they play in shaping optimal policy responses.

Our model works as follows. The society consists of four behavioral types of agents: unknown, infected, recovered, and dead. The unknown type consists of agents who lack immunity (are susceptible) against the infectious disease as well

[1] According to the meta-analysis of seroprevalence studies by Ioannidis (2021), the median IFR of COVID-19 is 0.27 percent. On the other hand, as we document in Section 4, the median CFR across more than 200 countries and regions is 1.35 percent. Thus, the reporting rate is an the order $0.27/1.35 = 20$ percent.

as those who are infected or recovered but unconfirmed due to imperfect testing or asymptomatic infections. Each period, alive agents take actions $a \in [0, 1]$, which we interpret as their overall level of economic activity. We assume that in the absence of an epidemic, agents prefer taking the highest action $a = 1$. During an epidemic, taking higher actions exposes oneself to the risk of infection. As a result, rational agents without confirmed immunity (the unknown type) optimally choose lower actions, i.e., they voluntarily practice social distancing. Known infected and recovered agents have no incentive to social distance and choose the highest action available to them. We define a perfect Bayesian Markov competitive equilibrium to be an allocation in which (i) agents form beliefs about their health status and optimize given the state variables (population shares of each type) and (ii) the evolution of beliefs and state variables are consistent with Bayes' rule and the collective behavior of agents. We obtain two main theoretical results. First, we prove the existence of a (pure strategy) perfect Bayesian Markov competitive equilibrium (Markov equilibrium for short). This result is important because to achieve the equilibrium, individuals and policymakers only need to form expectations about the future given a few state variables and do not require implausibly sophisticated coordination among them.

The equilibrium allocation is in general inefficient due to the externality caused by the actions of infected agents. Existing analyses of the pandemic have either focused upon allocations in which a planner dictates all activity in perpetuity,[2] or laissez-faire allocations in which no social distancing is imposed.[3] Although obviously informative, neither of these cases models the problem of a planner who was previously slow to act or who believes her capacity to enforce restrictions may dissipate over time. The recursive nature of our solution methods allows us to address this situation, as we compute equilibrium and efficient activity levels at every point in the state space. In order to highlight the role that imperfect enforcement plays in optimal policy responses, we distinguish between two types of externalities: static and dynamic. The static externality arises because the activity of infected agents affects the probability that susceptible agents become infected in the current period. The dynamic externality arises because the collective behavior of agents affects the evolution of the prevalence of the virus throughout the population. The interplay between the static and dynamic externalities is subtle as they can move in different directions: when individuals choose higher activity,

[2]Examples include Alvarez et al. (2020), Farboodi et al. (2020), and Kruse and Strack (2020), among others.

[3]An example is Toxvaerd (2020).

they increase the risk of infecting their fellow citizens today, but if they become infected, they reduce the risk they pose to others in the future.

This brings us to the second theoretical result. We prove that the difference between the static efficient (the optimal choice taking future prevalence as given) and equilibrium actions is bounded above by a number proportional to the fraction of unknown infected agents, a quantity that vanishes as the probability of diagnosis converges to 1. This shows that a government who can only enforce short-term lockdown policies will always wish to curb activity, but their incentive to do so vanishes with the fraction of unconfirmed cases. This observation is noteworthy because we also provide examples in which a government with unlimited enforcement power wishes to recommend *higher* activity than that which occurs in equilibrium. This phenomenon is observed in situations in which the vaccine is not expected to arrive until far into the future, so that herd immunity is the only possible end to the pandemic. In these circumstances a benevolent government may wish to *encourage* activity until the mass of susceptible agents is sufficiently small, before imposing social distancing to minimize the total number of infections (and hence deaths) necessary to reach herd immunity. It is in this sense that we show how policy prescriptions can depend crucially on enforcement capabilities.

The presence of diagnosed and undiagnosed infected agents implies that there are two tools for intervention available to the planner: activity recommendations for unknown agents (referred to as *lockdown* policies) and recommendations for known infected agents (referred to as *quarantine* policies). This combination of undiagnosed agents and the possibility of imperfect enforcement capabilities is precisely what makes the problem of the planner difficult. Indeed, if all infected agents could be immediately and costlessly quarantined, then the pandemic would likely not have had the immense economic and social impact that we have observed over the past year. In the majority of this paper, we therefore focus primarily on the recommendations to unknown agents, taking as given a fixed level of activity by the known infected agents.

To illustrate our theory as well as to study the optimal interventions, we calibrate our model to the current COVID-19 epidemic and conduct a number of robustness exercises. Due to endogenous social distancing, in the Markov equilibrium the infection curve flattens relative to the case of myopic behavior (the standard SIR model). We find that the planner's optimal interventions are significantly affected by the diagnosis rate of infections, the vaccine arrival rate, and the planner's ability to enforce policies over the long term. The welfare gains from lockdown policies tend to be small, reducing the welfare loss from the pandemic relative to the equilibrium by less than 10 percent. When the vaccine is expected to arrive within a year or two (as could be expected at the beginning of the COVID pandemic), the optimal reduction in activity begins earlier, is more gradual, and extends well beyond the date at which activity has returned to normal in the equilibrium allocation. In contrast, when the vaccine is expected to arrive decades into the future, the optimal policy is to *encourage* (discourage) agents to take high actions *before* (after) achieving herd immunity so that herd immunity is achieved quickly but unnecessary deaths are avoided. In contrast to lockdown policies, we find that quarantine policies are effective even with imperfect testing.

1.1 Related literature

Kermack and McKendrick (1927) present the basic mathematical framework for studying the evolution of infectious diseases. Models that build upon this framework assume that agents can be placed into categories based on their health status, that there is a fixed probability that an infected agent passes the infection to a susceptible agent when they meet, and that there is a fixed rate at which infected agents recover or die. In the simplest formulation, there are only susceptible and infected agents. These SI models are appropriate for infectious diseases that are incurable but not deadly.[4] SIS models are ones in which infected agents can recover, but when they do, they become susceptible to reinfection. SIR models are ones in which infected agents can recover (or die) and acquire lifelong immunity.[5]

Although mathematical epidemic models provide insights into the spread of infectious diseases, they often ignore individual choice or public policy.[6] There are several papers that modify the basic SIR model to allow for either government policies or individual decisions to influence the course of the epidemic. We describe the models as *non-strategic* if the government has the ability to mandate changes through quarantines, lockdowns, or other non-pharmaceutical interventions. We call the models *strategic* if individual agents independently decide levels of care

[4] An example is the Epstein-Barr virus (EBV) infection, which causes latent lifelong infections. Most people are infected during childhood and experience only mild symptoms such as fever. When infection occurs among adolescents and adults, EBV causes infectious mononucleosis (kissing disease) in 30–50 percent of the cases.

[5] It is not yet clear whether patients who recover from COVID-19 acquire lifelong immunity. Lifelong immunity is observed in measles and rubella, whereas immunity against influenza is relatively short-lived. However, SIR models are likely appropriate for modeling the current COVID-19 epidemic because once a sufficient fraction of the population acquires partial immunity, the situation will likely be similar to that of the seasonal influenza.

[6] Eksin et al. (2019) document that mathematical epidemic models that ignore behavioral changes have large forecast errors.

that influence exposure to the disease.

Sethi (1978) examines the problem of a planner who can choose to quarantine a fraction of infected agents in an SIS model. With linear payoffs and costs, he identifies a bang-bang solution in which the planner either quarantines all infected agents or none of them. More recently, Kruse and Strack (2020) incorporate social distancing in a non-strategic model of infections and study a version of the SIR model in which a social planner can, at a cost, influence the transmission rate. With linear cost, they show that socially optimal policies are bang-bang: the social planner reduces the transmission rate as much as possible or does not reduce the transmission rate at all. It is typically not optimal to reduce the transmission rate when the fraction of infected agents is small. For some of their analysis, Kruse and Strack (2020) assume that the planner can only impose social distancing for (no more than) a fixed length. With this restriction, lockdowns should start only after the number of infected agents reaches a threshold.

Turning to strategic models, Geoffard and Philipson (1996), Kremer (1996), and Auld (2003) study the extent to which strategic choices may undermine the effects of public policy regarding the spread of an infectious disease. These papers focus on HIV (human immunodeficiency virus) infections, where the heterogeneity of the population and the ability to select whom to interact with are of first-order importance. Reluga (2010), Chen et al. (2011), Fenichel et al. (2011), Chen (2012), Fenichel (2013), and Toxvaerd (2019) present strategic models of social distancing that predates the current COVID-19 epidemic. Fenichel (2013), which is an extended analysis of Fenichel et al. (2011), assumes that agents can select the intensity of their interaction with others and assumes that flow utility is a single-peaked concave function of this intensity. He contrasts socially optimal choices of these contact levels with privately selected values and points out that if the social planner cannot distinguish between groups (and therefore any restriction on interactions must apply to susceptible, infected, and recovered agents alike), then social welfare may be higher in the laissez-faire equilibrium than in the constrained social planner's problem. This possibility arises because the planner's intervention constrains the participation of recovered agents (who generate positive externalities) in addition to the participation of infected agents.

Chen et al. (2011) and Chen (2012) study a static game in which susceptible agents decide on their level of activities. This game may exhibit multiple equilibria, which are typically inefficient. Whether the susceptible agents are more or less active than the social optimum depends on the nature of the matching technology. More recently, Toxvaerd (2019) points out that in the presence of strategic agents, public policy interventions that lower the infectiousness of a disease may lower social welfare because agents respond to the change by increasing their own exposure.

Following the onset of the COVID-19 epidemic, a large number of papers have been written by economists. Since this literature is too large to review, we only discuss the subset of papers that focus on the theory and applications. Abel and Panageas (2020) study an optimal control problem as well as the laissez-faire equilibrium in an SIR model with population growth and show that a steady state exists and the disease becomes endemic no matter how large the cost from excess death is. Budish (2020) conceptualizes $\mathcal{R} \leq 1$ (effective reproduction number less than 1) as a constraint, discusses the optimal policy in a static setting with heterogeneous economic activities, and illustrates that cheap policies such as mask wearing go a long way in containing the virus spread with minimal welfare costs. Toxvaerd (2020) studies an SIR model with endogenous social distancing, which is similar to ours. Assuming linear utility and costs, he shows that susceptible agents either engage in no social distancing at all or social distance to maintain a target peak prevalence, which endogenously flattens the infection curve.

Relative to this small literature of theoretical strategic epidemic models, our main contribution is that we explicitly model imperfect testing and enforcement and systematically study the welfare implications and optimal policies.

2 Behavioral SIR model with imperfect testing

We introduce rational and potentially undiagnosed Bayesian agents into the basic Kermack and McKendrick (1927) Susceptible-Infected-Recovered (SIR) epidemic model.

2.1 Model

We consider an infectious disease that can be transmitted between agents in a society, which consists of a large (but finite) number of agents indexed by $n \in N = \{1, \dots, N\}$. Time is discrete, runs forever, and is denoted by $t = 0, \Delta, 2\Delta, \dots,$ where $\Delta > 0$ is the length of time in one period.

Agent types and information At each point in time, agents are categorized into several types based on their health status and information. An agent who does not have immunity against the infectious disease is called *susceptible* and denoted by S. An infected agent who is known (unknown) to be infected is denoted by I_k (I_u). An alive agent with immunity who is known (unknown) to be immune is called *removed* (or *recovered*) and denoted by R_k (R_u). A dead agent is denoted by D. The set of all types (health status) is denoted by

$$H \coloneqq \{S, I_k, I_u, R_k, R_u, D\}.$$

For instance, an agent could be I_k if he tests positive for antigen or he shows specific symptoms (is *symptomatic*), and I_u if no tests are available and he shows no specific symptoms (is *asymptomatic*). Similarly, an agent could be R_k if he tests positive for antibodies, he recovered from a past symptomatic infection and immunity is lifelong, or he is vaccinated. Thus the set of information types is

$$\{U \coloneqq S \cup I_u \cup R_u, I_k, R_k, D\},$$

where U denotes the *unknown* type. When no confusion arises, we refer to an I_k agent as "infected" and an R_k agent as "removed/recovered," without the qualifier "known." Importantly, we suppose that when an agent gets infected, with probability $\sigma \in (0, 1]$ the agent receives a signal that reveals the true health status (known infected, I_k). Otherwise, the agent becomes unknown infected (I_u). Although we refer to the signal as a "test," the signal could be literally a laboratory test as well as other information such as the presence of specific symptoms, knowledge of close contacts with confirmed cases, etc. We refer to the probability σ as the *diagnosis* rate. In Appendix B we show that a model with a single signal is observationally equivalent to a model with multiple signals with potentially heterogeneous fatality rates.

Let $N_h \subset N$ be the set of agents with health status $h \in H$. With a slight abuse of notation, we also use the same symbol h to denote the fraction of type h agents in the population, so $h = |N_h|/N$. The space of the aggregate state (type distribution) is denoted by

$$Z \coloneqq \left\{ z = (S, I_k, I_u, R_k, R_u, D) \,\middle|\, S + I_k + I_u + R_k + R_u + D = 1, Nz \in \mathbb{Z}_+^6 \right\}. \tag{2.1}$$

We suppose that the aggregate state is observable. (I_k, R_k, D are observable, and I_u and R_u can be inferred from a small scale random antigen and antibody testing.) The economy starts at $t = 0$ with some initial condition $z_0 \in Z$.

Actions and preferences At each point in time, each alive agent takes action $a \in A = [\underline{a}, 1]$, where the minimum action $\underline{a}$ satisfies $0 \leq \underline{a} < 1$. We interpret a as the economic activity level: loosely, $a = 1$ corresponds to following a normal life and $a = \underline{a}$ to being completely locked down (minimum activity level for subsistence). The utility function of an unknown (U) and known recovered (R_k) agent is denoted by $u : A \to \mathbb{R}$. The utility function of a known infected agent (I_k) is denoted by $u_I : A \to \mathbb{R}$. A dead agent receives the flow utility $u_D \in \mathbb{R}$.[7] Agents discount future payoffs with discount factor $\mathrm{e}^{-r\Delta}$, where $r > 0$ is the discount rate.

Disease transmission Agents meet each other randomly over time and transmit the infectious disease. If agents n, n' take actions $a_n, a_{n'}$, respectively, then agent n bumps into agent n' during a period with probability $\lambda \Delta a_n a_{n'}/N$, where $\lambda \in (0, 1/\Delta)$ is a parameter (meeting rate) that governs the level of social interaction with full activity ($a_n = a_{n'} = 1$). We take λ as given, which depends on how the society is organized (e.g., population density, whether workers commute by car or public transportation, whether consumers shop online or at physical stores, whether classes are taught remotely or in-person, etc.).

If agent n is susceptible ($n \in N_S$) and agent n' is infected ($n' \in N_{I_k} \cup N_{I_u}$), the infectious disease is transmitted from n' to n with probability $\tau \in (0, 1]$ conditional on n bumping into n' at time t.[8] We also take τ as given, which depends on how contagious the disease is as well as how the society is organized (e.g., how often people wash their hands, whether they wear masks, whether they greet others by bowing, shaking hands, hugging, or kissing, etc.).

We assume that Δ is small enough such that in any period an agent bumps into at most one other agent. Therefore if $a_h := (\sum_{n' \in N_h} a_{n'}) / |N_h|$ denotes the average action of type h, a particular susceptible agent n who takes action a gets infected with probability

$$\begin{aligned} \sum_{n' \in N_{I_k} \cup N_{I_u}} \tau \frac{\lambda \Delta a a_{n'}}{N} &= \tau \lambda \Delta a \left(\frac{\sum_{n' \in N_{I_k}} a_{n'}}{|N_{I_k}|} \frac{|N_{I_k}|}{N} + \frac{\sum_{n' \in N_{I_u}} a_{n'}}{|N_{I_u}|} \frac{|N_{I_u}|}{N} \right) \\ &= \beta \Delta (a_{I_k} I_k + a_{I_u} I_u) a, \end{aligned} \tag{2.2}$$

[7]More precisely, u_D is the flow utility of being dead anticipated by alive agents.

[8]Thus the act of "bumping into" is asymmetric between the members in a meeting. Figuratively, here is what happens if agent n bumps into n' (in the case of an upper respiratory infection such as COVID-19): (i) agent n meets agent n', (ii) agent n' sneezes into agent n's face (and transmits the disease to n with probability τ if n' is infected), and (iii) they part with each other.

where $\beta \coloneqq \tau\lambda$ is the baseline transmission rate and we have used the notation $h = |N_h|/N$ for $h = I_k, I_u$. The timing convention is that if an infection occurs at time t, the (previously susceptible, now infected) agent changes status to I_k or I_u at time $t + \Delta$. Since only susceptible agents are prone to infection, the expected fraction of the population that gets newly infected between time t and $t + \Delta$ is

$$\beta \Delta a_S S(a_{I_k} I_k + a_{I_u} I_u), \tag{2.3}$$

which is called *incidence* in epidemiology. The fraction of infected agents

$$|N_{I_k} \cup N_{I_u}|/N = I_k + I_u \tag{2.4}$$

is called *prevalence*.

Recovery and death An infected agent is removed (becomes no longer infected by either recovering or dying) with probability $\gamma\Delta$ each period, where $\gamma \in (0, 1/\Delta]$ is the removal rate. Conditional on being removed, a known infected (I_k) agent dies with probability $\delta \in (0, 1]$ and an unknown (I_u) agent always recover. The rationale for this assumption is that infected agents with more severe symptoms are more likely to get tested as well as to die. Letting δ_0 be the fatality rate among all (known and unknown) infected agents, since the diagnosis rate is σ, we have

$$\delta = \frac{\delta_0}{\sigma}. \tag{2.5}$$

Finally, recovered or dead agents remain in their corresponding states forever, implying that recovered agents acquire lifelong immunity (see Footnote 5). Again the timing convention is that if an infected agent is removed at time t, the agent changes status to R_k, R_u, or D at time $t + \Delta$.

In epidemiology, there are several notions of fatality rate, and it is important to understand the distinction. The fatality rate among all (known and unknown) infected cases (which corresponds to δ_0) is called the *infection fatality rate (IFR)*. The fatality rate among known (confirmed) infected cases (which corresponds to δ if the signal is a laboratory test) is called the *case fatality rate (CFR)*. The fatality rate among the entire population is called *mortality*. Clearly, by definition we have

$$\text{Mortality} \le \text{IFR} \le \text{CFR}.$$

Vaccine arrival We assume that a vaccine arrives at a Poisson rate $\nu \geq 0$, independent of everything else. Thus in our discrete-time setting, the probability that a vaccine arrives between time t and $t+\Delta$ is $1-\mathrm{e}^{-\nu\Delta}$. We assume that the vaccine is perfectly effective, perfectly safe, and has no cost. Thus once a vaccine arrives, all non-infected agents will be vaccinated and become immune (R_k). The vaccine is not a cure and hence has no effect on infected agents.

2.2 Assumptions

Throughout the rest of the paper, we maintain the following assumptions.

Assumption 1 (Utility function)**.** *The utility functions satisfy the following conditions: (i) $u : A = [\underline{a}, 1] \to \mathbb{R}$ is twice continuously differentiable and satisfies $u(1) = 0$, $u' > 0$, and $u'' < 0$, (ii) $u_I : A = [\underline{a}, 1] \to \mathbb{R}$ is continuous, strictly concave, and achieves a unique maximum at $a_I \in A$, and (iii) $u_D < u_I(a_I) \leq u(1)$.*

The assumption $u(1) = 0$ simplifies the algebra and is without loss of generality because we can shift the utility functions by a constant without affecting behavior. The assumptions $u_D < u_I(a_I) \leq u(1)$ simply imply that being asymptomatic is preferable to being symptomatic, which is in turn preferable to being dead. The condition that u_I is single-peaked at $a_I \in A$ implies that a potentially intermediate value of activity level (rest) is myopically optimal for symptomatic agents. This assumption can also be interpreted as altruism, sense of duty, or an enforcement of a quarantine policy.

Assumption 2 (Perfect competition)**.** *Agents view the evolution of the aggregate state z as exogenous and ignore the impact of their behavior on the aggregate state.*

Assumption 2 is necessary for analytical tractability and is reasonable when the number of agents N is large.

Assumption 3 (Consistency)**.** *On equilibrium paths, agents update their beliefs using the Bayes rule. Off equilibrium paths, unknown (U) agents believe they are susceptible with probability*

$$\mu(z) := \begin{cases} \frac{S}{S+I_u+R_u} & \text{if } S > 0, \\ 0 & \text{otherwise.} \end{cases} \tag{2.6}$$

The assumption that agents apply the Bayes rule may not be realistic because a pandemic such as COVID-19 is rare and agents may have difficulty forming beliefs when faced with an unprecedented situation. However, we focus on the Bayes rule because it provides a benchmark analysis. The assumption that we specify the off-equilibrium beliefs as in (2.6) suggests that our equilibrium concept would be a perfect Bayesian equilibrium (PBE). We discuss this point in more detail in Section 3.2.

3 Equilibrium analysis

This section defines and establishes the existence of equilibrium and characterizes individual behavior.

3.1 Individual behavior

We formalize the individual optimization problem recursively. Let $z \in Z$ be the aggregate state and V_h be the value function of type $h \in \{U, I_k, R_k, D\}$ agents.

Dead agents Because dead agents remain dead and their flow utility is u_D, we obtain

$$V_D = (1 - \mathrm{e}^{-r\Delta})u_D + \mathrm{e}^{-r\Delta}V_D \iff V_D = u_D.$$

Known recovered agents Because known recovered agents have lifelong immunity, their value function is constant and the Bellman equation is

$$V_{R_k} = \max_{a \in A} \left\{(1 - \mathrm{e}^{-r\Delta})u(a) + \mathrm{e}^{-r\Delta}V_{R_k}\right\}.$$

The optimal action is clearly $a_{R_k} = 1$ (full activity) and the value function is $V_{R_k} = u(1) = 0$ by Assumption 1.

Known infected agents By assumption, known infected agents are removed with probability $\gamma\Delta$, and conditional on removal, die with probability $\delta = \delta_0/\sigma$. Since the health status transitions are independent of the aggregate state and action, the value function is constant and the Bellman equation is

$$V_{I_k} = \max_{a \in A} \left\{(1 - \mathrm{e}^{-r\Delta})u_I(a) + \mathrm{e}^{-r\Delta}(\underbrace{(1 - \gamma\Delta)V_{I_k}}_{\text{stay infected}} + \underbrace{\gamma\Delta[(1 - \delta)V_{R_k} + \delta V_D])}_{\text{removal}}\right\}. \tag{3.1}$$

By Assumption 1 the function u_I is single-peaked at a_I, and so the optimal action is $a_{I_k} = a_I$. Since $V_{R_k} = u(1) = 0$ and $V_D = u_D$, (3.1) simplifies to

$$V_{I_k} = \frac{(1 - \mathrm{e}^{-r\Delta})u_I + \mathrm{e}^{-r\Delta}\gamma\Delta\delta u_D}{1 - \mathrm{e}^{-r\Delta}(1 - \gamma\Delta)}, \tag{3.2}$$

where $u_I := u_I(a_I)$. Note by Assumption 1 that we have $u_I = u_I(a_I) \le u(1) = 0$, so $V_{I_k} < 0 = V_{R_k}$.

Unknown agents Because unknown agents need to infer their health status, the analysis is more complicated. Suppose unknown agents adhere to a policy function $a_U : Z \to A$ and that they always have the belief (2.6). (We will verify that these assumptions are satisfied in equilibrium.) The policy $a_U(z)$, together with the mechanisms of disease transmission, symptom development, recovery, and death, generates transition probabilities $\{q(z, z')\}_{(z,z') \in Z^2}$ for the aggregate state conditional on no vaccine arrival. (Note that Z in (2.1) is a finite set.) By Assumption 2, agents view this law of motion as exogenous. Let $V_U(z)$ be the value function of an unknown agent who chooses the action optimally in this environment. By (2.2) and the analysis of known infected (I_k) agents, an agent taking full action ($a = 1$) gets infected with probability

$$p(z) := \beta\Delta(a_I I_k + a_U(z) I_u) \tag{3.3}$$

conditional on being susceptible. Noting that infection is known with probability σ and the vaccine arrives with probability $1 - \mathrm{e}^{-\nu\Delta}$, the Bellman equation is

$$\begin{aligned} V_U(z) = \max_{a \in A} \Bigg\{ &(1 - \mathrm{e}^{-r\Delta})u(a) + \underbrace{\mathrm{e}^{-r\Delta}\sigma\mu p a V_{I_k}}_{\text{known infection}} \\ &+ \underbrace{\mathrm{e}^{-r\Delta}(1 - \sigma\mu p a)}_{\text{stay unknown}} \mathrm{E}_z\big(\underbrace{\mathrm{e}^{-\nu\Delta} V_U(z')}_{\text{no vaccine}} + \underbrace{(1 - \mathrm{e}^{-\nu\Delta})V_{R_k}}_{\text{vaccine}}\big) \Bigg\}, \end{aligned} \tag{3.4}$$

where E_z denotes the expectation with respect to $\{q(z, z')\}$, $\mu = \mu(z)$ is given by (2.6), $p = p(z)$ is given by (3.3), and V_{I_k} is given by (3.2). Noting that $V_{R_k} = u(1) = 0$, (3.4) simplifies to

$$V_U(z) = \max_{a \in A} \left\{ (1 - \mathrm{e}^{-r\Delta})u(a) + \mathrm{e}^{-r\Delta}\,\mathrm{E}_z((1 - \sigma\mu p a)\mathrm{e}^{-\nu\Delta} V_U(z') + \sigma\mu p a V_{I_k}) \right\}. \tag{3.5}$$

The following proposition establishes the existence and uniqueness of V_U and provides some bounds on value functions.

Proposition 3.1 (Value functions)**.** *Fix a policy function $a_U : Z \to A$ of unknown agents. Then there exists a unique value function $V_U : Z \to \mathbb{R}$ satisfying the Bellman equation (3.5). Furthermore, the value functions satisfy the following inequalities:*

$$\begin{aligned} V_D = u_D &< \frac{(\mathrm{e}^{r\Delta}-1)u_I + \gamma\Delta\delta u_D}{\mathrm{e}^{r\Delta}-1+\gamma\Delta} = V_{I_k} \\ &< \frac{\mathrm{e}^{\nu\Delta}\sigma\beta\Delta}{\mathrm{e}^{(r+\nu)\Delta}-1+\sigma\beta\Delta} V_{I_k} \le V_U(z) \le V_{R_k} = 0. \end{aligned} \tag{3.6}$$

The proof of Proposition 3.1, as well as other longer proofs, is deferred to Appendix A. The inequality (3.6) is quite intuitive. In terms of flow utility, having no symptoms is better than having symptoms, which is better than death. Because the states R_k, D are absorbing and a known infected agent may recover or die, the inequalities $V_D < V_{I_k} < V_{R_k} = 0$ are immediate. The inequality $V_U \le V_{R_k}$ is also immediate because an unknown agent could get infected and generally chooses a lower action. The inequality $V_{I_k} < V_U$ follows from the fact that an unknown agent can always choose the myopic optimal action ($a = 1$, which generates flow utility $0 = u(1) \ge u_I(a_I) = u_I$) but gets infected only in the future.

The following proposition characterizes the best response of an unknown agent. To state the result, we define the inverse marginal utility function $\phi : \mathbb{R} \to [\underline{a}, 1]$ by

$$\phi(x) := \begin{cases} 1 & \text{if } x \le u'(1), \\ (u')^{-1}(x) & \text{if } u'(1) < x < u'(\underline{a}), \\ \underline{a} & \text{if } x \ge u'(\underline{a}). \end{cases} \tag{3.7}$$

Proposition 3.2 (Best response of U agents)**.** *Fix a policy function $a_U : Z \to A$ and a continuation value $V_U : Z \to \mathbb{R}$ of unknown agents. Then the best response of an unknown agent is*

$$a^* = \phi\left(\frac{\sigma\mu(z)p(z)}{\mathrm{e}^{r\Delta}-1}(\mathrm{E}_z\, \mathrm{e}^{-\nu\Delta}V_U(z') - V_{I_k})\right), \tag{3.8}$$

where $\mu(z)$ and $p(z)$ are given by (2.6) and (3.3), respectively.

Because u' is strictly decreasing, ϕ in (3.7) is decreasing. As a result, we immediately obtain the following corollaries. Corollary 3.3 provides comparative statics results in a partial equilibrium setting (where the policy function $a_U(z)$ and the continuation value $V_U(z)$ are exogenously given).

Corollary 3.3 (Comparative statics)**.** *The best response of an unknown agent a^* in (3.8) increases if (i) ν decreases (vaccine arrival less likely), (ii) $\mu(z)$ or $p(z)$ decrease (infection less likely), (iii) δ_0 decreases (death less likely), and (iv) u_I or u_D increase (infection and death less scary).*

The result that a^* decreases in $\mu(z), p(z), \delta_0$ and increases in u_I, u_D is intuitive: increasing $\mu(z), p(z), \delta_0$ or decreasing u_I, u_D makes the risk of infection or death higher, which makes agents take more precautions. The fact that the possibility of vaccine arrival (higher ν) makes agents take more precautions is also intuitive: when agents expect a vaccine to arrive soon, being locked down is less painful because it is expected to be short. As we shall see later in our numerical analysis, the vaccine arrival rate ν significantly affects the socially efficient action. The next corollary shows that when prevalence is sufficiently low, agents take no precautions.

Corollary 3.4 (Full activity with sufficiently low prevalence)**.** *Let $I := I_k + I_u$ be the prevalence defined in (2.4). There exists $\bar{I} > 0$ such that for all policy function $a_U(z)$, we have $a^* = 1$ whenever $I < \bar{I}$. In particular, we can take*

$$\bar{I} = -\frac{e^{r\Delta} - 1}{\Delta} \frac{u'(1)}{\sigma \beta V_{I_k}}, \tag{3.9}$$

where $V_{I_k} < 0$ is given by (3.2).

Letting $\bar{I}$ as in (3.9) and using (3.6), $u_I \leq 0$, and (2.5), we obtain

$$\bar{I} \leq \frac{e^{r\Delta} - 1}{\Delta} \left(\frac{e^{r\Delta} - 1}{\gamma \Delta} + 1 \right) \frac{u'(1)}{\beta \delta_0 (-u_D)}. \tag{3.10}$$

Since $u_D < 0$, the right-hand side of (3.10) is increasing in $\nu, r, u'(1), u_D$ and decreasing in β, γ, δ_0. Thus impatience, a higher marginal utility of full action, less disutility from death, a lower transmission rate, a longer duration of infections, and lower mortality all make agents more likely to take full action. The intuition is identical to Corollary 3.3.

3.2 Definition and existence of equilibrium

Our equilibrium concept is the (pure strategy) perfect Bayesian Markov competitive equilibrium defined below. Here "perfect Bayesian" means that agents update beliefs on equilibrium paths using the Bayes rule as in Assumption 3; "Markov" means that the optimal actions agents choose are functions of state variables; and "competitive" means that agents view the evolution of aggregate state variables as exogenous as in Assumption 2.

Definition 3.5 (Markov equilibrium)**.** A (pure strategy) *perfect Bayesian Markov competitive equilibrium*, or *Markov equilibrium* for short, consists of unknown agents' belief $\mu(z)$ of being susceptible, transition probabilities $\{q(z, z')\}_{z,z' \in Z}$ for the aggregate state, value functions $\{V_h(z)\}_{h=U,I_k,R_k,D}$, and policy functions $\{a_h(z)\}_{h=U,I_k,R_k}$ such that

(i) (Consistency) The belief $\mu(z)$ satisfies the Bayes rule on equilibrium paths and is given by (2.6) off equilibrium paths; the transition probabilities $\{q(z, z')\}$ are consistent with individual actions and the mechanisms of disease transmission, symptom development, recovery, and death,

(ii) (Sequential rationality)

(U) $V_U(z)$ satisfies the Bellman equation (3.5) and $a = a_U(z)$ achieves the maximum, where E_z denotes the conditional expectation using $\{q(z, z')\}$ and $p = p(z)$ is as in (3.3),

(I_k) $V_{I_k}(z)$ is as in (3.2) and $a_{I_k}(z) = a_I$,

(R_k) $V_{R_k}(z) = 0$ and $a_{R_k}(z) = 1$,

(D) $V_D(z) = u_D$.

Note that Definition 3.5 only describes the society before vaccine arrival. Once the vaccine arrives, because no new infections occur by assumption, it is optimal for all agents to take their myopic optimal action ($a = 1$ for $h = U, R_k$ and $a = a_I$ for $h = I_k$) and the problem becomes trivial.

The astute reader may wonder why we adopt the notion of perfect Bayesian equilibrium and specify that beliefs is given by (2.6) even off the equilibrium path. The belief $\mu(z)$ equals the posterior belief if unknown agents have a common prior, take identical actions, and learn that the aggregate state is z, but an unknown Bayesian agent who is contemplating deviating from the equilibrium action $a_U(z)$ generally has a different posterior. The primary reason for our choice of solution concept is that of tractability, as we wish to study the role of forward-looking agents that are uncertain about their health status without obscuring the analysis with technicalities.[9] The notion of perfect Bayesian equilibria allows agents to be

[9]To be more specific, in our discrete-time setting, insisting on applying the Bayes rule off the equilibrium path leads to non-concave maximization problems, and hence the existence of equilibrium is not warranted.

forward-looking and rational, if somewhat "forgetful." In Appendix C we extend the model to the case with perfect recall and show that our results are robust.

The following theorem establishes the existence of equilibrium.

Theorem 3.6 (Existence of equilibrium)**.** *Suppose Assumptions 1–3 hold. Then there exists a pure strategy perfect Bayesian Markov competitive equilibrium, where the belief $\mu(z)$ always satisfies* (2.6)*.*

The idea of the proof of Theorem 3.6 is to start with a guess of equilibrium policy $a_U(z)$, update it as the (necessarily unique) maximizer in (3.5), show that this updating rule is continuous, and apply the Brouwer fixed-point theorem to establish a fixed-point of this operation.

3.3 Externalities and efficiency

In this section we study the source of externalities in the model and the efficiency properties of the equilibrium. Let $a_U(z)$ and $a_{I_k}(z)$ be any policy functions for unknown agents and known infected agents exogenously chosen by the social planner. Since the behavior of known recovered agents does not affect the aggregate state dynamics, without loss of generality we set $a_{R_k}(z) = 1$, which is individually optimal. Suppose the planner wishes to choose $(a_U(z), a_{I_k}(z))$ to maximize social welfare. In general, individual agents have a strong incentive to deviate from such recommendations and choose the individually optimal actions characterized by Proposition 3.2. Thus let $\lambda \geq 0$ be the hazard rate of failing to enforce the recommended policy $(a_U(z), a_{I_k}(z))$, and suppose that once the policy implementation fails, the society reverts back to the Markov equilibrium characterized in Theorem 3.6. Letting $V_h^\lambda(z; a_U, a_I)$ be the value function of type h agents in this environment, since the probability of reverting to equilibrium is $1 - \mathrm{e}^{-\lambda\Delta}$, we have

$$
\begin{aligned}
V_U^\lambda(z) =& (1-\mathrm{e}^{-\lambda\Delta})V_U(z) + \mathrm{e}^{-\lambda\Delta}\big[(1-\mathrm{e}^{-r\Delta})u(a_U) \\
&+ \mathrm{e}^{-r\Delta}\,\mathrm{E}_z((1-\sigma\mu p a_U)\mathrm{e}^{-\nu\Delta}V_U^\lambda(z') + \sigma\mu p a_U V_{I_k}^\lambda(z'))\big], \qquad (3.11\text{a}) \\
V_{I_k}^\lambda(z) =& (1-\mathrm{e}^{-\lambda\Delta})V_{I_k} + \mathrm{e}^{-\lambda\Delta}\big[(1-\mathrm{e}^{-r\Delta})u_I(a_{I_k}) \\
&+ \mathrm{e}^{-r\Delta}\,\mathrm{E}_z((1-\gamma\Delta)V_{I_k}^\lambda(z') + \gamma\Delta((1-\delta)V_{R_k} + \delta V_D))\big], \qquad (3.11\text{b})
\end{aligned}
$$

where $V_U(z)$ and V_{I_k} are the equilibrium value functions in Theorem 3.6, $\mu = \mu(z)$ and $p = p(z)$ are given by (2.6) and (3.3), and we write $a_h = a_h(z)$ and $V_h^\lambda(z) = V_h^\lambda(z; a_U, a_{I_k})$. By the standard contraction mapping argument, V_h^λ is well-defined.

The utilitarian social welfare associated with the policies $(a_U(z), a_{I_k}(z))$ is then

$$\begin{aligned}\mathcal{W}^\lambda(z) &:= (S + I_u + R_u)V_U^\lambda(z) + I_k V_{I_k}^\lambda(z) + R_k V_{R_k}(z) + DV_D(z) \\ &= (S + I_u + R_u)V_U^\lambda(z) + I_k V_{I_k}^\lambda(z) + Du_D, \end{aligned} \tag{3.12}$$

where we have used $V_{R_k} = 0$ and $V_D = u_D$. Noting that u_D is constant and D depends only on z (and not on the policies $(a_U(z), a_{I_k}(z))$), conditional on the aggregate state z, we can rank social welfare associated with particular policy functions by using

$$W^\lambda(z) := UV_U(z)^\lambda + I_k V_{I_k}^\lambda(z) \tag{3.13}$$

instead of $\mathcal{W}^\lambda$, where $U := S + I_u + R_u$ denotes the fraction of U agents.

Given the hazard rate λ, the optimal policies $(a_U(z), a_{I_k}(z))$ maximize (3.13). The Markov equilibrium in Definition 3.5 is generally inefficient because the equilibrium policies $(a_U^*(z), a_{I_k}^*(z))$ do not maximize the welfare criterion (3.13) due to externalities. We identify two types of externalities, static and dynamic. The static externality refers to the fact that when (known or unknown) infected agents take higher actions, they infect susceptible agents with some probability and affect the *current* value function of unknown agents, even if the future value functions and the aggregate state transitions remain the same. The dynamic externality refers to the fact that the collective behavior of agents affects the aggregate state transitions and hence future value functions. Because studying the dynamic effect analytically is challenging, we focus on the static effect in this section, and study the dynamic externality in our numerical exercises.

To study the static externality, imagine that the social planner can intervene to alter agents' current actions (a_U, a_{I_k}), fixing the transition probabilities $\{q(z, z')\}$ of the aggregate state as well as the next period's value functions $\{V_h(z')\}_{h=U,I_k,R_k,D}$. Using (3.1), (3.5), (3.13), and noting that $V_{R_k} = 0$, we can define the objective function of the planner that seeks to eliminate the static externality by

$$\begin{aligned} &W(a_U, a_{I_k}, z) := \\ &\quad U\left[(1 - \mathrm{e}^{-r\Delta})u(a_U) + \mathrm{e}^{-r\Delta}\,\mathrm{E}_z((1 - \sigma\mu p a_U)\mathrm{e}^{-\nu\Delta}V_U(z') + \sigma\mu p a_U V_{I_k}(z'))\right] \\ &\qquad + I_k\left[(1 - \mathrm{e}^{-r\Delta})u_I(a_{I_k}) + \mathrm{e}^{-r\Delta}\,\mathrm{E}_z((1 - \gamma\Delta)V_{I_k}(z') + \gamma\Delta\delta V_D)\right]. \end{aligned} \tag{3.14}$$

Restricting the action of known infected agents can be interpreted as quarantine. Restricting the action of unknown agents can be interpreted as lockdown. Hence we introduce the following definition.

Definition 3.7 (Static efficient actions)**.** Fix transition probabilities $\{q(z,z')\}$ and value functions V_U, V_{I_k}. Given a policy function $a_U(z)$ of unknown agents, we say that the *quarantine policy* $a^\dagger_{I_k}(z)$ is *static efficient* if $a_{I_k} = a^\dagger_{I_k}(z)$ achieves the maximum of (3.14). Similarly, given a policy function $a_{I_k}(z)$ of known infected agents, we say that the *lockdown policy* $a^\dagger_U(z)$ is *static efficient* if $a_U = a^\dagger_U(z)$ achieves the maximum of (3.14).

The reason we use the qualifier "static" can be understood as follows. Imagine that a benevolent government implements some lockdown policy when faced with an epidemic. The optimal policy then obviously depends on the scale and duration of the intervention. If the intervention is implemented on a large scale, the dynamic effects cannot be ignored. Furthermore, because individuals have an incentive to deviate from lockdown policies, the optimal policies when the government has unlimited enforcement power may be different from a short-term policy. The policies in Definition 3.7 are efficient in the sense that the government seeks to implement the optimal policy on a small scale when anticipating the continuation values V_U, V_{I_k} (which are the equilibrium value functions if the government can commit only for one period).

The following proposition characterizes the static efficient quarantine policy.

Proposition 3.8 (Static efficient quarantine policy)**.** *Suppose the utility function of known infected agents u_I is continuously differentiable, strictly concave, with inverse marginal utility function ϕ_I analogous to* (3.7)*. Then the static efficient quarantine policy is given by*

$$a^\dagger_{I_k}(z) = \phi_I\left(\frac{\sigma\beta\Delta}{\mathrm{e}^{r\Delta}-1}\mu(z)a_U(z)U\,\mathrm{E}_z(\mathrm{e}^{-\nu\Delta}V_U(z') - V_{I_k}(z'))\right). \tag{3.15}$$

Proof. Let W be as in (3.14). Using $p = \beta\Delta(a_{I_k}I_k + a_U I_u)$, we obtain

$$\begin{aligned}\frac{\partial W}{\partial a_{I_k}} &= -U\mathrm{e}^{-r\Delta}\sigma\mu\beta\Delta I_k a_U\,\mathrm{E}_z(\mathrm{e}^{-\nu\Delta}V_U(z') - V_{I_k}(z')) + I_k(1-\mathrm{e}^{-r\Delta})u_I'(a_{I_k}) \\ &= I_k\mathrm{e}^{-r\Delta}\Delta\left(\frac{\mathrm{e}^{r\Delta}-1}{\Delta}u_I'(a_{I_k}) - \sigma\mu\beta a_U U\,\mathrm{E}_z(\mathrm{e}^{-\nu\Delta}V_U(z') - V_{I_k}(z'))\right).\end{aligned}$$

The rest of the proof is the same as Proposition 3.2. ∎

Proposition 3.8 has several implications. First, (3.15) does not depend on the fraction of known infected agents I_k except through the continuation values.

This is because the welfare gain from reducing new infections and the welfare loss from restricting infected agents' actions are both proportional to I_k, as we can see from the proof of Proposition 3.8. Second, under normal circumstances the planner seeks to quarantine known infected agents intensively. To see this, using (3.3), note that the individually optimal action of unknown agents (3.8) can be rewritten as

$$a_U^* = \phi\left(\frac{\sigma\beta\Delta}{\mathrm{e}^{r\Delta}-1}\mu(z)(a_{I_k}I_k + a_U I_u)\,\mathrm{E}_z(\mathrm{e}^{-\nu\Delta}V_U(z') - V_{I_k}(z'))\right). \tag{3.16}$$

Assuming that infected and unknown agents have identical preferences (so $u_I = u$), the only difference between (3.15) and (3.16) is that the argument of the former is proportional to $a_U U$, whereas the argument of the latter is proportional to $a_{I_k}I_k + a_U I_u$. In normal situations, the fraction of actively infected agents in the society is rather small, so $I_k, I_u \ll U$. Then the argument of (3.15) is a much bigger number than that of (3.16), and because ϕ is decreasing, we would have $a_{I_k}^\dagger \ll a_U^*$.

We next study the optimal intervention for unknown agents. Suppose all known infected agents take action $a_I \in A$, which can be interpreted as an exogenous quarantine policy by the discussion after Assumption 1. Let $a_U^*(z)$ be the equilibrium policy of unknown agents and V_U, V_{I_k} the corresponding value functions established in Theorem 3.6. The following theorem provides a bound for the static efficient lockdown policy, which is our main theoretical result.

Theorem 3.9 (Static efficient lockdown policy). *Let $a_U^*(z)$ be the equilibrium policy of unknown agents and $m = \min_{a\in A}|u''(a)| > 0$. There exists a unique static efficient lockdown policy $a_U^\dagger(z)$, which satisfies*

$$\begin{aligned} 0 &\le a_U^*(z) - a_U^\dagger(z) \\ &\le \min\left\{\frac{\sigma\beta\Delta}{m(\mathrm{e}^{r\Delta}-1)}\mu(z)a_U^*(z)I_u\,\mathrm{E}_z(\mathrm{e}^{-\nu\Delta}V_U(z') - V_{I_k}), 1-\underline{a}\right\}. \end{aligned} \tag{3.17}$$

In particular, if $\sigma = 1$ (so $I_u = 0$), then $a_U^\dagger(z) = a_U^*(z)$.

Theorem 3.9 has several implications. First, the inequality $a_U^* \ge a_U^\dagger$ implies that the equilibrium action of unknown agents is too high relative to the static efficient lockdown policy. This is intuitive, because a subset of unknown agents are infected and impose a negative externality on others by going out and infecting other susceptible agents. Second, and more interestingly, the difference between the static efficient and equilibrium actions is bounded above by a number proportional to the fraction of unknown infected agents. This quantity is typically small, and approaches zero as $\sigma \uparrow 1$. A government who can only enforce short-term lockdown policies will therefore always (weakly) wish to curb activity, but their incentive to do so vanishes with the fraction of unconfirmed cases. This is particularly noteworthy because we shall later provide examples in which a government with unlimited enforcement power would wish to do the *reverse*, and recommend higher activity than that which occurs in equilibrium. It is in this sense that policy prescriptions can depend crucially on the enforcement capabilities of the government.

3.4 Equilibrium dynamics

This section studies the equilibrium dynamics. Given arbitrary policy functions $\{a_h(z)\}_{h=U,I_k,R_k}$, we define the transition probabilities $\{q(z,z')\}$ induced by these policies and the mechanisms of disease transmission, symptom development, recovery, and death. To simplify the notation, let $z_t = z$, $S_t = S$, $S_{t+\Delta} = S_{+\Delta}$, etc. Using (2.2), it is straightforward to show that

$$\mathrm{E}_z(S_{+\Delta} - S) = -\beta\Delta a_U(z)S(a_{I_k}(z)I_k + a_U(z)I_u), \tag{3.18a}$$
$$\mathrm{E}_z(I_{k,+\Delta} - I_k) = \sigma\beta\Delta a_U(z)S(a_{I_k}(z)I_k + a_U(z)I_u) - \gamma\Delta I_k, \tag{3.18b}$$
$$\mathrm{E}_z(I_{u,+\Delta} - I_u) = (1-\sigma)\beta\Delta a_U(z)S(a_{I_k}(z)I_k + a_U(z)I_u) - \gamma\Delta I_u. \tag{3.18c}$$

(We omit the dynamics for R_k, R_u, D agents because they do not depend on policy functions.) We can simplify these equations further if we consider the limit $N \to \infty$ and apply the strong law of large numbers. In the large population limit, letting $I = I_k + I_u$ be the fraction of infected agents, we have $I_k = \sigma I$ and $I_u = (1-\sigma)I$. Therefore adding (3.18b) and (3.18c), we obtain the system of deterministic difference equations

$$\frac{S_{+\Delta} - S}{\Delta} = -\beta a_U(z)S(\sigma a_{I_k}(z) + (1-\sigma)a_U(z))I, \tag{3.19a}$$
$$\frac{I_{+\Delta} - I}{\Delta} = (\beta a_U(z)S(\sigma a_{I_k}(z) + (1-\sigma)a_U(z)) - \gamma)I = \gamma(\mathcal{R}(z) - 1)I, \tag{3.19b}$$

where

$$\mathcal{R}(z) := \frac{\beta}{\gamma}a_U(z)S(\sigma a_{I_k}(z) + (1-\sigma)a_U(z)) \tag{3.20}$$

is known as the *effective reproduction number* in epidemiology. At the early stage of the epidemic, by definition we have $S \approx 1$ and $I \approx 0$. Therefore by Corollary

3.4 (and assuming $u_I(a) = u(a)$), the equilibrium policies are $a_U(z) = a_{I_k}(z) = 1$, and (3.20) reduces to

$$\mathcal{R}(z) = \beta/\gamma =: \mathcal{R}_0, \tag{3.21}$$

which is known as the *basic reproduction number*. When unknown agents are myopic (so $a_U(z) = a_{I_k}(z) = 1$), (3.19) reduces to

$$\frac{S_{+\Delta} - S}{\Delta} = -\beta S I, \tag{3.22a}$$

$$\frac{I_{+\Delta} - I}{\Delta} = (\beta S - \gamma) I, \tag{3.22b}$$

which is the usual Kermack and McKendrick (1927) Susceptible-Infected-Recovered (SIR) model in discrete time. Since $a_U(z), a_{I_k}(z) \leq 1$, by (3.19b), (3.20), and (3.21), we always have

$$\frac{I_{+\Delta} - I}{\Delta} \leq \gamma(\mathcal{R}_0 S - 1) I.$$

Since by (3.19a) the fraction of susceptible agents S always decreases over time, the fraction of infected agents I decreases over time once

$$\mathcal{R}_0 S \leq 1 \iff S \leq 1/\mathcal{R}_0 \tag{3.23}$$

holds, where $\mathcal{R}_0$ is the basic reproduction number in (3.21). We say that society has achieved *herd immunity* if (3.23) holds.

4 Numerical analysis

In this section we use a numerical example calibrated to the COVID-19 epidemic to study how the agents' optimizing behavior, diagnosis rate, and lockdown policies affect the epidemic dynamics.

4.1 Model specification

One period corresponds to a day and we assume 5 percent annual discounting, so $r = 0.05/365.25$ and $\Delta = 1$. Because COVID-19 vaccines became available about one year after the start of the epidemic, we set the vaccine arrival rate to $\nu = 1/365.25$. Unless otherwise stated, we suppose that the social planner has unlimited enforcement power ($\lambda = 0$ in (3.11)–(3.13)).

We set the epidemic parameters from the medical literature. The daily transmission rate is $\beta = 1/5.4$ based on the median serial interval (the number of days it takes for an infected person to transmit the disease to another person) in the meta-analysis of Rai et al. (2021). The daily recovery rate is $\gamma = 1/13.5$ based on the median infectious period estimated in You et al. (2020). These numbers imply that the basic reproduction number in (3.21) is $\mathcal{R}_0 = \beta/\gamma = 13.5/5.4 = 2.5$, which equals the current best estimate used by the Centers for Disease Control and Prevention (CDC).[10] The infection fatality rate (IFR) is $\delta_0 = 0.0027$, which is the median value in the meta-analysis of Ioannidis (2021) on studies that use seroprevalence data.[11]

The choice of the diagnosis rate σ is more controversial. One possibility is to estimate the case fatality rate (CFR) δ and set $\sigma = \delta_0/\delta$ based on (2.5). Using the data on the cumulative number of reported cases and deaths provided by the Johns Hopkins University Center for Systems Science and Engineering,[12] we find that the median CFR across all (more than 200) regions is $\delta = 0.0134$. This would imply $\sigma = \delta_0/\delta = 0.2$ (20 percent). However, this calculation ignores other information such as the presence of symptoms. Another possibility is to set σ as the fraction of symptomatic agents. Based on the case study of the cruise ship Diamond Princess, which experienced a COVID-19 outbreak in February 2020 and whose passengers were all tested, Mizumoto et al. (2020) document that about 50 percent of confirmed cases were asymptomatic. Noting that the symptoms of COVID-19 are often similar to those of other upper respiratory infections and not specific enough to confirm the diagnosis, as a compromise, we choose an intermediate value $\sigma = 0.4$ for the diagnosis rate in our baseline analysis.

We set the minimum action to $\underline{a} = 0.01$ (which is somewhat arbitrary but never binds in simulations). The utility function of a symptom-free agent exhibits constant relative risk aversion (CRRA) of $\alpha > 0$, so

$$u(a) = \begin{cases} \frac{a^{1-\alpha}-1}{1-\alpha}, & (0 < \alpha \neq 1) \\ \log a, & (\alpha = 1) \end{cases} \tag{4.1}$$

which satisfies Assumption 1. We suppose $u_I = u$, so the optimal action for known infected agents is $a_I = 1$. Although this assumption is unrealistic because

[10] https://www.cdc.gov/coronavirus/2019-ncov/hcp/planning-scenarios.html

[11] Unlike the case fatality rate (CFR), which is defined as the reported number of deaths divided by the reported number of cases, the estimation of the IFR is complicated by the fact that cases and deaths may be underreported. The true number of cases can be estimated from random antibody testing, which is called a seroprevalence survey. If we assume that the underreporting of deaths is not severe, then the IFR can be estimated by dividing the reported number of deaths by the estimated number of cases.

[12] See Dong et al. (2020), data found at: https://github.com/CSSEGISandData/COVID-19

infected agents may be incapacitated, altruistic, or quarantined, it provides the most conservative (worst case) analysis. For numerical illustrations, we set $\alpha = 1$ (log utility). We calibrate the flow utility from death to be $u_D = -12.22$ based on the case study from Sweden, which did not introduce mandatory lockdowns (see Appendix E for details). Finally, we set the initial condition to $I_0 = 10^{-6}$, $S_0 = 1 - I_0$, $R_0 = 0$, and $D_0 = 0$.

4.2 Equilibrium, static efficient, and efficient actions

We solve for the perfect Bayesian Markov competitive equilibrium using the algorithm discussed in Appendix D. As a point of comparison, we also solve for the myopic allocation in which all agents choose $a = 1$, as well as the static efficient and efficient actions, which correspond to setting $\lambda = \infty$ and $\lambda = 0$ in (3.11)–(3.13), respectively. Figure 1 shows the epidemic dynamics studied in Section 3.4 for the myopic equilibrium, which is the standard Kermack and McKendrick (1927) SIR model. Here and elsewhere, the vertical dashed line indicates the first date of achieving herd immunity as defined by (3.23).

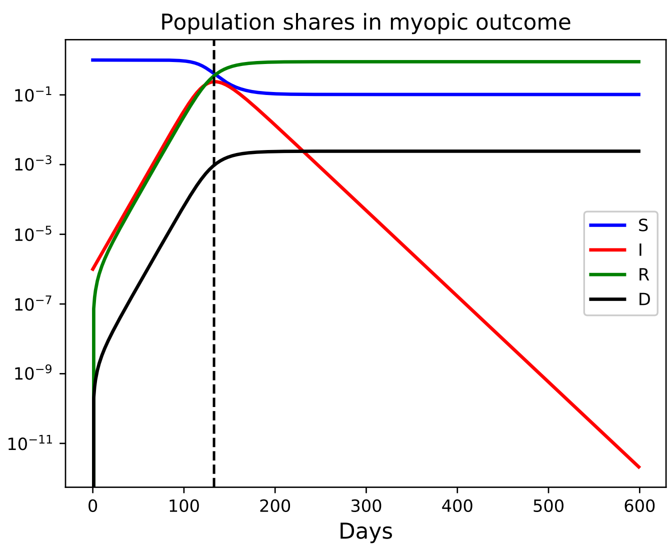

Figure 1: Epidemic dynamics in myopic equilibrium.

When agents are myopic and choose $a = 1$, as is well known, the epidemic has two phases: in the first phase, the fraction of infected agents initially grows exponentially (at a daily rate of approximately $\beta - \gamma = \gamma(\mathcal{R}_0 - 1)$ by setting $S = 1$ in (3.22b)) until society achieves herd immunity (peak prevalence is 23.9percent); in the second phase, the fraction of infected agents declines exponentially at a daily rate of γ. The epidemic dynamics significantly change once we introduce optimizing behavior. Figure 2 shows the epidemic dynamics (left panels) and contour plots of recommended actions over the state space (right panels). The top and bottom panels are for the Markov equilibrium and efficient allocations, respectively.

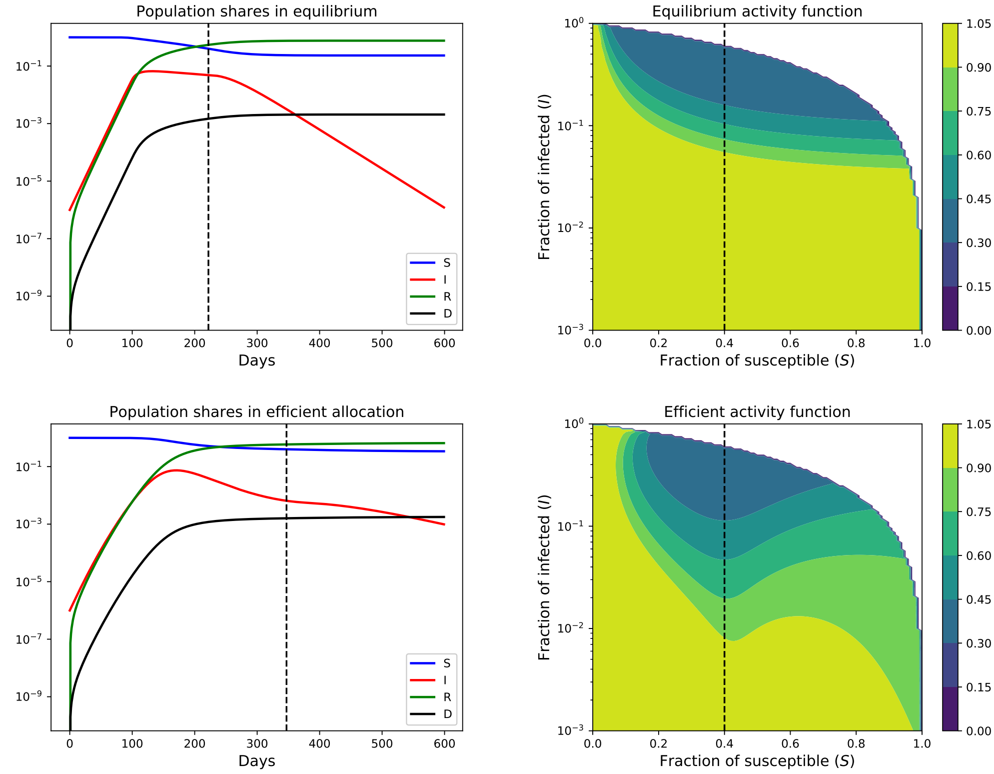

Figure 2: Epidemic dynamics and recommended actions.

We now make a few observations regarding efficient and equilibrium allocations. First, when agents are forward-looking, the epidemic has three phases: during the first phase, the disease spreads freely and we see exponential growth (peak prevalence is 6.62percent); during the second phase, unknown agents voluntarily practice social distancing (set $a_U^* < 1$) and the spread of the disease is endogenously mitigated; during the third phase, society achieves herd immunity and the fraction of infected agents declines exponentially. Second, the epidemic dynamics and recommended action for the solution to the planner's problem (bottom panels) are qualitatively different from the Markov equilibrium. The planner that can enforce and commit to a lockdown policy finds it optimal to substantially reduce the action of unknown agents (bottom right). In the resulting dynamics (bottom left), the fraction of infected agents both grows and declines more slowly, and it takes almost 50 percent longer to reach herd immunity. Figure 3 plots the time paths of recommended actions (equilibrium and efficient) corresponding to the dynamics in the left panels of Figure 2. As expected, we see that the efficient action is almost everywhere below that of the equilibrium action, implying that the planner wishes to curb activity more quickly and for longer than do the individual agents. Further, compared with the equilibrium allocation, the reduction in activity in the efficient allocation begins earlier, is more gradual, and extends well beyond the date at which activity has returned to normal in the equilibrium allocation.

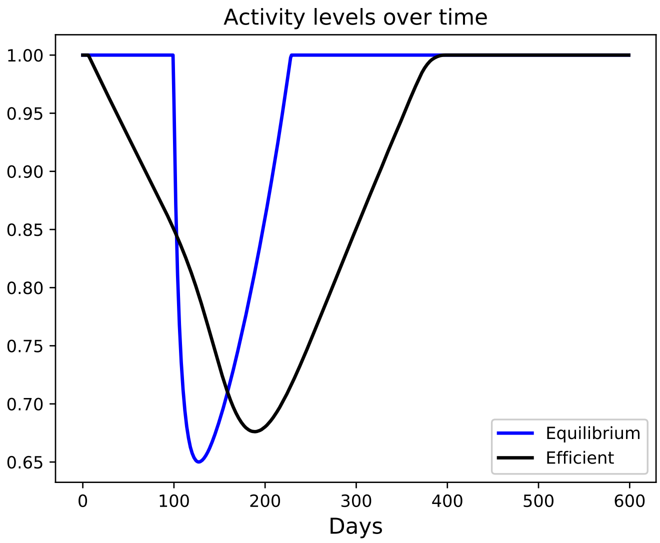

Figure 3: Recommended action over time.

However, as we noted in the introduction, the efficient and equilibrium time paths plotted in Figure 3 do not necessarily provide adequate guidance to a government that was either slow to react to the initial infection or one that has limited ability to carry out its desired intervention. We emphasize that Figure 3 ought to be interpreted as plotting the activity levels across two fictitious countries: one that pursued no non-pharmaceutical interventions at all, and one that followed the optimal path from the outset of the virus. We therefore believe it is instructive to complement the foregoing and examine how the recommendations of governments with varying degrees of enforcement ability vary along the *same* path for the state variables.

To analyze this point and to illustrate the theoretical analysis in Section 3.3, Figure 4 plots the equilibrium level of activity together with the static efficient action ($\lambda = \infty$) and the action the planner would recommend ($\lambda = 0$) along the equilibrium path. As Theorem 3.9 suggests, the static efficient path is everywhere below that of the equilibrium path, and so a government who can only enforce lockdowns for a short period of time would unambiguously wish to do so. However, it does not follow that a government with the ability to perfectly enforce activity levels in perpetuity would wish to reduce activity. Indeed, in this example the difference between the efficient and equilibrium activity levels cannot be unambiguously signed, for at some points in the middle of the pandemic the planner's activity choice exceeds that of the agents in equilibrium. We interpret this last observation to illustrate that relative to the efficient allocation, the competitive equilibrium exhibits inefficiently volatile consumption, with abrupt changes over the time that the planner may wish to avoid with more gradual increases and decreases in activity.

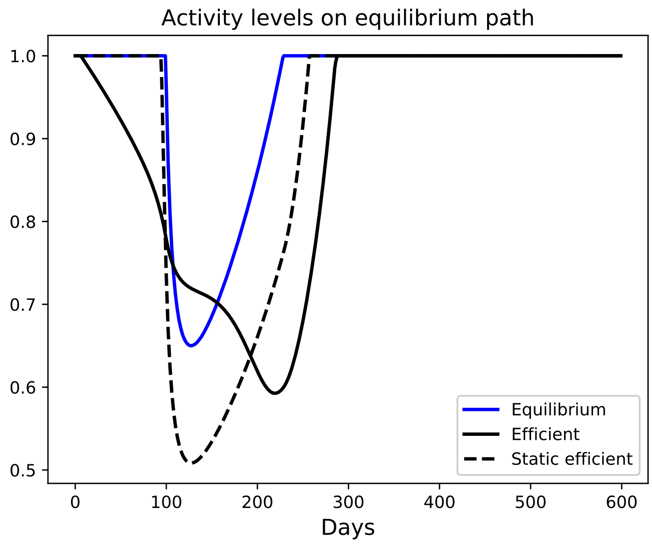

Figure 4: Recommended action on equilibrium path.

4.3 The importance of unknown infected agents

We have so far assumed that unknown infected agents account for $1 - \sigma = 60$ percent of all infected agents. To illustrate the role of imperfect testing and reporting (and to therefore relate our results to the existing literature), we now investigate disease dynamics for the case in which $\sigma = 1$. Figure 5 shows the epidemic dynamics (left panels), contour plots of recommended actions over the state space (right panels), and the recommended actions over time (bottom panel). The top panels are for the Markov equilibrium and the bottom panels are for the efficient actions.

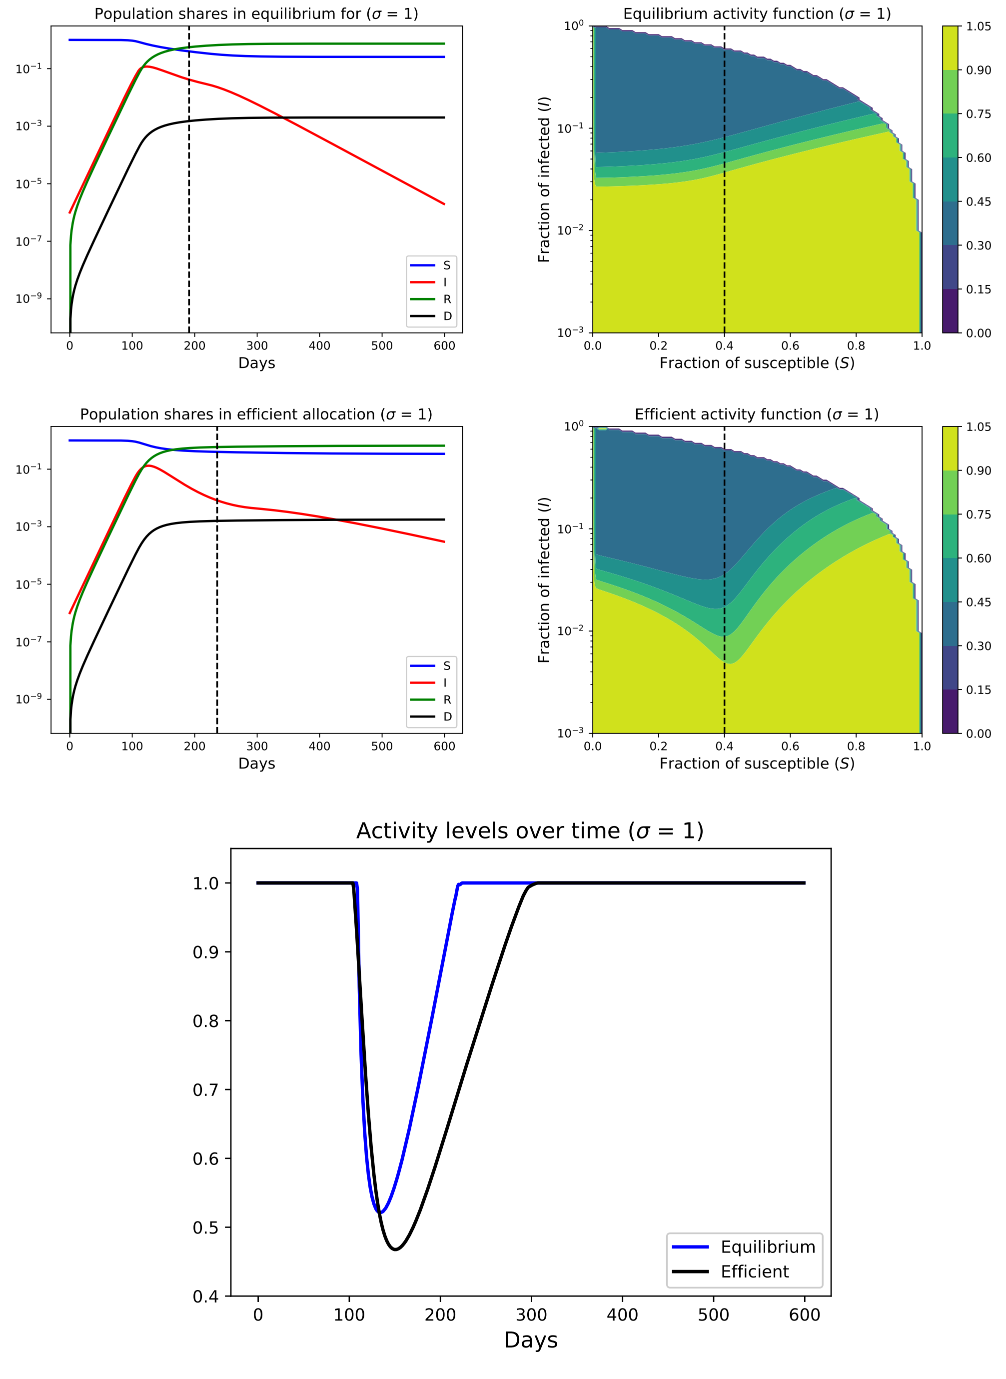

Figure 5: Epidemic dynamics and recommended actions with perfect reporting ($\sigma = 1$).

When $\sigma = 1$, all infected agents are known, and by assumption they all take action $a_I = 1$. Because the average action of infected agents is higher, the negative externality from infected agents is larger, and so the unknown (in this case, susceptible) agents reduce their actions more relative to the case with $\sigma < 1$. The resulting efficient lockdown policy appears qualitatively different from the case with imperfect testing, with the planner essentially delaying any reduction in activity until the date at which the agents begin social distancing in equilibrium. Further, in this case the efficient and equilibrium paths appear qualitatively similar, except for the fact that the latter is more pronounced in its departure from "normality."

To further our understanding of these comparative statics, we compute the welfare cost and death toll across different specifications for the diagnosis rate σ. For welfare, we use the utilitarian social welfare $\mathcal{W}$ in (3.12) and apply the inverse utility function u^{-1} to convert into units of activity. Since welfare without an epidemic equals the highest action 1, we can compute the welfare cost of the epidemic given the current state $z \in Z$ as

$$C(z) \coloneqq 1 - u^{-1}(\mathcal{W}(z)). \tag{4.2}$$

If we identify activity with current output, we can interpret $C(z)$ as the fraction of *permanent* consumption that society is willing to give up to avoid the epidemic. For the death toll, we use the cumulative death (per 100,000 population) after 600 days conditional on no vaccine arrival. Table 1 shows the results, where the columns labeled "Myopic," "PBE," and "SPP" correspond to the myopic equilibrium (standard SIR model), perfect Bayesian Markov competitive equilibrium, and the social planner's problem (efficient action), respectively. The epidemic has a large welfare cost of about a 1.8 percent reduction in permanent consumption. Interestingly, for all reporting levels examined, the difference in welfare is larger between the equilibrium and myopic allocation than between the equilibrium and efficient allocation, implying that the welfare gains from the optimal lockdown policy is modest.

	Welfare loss (percent)			Death toll (per 100,000)		
σ	Myopic	PBE	SPP	Myopic	PBE	SPP
0.1	2.04	1.80	1.68	242	209	172
0.2	2.04	1.81	1.70	242	208	175
0.4	2.04	1.82	1.74	242	207	178
0.7	2.03	1.83	1.76	242	204	179
1.0	2.04	1.81	1.76	242	201	178

Table 1: Welfare loss and death toll.

4.4 The importance of quarantine policies

The analysis so far assumes that known infected (I_k) agents choose the highest action $a_I = 1$. Although this is unrealistic because infected agents may self-isolate due to incapacitation or altruism, it provides the most conservative analysis (worst-case scenario). As a complementary analysis, we solve the model assuming that known infected agents choose $a_I = 0.4$ yet enjoy the highest possible utility ($u_I = 0$). The choice of $a_I = 0.4$ is motivated by the empirical study of He et al. (2020), who document that 44 percent of secondary cases were infected during the presymptomatic stage (prior to diagnosis) of the primary cases. This assumption corresponds to the immediate isolation of the infected cases upon diagnosis, which provides the best-case analysis.

Table 2 shows the welfare cost and death toll with maximal quarantine. In each case, the welfare gains from quarantine are substantial even with a relatively low diagnosis rate σ. In contrast, the additional welfare gains from the optimal lockdown policy are modest. As is intuitive, the difference between the high and low quarantine allocations is particularly stark when the diagnosis rate is high. For $\sigma = 1$, individuals take no precautions in either the efficient or equilibrium allocations, since the level of infections necessary to reach herd immunity is very small.

	Welfare loss (percent)			Death toll (per 100,000)		
σ	Myopic	PBE	SPP	Myopic	PBE	SPP
0.1	1.94	1.69	1.53	236	204	171
0.2	1.80	1.57	1.41	229	197	163
0.4	1.45	1.28	1.11	208	182	143
0.7	0.65	0.64	0.52	149	148	102
1.0	0.00	0.00	0.00	0	0	0

Table 2: Welfare loss and death toll with quarantine.

4.5 The importance of vaccine arrival

Our model has so far assumed that a vaccine is expected to arrive within one year. Consequently, it is likely that the vaccine will have arrived before the attainment of herd immunity. How much does this assumption affect the optimal policies? Should we expect qualitatively similar features to emerge when a vaccine is unlikely to be forthcoming? As an extreme example, we solve for the equilibrium and efficient actions when the vaccine arrives very far into the future ($T = 100$ years).

Figure 6 shows the epidemic dynamics (left panels), contour plots of recommended actions over the state space (right panels), and the recommended action over time (bottom panel).

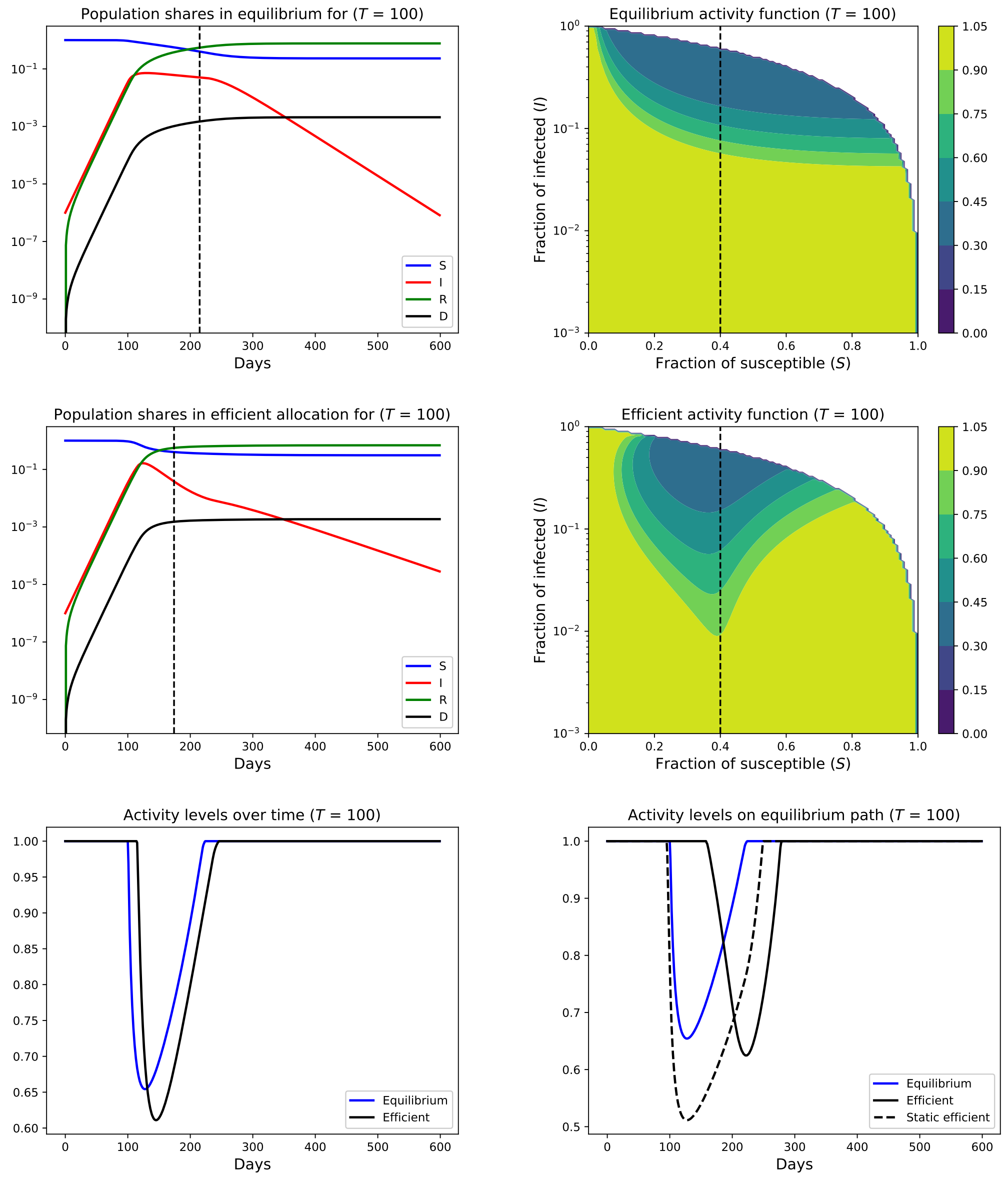

Figure 6: Epidemic dynamics and recommended actions with $T = 100$.

Compared to the case with vaccines (Figure 2, where $\nu = 1/365.25$), the equilibrium dynamics and action (top panels) without vaccines are essentially identical. This is because the possibility of vaccine arrival changes the argument of the individually optimal action (3.8) by a factor of $\mathrm{e}^{-\nu\Delta}$, which is close to 1 even if ν is large such as $1/365.25$. However, the possibility of vaccine arrival significantly affects the optimal lockdown policy. Comparing the efficient actions in the bottom right panels of Figures 2 and 6, we can see that when a vaccine is not expected, the planner *encourages* (discourages) agents to take high action *before* (after) achieving herd immunity, similar to the case with perfect reporting (Figure 5). Although this result may be counterintuitive, the intuition is actually straightforward. In the absence of a vaccine, the only way to end the epidemic is to achieve herd immunity. Therefore the planner wishes to initially exacerbate the epidemic to reach herd immunity quickly (thus reducing the cost of low action) and then mitigate the epidemic after herd immunity to reduce unnecessary deaths. We can clearly see this from the bottom right panel of Figure 6, where we plot the recommended action along the equilibrium path (not over time), meaning that we are plotting the recommended actions at the same points in the state space.

	Welfare loss (percent)			Death toll (per 100,000)		
σ	Myopic	PBE	SPP	Myopic	PBE	SPP
0.1	2.85	2.84	2.59	242	209	185
0.2	2.85	2.83	2.59	242	209	185
0.4	2.84	2.81	2.59	242	207	186
0.7	2.83	2.74	2.57	242	205	186
1.0	2.84	2.63	2.53	242	202	183

Table 3: Welfare loss and death toll with late vaccine arrival.

Table 3 shows the welfare cost and death toll when the vaccine is expected to take $T = 100$ years to arrive. As expected, the welfare cost of the epidemic is higher since it takes longer for the vaccine to arrive. More interestingly, the bottom panels of Figures 5 and 6 show that the optimal policy is to *encourage* (discourage) agents to take high action *before* (after) achieving herd immunity. Our intuition for this is that when it will be a long time until a vaccine arrives, the planner wishes to reduce abrupt fluctuations in activity and also ensure that herd immunity is not excessively "overshot" in order to minimize deaths. Figure 7 reinforces this intuition and shows that the share of susceptible agents is reduced more rapidly and plateaus more smoothly in the efficient allocation, so as to effect a "smooth landing" out of the pandemic.

5 Concluding remarks

In this paper we have theoretically studied optimal epidemic control in an equilibrium model with imperfect testing and enforcement. We proved the existence of a perfect Bayesian Markov competitive equilibrium, and showed that a government

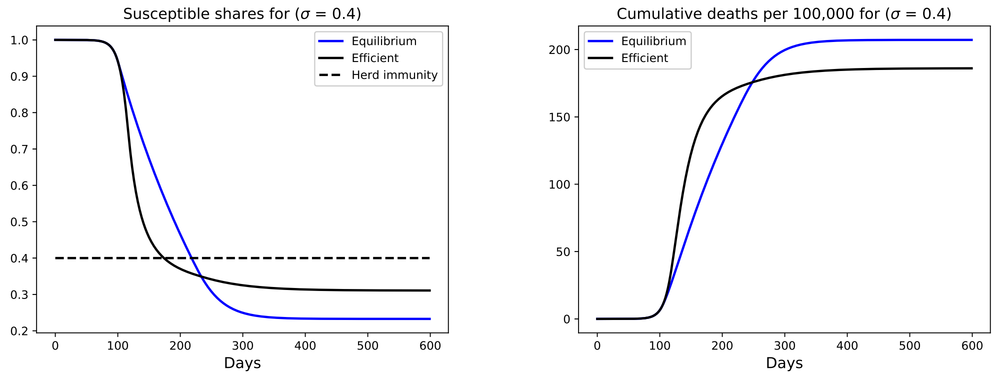

Figure 7: Efficient and equilibrium susceptible shares and deaths with $T = 100$.

that can only enforce lockdowns for a short period of time will wish to (weakly) reduce activity, but that its incentive to do vanishes if testing is perfect. We then showed through numerical examples that optimal policy is highly dependent upon the government's ability to enforce social distancing over the long term, and that if a vaccine is not expected to ever arrive, the optimal activity recommended by the government may *exceed* the equilibrium activity levels until the population of susceptible agents is sufficiently low. In contrast, when a vaccine is expected to arrive within a few years (as could reasonably be expected at the outset of the COVID-19 epidemic), the optimal policy typically involves an immediate and prolonged enforcement of social activity.

References

Abel, A. and Panageas, S. Social Distancing, Vaccination and the Paradoxical Optimality of an Endemic Equilibrium. Technical Report w27742, National Bureau of Economic Research, Cambridge, MA, August 2020.

Alvarez, F., Argente, D., and Lippi, F. A Simple Planning Problem for COVID-19 Lockdown. Technical Report w26981, National Bureau of Economic Research, Cambridge, MA, April 2020.

Auld, M. Choices, beliefs, and infectious disease dynamics. *Journal of Health Economics*, 22(3):361–377, May 2003. ISSN 01676296. doi:10.1016/S0167-6296(02)00103-0.

Budish, E. Maximize Utility subject to R≤1: A Simple Price-Theory Approach to Covid-19 Lockdown and Reopening Policy. Technical Report w28093, National Bureau of Economic Research, Cambridge, MA, November 2020.

Chen, F. A mathematical analysis of public avoidance behavior during epidemics using game theory. *Journal of Theoretical Biology*, 302:18–28, June 2012. ISSN 00225193. doi:10.1016/j.jtbi.2012.03.002.

Chen, F., Jiang, M., Rabidoux, S., and Robinson, S. Public avoidance and epidemics: Insights from an economic model. *Journal of Theoretical Biology*, 278 (1):107–119, June 2011. ISSN 00225193. doi:10.1016/j.jtbi.2011.03.007.

Dong, E., Du, H., and Gardner, L. An interactive web-based dashboard to track COVID-19 in real time. *The Lancet. Infectious Diseases*, 20(5):533–534, May 2020. ISSN 1474-4457. doi:10.1016/S1473-3099(20)30120-1.

Eksin, C., Paarporn, K., and Weitz, J. S. Systematic biases in disease forecasting – The role of behavior change. *Epidemics*, 27:96–105, June 2019. ISSN 17554365. doi:10.1016/j.epidem.2019.02.004.

Farboodi, M., Jarosch, G., and Shimer, R. Internal and External Effects of Social Distancing in a Pandemic. Technical Report w27059, National Bureau of Economic Research, Cambridge, MA, April 2020.

Fenichel, E. P., Castillo-Chavez, C., Ceddia, M. G., Chowell, G., Parra, P. A. G., Hickling, G. J., Holloway, G., Horan, R., Morin, B., Perrings, C., Springborn, M., Velazquez, L., and Villalobos, C. Adaptive human behavior in epidemiological models. *Proceedings of the National Academy of Sciences*, 108(15):6306–6311, April 2011. ISSN 0027-8424, 1091-6490. doi:10.1073/pnas.1011250108.

Fenichel, E. P. Economic considerations for social distancing and behavioral based policies during an epidemic. *Journal of Health Economics*, 32(2):440–451, March 2013. ISSN 01676296. doi:10.1016/j.jhealeco.2013.01.002.

Geoffard, P.-Y. and Philipson, T. Rational Epidemics and Their Public Control. *International Economic Review*, 37(3):603, August 1996. ISSN 00206598. doi:10.2307/2527443.

Hall, R. E., Jones, C. I., and Kleneow, P. J. Trading Off Consumption and COVID-19 Deaths. *Quarterly Review*, 42(1), June 2020. ISSN 02715287. doi:10.21034/qr.4211.

He, X., Lau, E. H. Y., Wu, P., Deng, X., Wang, J., Hao, X., Lau, Y. C., Wong, J. Y., Guan, Y., Tan, X., Mo, X., Chen, Y., Liao, B., Chen, W., Hu, F., Zhang, Q., Zhong, M., Wu, Y., Zhao, L., Zhang, F., Cowling, B. J., Li, F., and Leung,

G. M. Temporal dynamics in viral shedding and transmissibility of COVID-19. *Nature Medicine*, 26(5):672–675, May 2020. ISSN 1078-8956, 1546-170X. doi:10.1038/s41591-020-0869-5.

Ioannidis, J. P. A. Infection fatality rate of COVID-19 inferred from seroprevalence data. *Bulletin of the World Health Organization*, 99(1):19–33F, January 2021. ISSN 0042-9686. doi:10.2471/BLT.20.265892.

Kermack, W. O. and McKendrick, A. G. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 115(772):700–721, August 1927. ISSN 0950-1207, 2053-9150. doi:10.1098/rspa.1927.0118.

Kremer, M. Integrating Behavioral Choice into Epidemiological Models of AIDS. *The Quarterly Journal of Economics*, 111(2):549–573, May 1996. ISSN 0033-5533, 1531-4650. doi:10.2307/2946687.

Kruse, T. and Strack, P. Optimal Control of an Epidemic through Social Distancing. *SSRN Electronic Journal*, 2020. ISSN 1556-5068. doi:10.2139/ssrn.3581295.

Kushner, H. J. and Dupuis, P. G. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer US, New York, NY, 1992. ISBN 978-1-4684-0443-2 978-1-4684-0441-8. doi:10.1007/978-1-4684-0441-8.

Mizumoto, K., Kagaya, K., Zarebski, A., and Chowell, G. Estimating the asymptomatic proportion of coronavirus disease 2019 (COVID-19) cases on board the Diamond Princess cruise ship, Yokohama, Japan, 2020. *Eurosurveillance*, 25(10), March 2020. ISSN 1560-7917. doi:10.2807/1560-7917.ES.2020.25.10.2000180.

Rai, B., Shukla, A., and Dwivedi, L. K. Estimates of serial interval for COVID-19: A systematic review and meta-analysis. *Clinical Epidemiology and Global Health*, 9:157–161, January 2021. ISSN 22133984. doi:10.1016/j.cegh.2020.08.007.

Reluga, T. C. Game Theory of Social Distancing in Response to an Epidemic. *PLoS Computational Biology*, 6(5):e1000793, May 2010. ISSN 1553-7358. doi:10.1371/journal.pcbi.1000793.

Sethi, S. P. Optimal Quarantine Programmes for Controlling an Epidemic Spread. *Journal of the Operational Research Society*, 29(3):265–268, March 1978. ISSN 0160-5682, 1476-9360. doi:10.1057/jors.1978.55.

Toxvaerd, F. RATIONAL DISINHIBITION AND EXTERNALITIES IN PREVENTION. *International Economic Review*, 60(4):1737–1755, November 2019. ISSN 0020-6598, 1468-2354. doi:10.1111/iere.12402.

Toxvaerd, F. Equilibrium Social Distancing. March 2020. doi:10.17863/CAM.52489.

You, C., Deng, Y., Hu, W., Sun, J., Lin, Q., Zhou, F., Pang, C. H., Zhang, Y., Chen, Z., and Zhou, X.-H. Estimation of the time-varying reproduction number of COVID-19 outbreak in China. *International Journal of Hygiene and Environmental Health*, 228:113555, July 2020. ISSN 14384639. doi:10.1016/j.ijheh.2020.113555.

A Proofs

To prove Proposition 3.1, we note the following simple lemma.

Lemma A.1. *Let X be a complete metric space and $T : X \to X$ be a contraction with fixed point $x^* \in X$. If $\emptyset \neq X_1 \subset X$ is closed and $TX_1 \subset X_1$, then $x^* \in X_1$.*

Proof. Take any $x_0 \in X_1$ and define $x_n = T^n x_0$ for $n \in \mathbb{N}$. Since by assumption $TX_1 \subset X_1$, by induction we have $x_n \in X_1$ for all n. By the contraction mapping theorem, we have $x_n \to x^*$. Since X_1 is closed, we have $x^* \in X_1$. □

Proof of Proposition 3.1. Let $X = \mathbb{R}^Z$. For $V \in X$, define $TV(z)$ by the right-hand side of (3.5), where $V_U = V$. Since $A = [\underline{a}, 1]$ is nonempty and compact, $u : A \to \mathbb{R}$ is continuous, and Z is a finite set, the maximum is achieved and $T : X \to X$ is well-defined. Let us show that T is a contraction by verifying Blackwell's sufficient conditions. If $V_1, V_2 \in X$ and $V_1 \leq V_2$ pointwise, since $0 \leq \sigma\mu pa \leq 1$, it follows from the definition of T that

$$\begin{aligned}(TV_1)(z) =& \max_{a\in A}\left\{(1-\mathrm{e}^{-r\Delta})u(a) + \mathrm{e}^{-r\Delta}\,\mathrm{E}_z((1-\sigma\mu pa)\mathrm{e}^{-\nu\Delta}V_1(z') + \sigma\mu pa V_{I_k})\right\}\\ \leq& \max_{a\in A}\left\{(1-\mathrm{e}^{-r\Delta})u(a) + \mathrm{e}^{-r\Delta}\,\mathrm{E}_z((1-\sigma\mu pa)\mathrm{e}^{-\nu\Delta}V_2(z') + \sigma\mu pa V_{I_k})\right\}\\ =&(TV_2)(z),\end{aligned}$$

so T is monotone. If $V \in X$ and $c \geq 0$ is any constant, then

$$\begin{aligned}(T(V+c))(z) &= \max_{a \in A}\{(1-\mathrm{e}^{-r\Delta})u(a) + \mathrm{e}^{-r\Delta}\,\mathrm{E}_z((1-\sigma\mu pa)\mathrm{e}^{-\nu\Delta}(V(z')+c) + \sigma\mu pa V_{I_k})\} \\ &\leq \max_{a \in A}\{(1-\mathrm{e}^{-r\Delta})u(a) + \mathrm{e}^{-r\Delta}\,\mathrm{E}_z((1-\sigma\mu pa)\mathrm{e}^{-\nu\Delta}V(z') + \sigma\mu pa V_{I_k})\} + \mathrm{e}^{-(r+\nu)\Delta}c \\ &= (TV)(z) + \mathrm{e}^{-(r+\nu)\Delta}c,\end{aligned}$$

so the discounting property holds with modulus $\mathrm{e}^{-(r+\nu)\Delta} < 1$. Therefore T is a contraction mapping and has a unique fixed point $V_U \in X$. Because $u : A = [\underline{a}, 1] \to \mathbb{R}$ is continuous (hence bounded), it is straightforward to verify the transversality condition. Therefore V_U is the value function.

Next, let us show the inequalities (3.6). Noting that $u_D < u_I \leq 0$, $r > 0$, $\gamma > 0$, $\delta \in (0, 1]$, and V_{I_k} is given by (3.2), we have

$$V_D = u_D < \frac{(\mathrm{e}^{r\Delta}-1)u_I + \gamma\Delta\delta u_D}{\mathrm{e}^{r\Delta}-1+\gamma\Delta} = \frac{(1-\mathrm{e}^{-r\Delta})u_I + \mathrm{e}^{-r\Delta}\gamma\Delta\delta u_D}{1-\mathrm{e}^{-r\Delta}(1-\gamma\Delta)} = V_{I_k} < 0.$$

To obtain the lower bound for V_U, Define the constants

$$c_1 := \frac{\sigma\beta\Delta}{\mathrm{e}^{(r+\nu)\Delta}-1+\sigma\beta\Delta} \in (0,1) \tag{A.1}$$

and $V_1 := c_1 \mathrm{e}^{\nu\Delta} V_{I_k}$. Note that

$$\begin{aligned}c_1\mathrm{e}^{\nu\Delta} < 1 &\iff \frac{\mathrm{e}^{\nu\Delta}\sigma\beta\Delta}{\mathrm{e}^{(r+\nu)\Delta}-1+\sigma\beta\Delta} < 1 \\ &\iff \mathrm{e}^{\nu\Delta}\sigma\beta\Delta < \mathrm{e}^{(r+\nu)\Delta}-1+\sigma\beta\Delta \\ &\iff (\mathrm{e}^{\nu\Delta}-1)\sigma\beta\Delta < \mathrm{e}^{(r+\nu)\Delta}-1,\end{aligned}$$

which is true because $\sigma \in (0, 1]$, $\beta\Delta \in [0, 1]$, $r > 0$, and $\nu \geq 0$. Since $c_1\mathrm{e}^{\nu\Delta} < 1$ and $V_{I_k} < 0$, it follows that

$$V_{I_k} < c_1\mathrm{e}^{\nu\Delta}V_{I_k} = V_1 = \frac{\mathrm{e}^{\nu\Delta}\sigma\beta\Delta}{\mathrm{e}^{(r+\nu)\Delta}-1+\sigma\beta\Delta}V_{I_k},$$

which is part of (3.6). Let us next show $V_U \geq V_1$. To this end, define the (nonempty closed) subset of X by $X_1 = \{V \in X \mid V \geq V_1\}$. By Lemma A.1, to show $V_U \geq V_1$, it suffices to show $TX_1 \subset X_1$. If $V \in X_1$, then by the definition of

T, we have

$$\begin{aligned}(TV)(z) &= \max_{a\in A}\left\{(1-\mathrm{e}^{-r\Delta})u(a)+\mathrm{e}^{-r\Delta}\,\mathrm{E}_z((1-\sigma\mu pa)\mathrm{e}^{-\nu\Delta}V(z')+\sigma\mu paV_{I_k})\right\}\\ &\ge \max_{a\in A}\left\{(1-\mathrm{e}^{-r\Delta})u(a)+\mathrm{e}^{-r\Delta}\,\mathrm{E}_z((1-\sigma\mu pa)c_1V_{I_k}+\sigma\mu paV_{I_k})\right\}. \quad (\mathrm{A}.2)\end{aligned}$$

Setting $a=1$ in the right-hand side of (A.2) and noting that $u(1)=0$, we obtain

$$(TV)(z)\ge \mathrm{e}^{-r\Delta}((1-\sigma\mu p)c_1+\sigma\mu p)V_{I_k}.$$

Therefore to show $TX_1\subset X_1$, noting that $V_{I_k}<0$, it suffices to show that

$$\mathrm{e}^{-r\Delta}((1-\sigma\mu p)c_1+\sigma\mu p)\le c_1\mathrm{e}^{\nu\Delta}. \quad (\mathrm{A}.3)$$

By (3.3), $a_U(z)\le 1$, and the definition of Z in (2.1), we have

$$p=\beta\Delta(a_II_k+a_U(z)I_u)\le\beta\Delta.$$

Since $\mu\in[0,1]$ and $c_1\in(0,1)$, it follows that

$$\mathrm{e}^{-r\Delta}((1-\sigma\mu p)c_1+\sigma\mu p)\le \mathrm{e}^{-r\Delta}((1-\sigma\beta\Delta)c_1+\sigma\beta\Delta)=c_1\mathrm{e}^{\nu\Delta}$$

by (A.1), which implies (A.3). Therefore $TX_1\subset X_1$, as desired.

To show $V_U(z)\le V_{R_k}=0$, define the (nonempty closed) subset of X by $X_2:=\{V\in X\mid V\le 0\}$. If $V\in X_2$,

$$\begin{aligned}(TV)(z) = \max_{a\in A}&\{(1-\mathrm{e}^{-r\Delta})u(a)\\ &+\mathrm{e}^{-r\Delta}\,\mathrm{E}_z((1-\sigma\mu pa)\mathrm{e}^{-\nu\Delta}V(z')+\sigma\mu paV_{I_k})\}\le 0\end{aligned}$$

because $u(a)\le u(1)=0$, $V\le 0$, and $V_{I_k}<0$. Therefore $TX_2\subset X_2$, which implies $V_U\le 0$ by Lemma A.1. □

Proof of Proposition 3.2. Let

$$f(a):=(1-\mathrm{e}^{-r\Delta})u(a)+\mathrm{e}^{-(r+\nu)\Delta}\,\mathrm{E}_z((1-\sigma\mu pa)V_U(z')+\sigma\mu paV_{I_k})$$

be the expression inside the braces in (3.5). Since by Assumption 2 individual agents view the next period's state z' as exogenous, $V_U(z')$ does not depend on a

and f is strictly concave. Furthermore,

$$f'(a) = (1 - \mathrm{e}^{-r\Delta})u'(a) - \mathrm{e}^{-r\Delta}\sigma\mu p(\mathrm{E}_z\,\mathrm{e}^{-\nu\Delta}V_U(z') - V_{I_k}).$$

Hence considering the cases $f'(\underline{a}) \gtrless 0$ and $f'(1) \gtrless 0$, the optimal action is given by (3.8). $\square$

Proof of Corollary 3.3. Since $0 \geq V_U > V_{I_k}$ by (3.6) and $\mathrm{e}^{-\nu\Delta} \leq 1$, we always have $\mathrm{E}_z\,\mathrm{e}^{-\nu\Delta}V_U(z') - V_{I_k} > 0$. By (3.6) and (3.8), the best response can be rewritten as $a^* = \phi(x)$, where

$$x := \frac{\sigma\mu(z)p(z)}{\mathrm{e}^{r\Delta} - 1}\left(\mathrm{E}_z\,\mathrm{e}^{-\nu\Delta}V_U(z') - \frac{(\mathrm{e}^{r\Delta} - 1)u_I + \gamma\Delta\delta u_D}{\mathrm{e}^{r\Delta} - 1 + \gamma\Delta}\right) > 0. \tag{A.4}$$

Since $r > 0$, $u_D < u_I \leq 0$, and $V_U \leq 0$, it is straightforward to show that x is decreasing in u_I, u_D and increasing in $\nu, \mu(z), p(z), \delta$. Therefore the desired comparative statics results hold because ϕ is decreasing. $\square$

Proof of Corollary 3.4. Let x in (A.4) be the argument of ϕ in (3.8). Using (2.5), (3.2), and the bound $V_U \leq 0$ in (3.6), we can bound x from above as

$$x \leq \frac{\sigma\mu(z)p(z)}{\mathrm{e}^{r\Delta} - 1}(-V_{I_k}). \tag{A.5}$$

Using $\mu(z) \leq 1$, $a_U(z) \leq 1$, and

$$p(z) = \beta\Delta(a_I I_k + a_U(z)I_u) \leq \beta\Delta(I_k + I_u) = \beta\Delta I$$

by (3.3), it follows from (A.5) and $V_{I_k} < 0$ that

$$x \leq -\frac{\sigma\beta\Delta}{\mathrm{e}^{r\Delta} - 1}IV_{I_k}. \tag{A.6}$$

Therefore if $I \leq \bar{I}$, where $\bar{I}$ is as in (3.9), it follows from (A.6) that $x \leq u'(1)$. Therefore $a^* = \phi(x) = 1$ by (3.7). $\square$

To prove Theorem 3.6, we consider a general stochastic game with complete information and finitely many competitive agents denoted by $n \in N = \{1, \dots, N\}$. Suppose that there are finitely many agent types indexed by $h \in H = \{1, \dots, H\}$. At each point in time, the aggregate state of the economy is denoted by $z \in Z$, which is a finite set. A type h agent takes action $x \in A_h$. Let $A = \bigotimes_{h=1}^{H} A_h$.

Let $u_h(x, a, z)$ be the flow utility of a type h agent when the agent takes action $x \in A_h$, the average action of other agents is $a \in A$, and the aggregate state is $z \in Z$. Type h agents discount future utility with discount factor $\rho_h \in [0, 1)$. Time is denoted by $t = 0, 1, \dots$. The aggregate state z evolves stochastically depending on agents' average actions and the current state. Let $q(a, z, z')$ be the probability of $z_{t+1} = z'$ conditional on $(a_t, z_t) = (a, z) \in A \times Z$. When computing $q(a, z, z')$, agents take the average action a as given and ignore the impact of their action x on a (competitive behavior). An agent's type evolves stochastically depending on an agent's own action, the other agents' average actions, and the current state. Let $p_{hh'}(x, a, z)$ be the probability that a type h agent switches to type h' next period when the state is $z \in Z$, he takes action $x \in A_h$, the average action is $a \in A$, and the current state is $z \in Z$. Value functions are denoted by $V = (V_1, \dots, V_H)$, where $V_h : Z \to \mathbb{R}$. Reaction functions are denoted by $\theta = (\theta_1, \dots, \theta_H)$, where $\theta_h : Z \to A_h$.

Definition A.2. A (pure strategy) *Markov competitive equilibrium* is a pair (V, θ) of value and reaction functions such that

(i) (Consistency) The transition probability of the aggregate state is consistent with individual behavior, so $\Pr(z' \mid z) = q(\theta(z), z, z')$.

(ii) (Sequential rationality) For each $h \in H$ and $z \in Z$, the Bellman equation

$$V_h(z) = \max_{x \in A_h} \left\{ (1 - \rho_h) u_h(x, \theta(z), z) + \rho_h \sum_{h'=1}^{S} p_{hh'}(x, \theta(z), z) \, \mathrm{E}_z \, V_{h'}(z') \right\} \tag{A.7}$$

holds, and $x = \theta_h(z)$ achieves the maximum.

Theorem A.3. *Suppose the following assumptions hold:*

(i) For each $h \in H$, the action set A_h is a nonempty compact convex subset of some Euclidean space.

(ii) For each $h \in H$, the payoff function $u_h : A_h \times A \times Z \to \mathbb{R}$ is continuous in (x, a, z) and strictly concave in x.

(iii) The transition probability $q : A \times Z \times Z \to [0, 1]$ is continuous.

(iv) For each $h, h' \in H$, the transition probability $p_{hh'} : A_h \times A \times Z \to [0, 1]$ is continuous in (x, a, z) and affine in x.

Then there exists a pure strategy Markov competitive equilibrium.

To prove Theorem A.3, we need the following lemma.

Lemma A.4. *Let (X, d) be a complete metric space and Θ be a topological space. Endow $X \times \Theta$ with the product topology. Suppose $T : X \times \Theta \to X$ is continuous and there exists $\beta \in [0, 1)$ such that*

$$d(T(x, \theta), T(y, \theta)) \leq \beta d(x, y) \tag{A.8}$$

for all $x, y \in X$ and all $\theta \in \Theta$. Then the following statements are true:

(i) For each $\theta \in \Theta$, there exists a unique $x^*(\theta) \in X$ such that $T(x^*(\theta), \theta) = x^*(\theta)$.

(ii) $x^* : \Theta \to X$ is continuous.

Proof. The first claim is immediate from the contraction mapping theorem.

To show the second claim, for simplicity write $T_\theta x := T(x, \theta)$. Fix any $(x_0, \theta) \in X \times \Theta$ and define $x_n = T_\theta^n x_0$ for $n \in \mathbb{N}$. Then by the triangle inequality and the contraction property (A.8), we have

$$\begin{aligned} d(x_n, x_0) &\leq d(x_n, x_{n-1}) + \cdots + d(x_1, x_0) \\ &\leq (\beta^{n-1} + \cdots + 1) d(x_1, x_0) \\ &\leq \frac{1}{1 - \beta} d(x_1, x_0) = \frac{1}{1 - \beta} d(T_\theta x_0, x_0). \end{aligned}$$

Letting $n \to \infty$ and noting that $x_n \to x^*(\theta)$ by the contraction mapping theorem, we have

$$d(x^*(\theta), x_0) \leq \frac{1}{1 - \beta} d(T_\theta x_0, x_0). \tag{A.9}$$

Since $x_0 \in X$ and $\theta \in \Theta$ are arbitrary in (A.9), set $\theta = \theta'$ and $x_0 = x^*(\theta)$, where $\theta, \theta' \in \Theta$ are arbitrary. Since by definition $T_\theta x_0 = x_0$, it follows that

$$\begin{aligned} d(x^*(\theta'), x^*(\theta)) &\leq \frac{1}{1 - \beta} d(T_{\theta'} x_0, x_0) = \frac{1}{1 - \beta} d(T_{\theta'} x_0, T_\theta x_0) \\ &= \frac{1}{1 - \beta} d(T(x_0, \theta'), T(x_0, \theta)), \end{aligned} \tag{A.10}$$

where $x_0 = x^*(\theta)$. Since T is continuous and $x_0 = x^*(\theta)$ depends only on θ, for any $\epsilon > 0$ there exists an open neighborhood U of θ such that

$$d(T(x_0, \theta'), T(x_0, \theta)) < (1 - \beta)\epsilon \tag{A.11}$$

whenever $\theta' \in U$. Combining (A.10) and (A.11), for all $\theta' \in U$ we have

$$d(x^*(\theta'), x^*(\theta)) < \epsilon,$$

so $x^* : \Theta \to X$ is continuous. $\square$

Proof of Theorem A.3. Let $\mathcal{V} = bc(H \times Z)$ be the space of bounded continuous functions from $H \times Z$ to $\mathbb{R}$ equipped with the supremum norm. Then $(\mathcal{V}, \|\cdot\|)$ is a Banach space. (Indeed, it is just the Euclidean space $\mathbb{R}^{HZ}$ because H, Z are finite sets.) For $(V, \theta) \in \mathcal{V} \times A^Z$, define

$$(T_\theta V)_h(z) = \max_{x \in A_h} \left\{ (1 - \rho_h) u_h(x, \theta(z), z) + \rho_h \sum_{h', z'} p_{hh'}(x, \theta(z), z) q(\theta(z), z, z') V_{h'}(z') \right\}. \quad \text{(A.12)}$$

A straightforward application of the maximum theorem implies that $(T_\theta V)_h(z)$ is continuous in θ. Since H is a finite set, we have $\rho := \max_h \rho_h \in [0, 1)$. It is then straightforward to verify Blackwell's sufficient conditions, and for fixed $\theta \in A^Z$, the map $\mathcal{V} \ni V \mapsto T_\theta V$ is a contraction mapping with modulus ρ. It follows from Lemma A.4 that T_θ has a unique fixed point

$$V^*(\theta) = \{V_h^*(z, \theta)\}_{(h,z) \in H \times Z} \in \mathcal{V},$$

which is continuous in θ.

Since by assumption $u_h(x, a, z)$ is continuous and strictly concave in x and $p_{hh'}(x, a, z)$ is affine in x, the objective function inside the braces of (A.12) (where $V_h(z) = V_h^*(z, \theta)$) is continuous and strictly concave in x. Therefore there exists a unique maximizer, which we call $x = a_h^*(z, \theta)$. By the maximum theorem, $a_h^*(z, \theta)$ is continuous in θ. Since H, Z are finite sets, we may view $a^* : A^Z \to A^Z$ as a continuous map. Since A is nonempty, compact, and convex, by Brouwer's fixed point theorem, a^* has a fixed point. Letting $\theta = (\theta_h(z))$ be this fixed point, it is clear that all conditions in Definition A.2 are satisfied. $\square$

Proof of Theorem 3.6. We apply Theorem A.3.

The agent type is denoted by $h \in H = \{U, I_k, R_k, D\}$, which is finite. The aggregate state is denoted by $z \in Z$ defined in (2.1), which is finite. Define the action set of type h agents by $A_h = [\underline{a}, 1]$ if $h \neq D$ and $A_D = \{0\}$, which are nonempty compact convex subsets of $\mathbb{R}$. Define the utility function $u_h : A_h \times A \times$

$Z \to \mathbb{R}$ by

$$\begin{aligned} u_U(x,a,z) &= u_{R_k}(x,a,z) = u(x), \\ u_{I_k}(x,a,z) &= u_I(x), \\ u_D(x,a,z) &= u_D. \end{aligned}$$

Note that each u_h is continuous in (x, a, z) and strictly concave in $x \in A_h$. (Although u_D is a constant function, it is strictly concave in x because its domain $A_D = \{0\}$ is a singleton.)

Define the transition probabilities of individual states $p_{hh'}(x, a, z)$ as follows. For $h = R_k, D$ agents, because they remain in their corresponding state forever, the transition probabilities are 0 or 1, which are clearly continuous in all variables and affine in x. For I_k agents, by assumption they die with probability $\gamma\Delta\delta$. Hence $p_{I_kD}(x, a, z) = \gamma\Delta\delta$, which is continuous in all variables and constant (hence affine) in x. The same is true for $p_{Ih'}$ for any $h' \in H$. By (2.2) and (3.5), the transition probabilities of U agents are

$$\begin{aligned} p_{UU}(x,a,z) &= \mathrm{e}^{-\nu\Delta}(1 - \sigma\mu(z)\beta\Delta(a_{I_k}I_k + a_U I_u)x), \\ p_{UI_k}(x,a,z) &= \sigma\mu(z)\beta\Delta(a_{I_k}I_k + a_U I_u)x, \\ p_{UR_k}(x,a,z) &= (1 - \mathrm{e}^{-\nu\Delta})(1 - \sigma\mu(z)\beta\Delta(a_{I_k}I_k + a_U I_u)x), \\ p_{UD}(x,a,z) &= 0, \end{aligned}$$

which are continuous in all variables and affine in x. Finally, the transition probability for the aggregate state $q : A \times Z \times Z \to [0, 1]$ is clearly continuous because it is determined by the current state $z \in Z$ (which is finite) and individual's actions, whose transition probabilities are all continuous. The existence of equilibrium in the sense of Definition A.2 then follows from Theorem A.3. The resulting value and policy functions clearly satisfy Definition 3.5. Because U agents take identical actions, their beliefs always satisfy (2.6). $\square$

We need the following lemma to prove Theorem 3.9.

Lemma A.5. *Let $I \subset \mathbb{R}$ be an interval, $f : I \to \mathbb{R}$ continuously differentiable with $f' \neq 0$ on I, and $g : I \to \mathbb{R}$ continuous. If $f(a) = g(b)$ for some $a, b \in I$, then*

$$|a - b| \leq \frac{|f(b) - g(b)|}{\min_{x \in [a,b]} |f'(x)|}.$$

Proof. By the mean value theorem, there exists $\theta \in (0,1)$ such that

$$f(b) - f(a) = f'(c)(b-a),$$

where $c = (1-\theta)a + \theta b$. Since by assumption $f(a) = g(b)$, we obtain

$$|f(b) - g(b)| = |f(b) - f(a)| = |f'(c)(b-a)| \geq m\,|a-b|\,,$$

where $m = \min_{x\in[a,b]} |f'(x)| > 0$ since f' is continuous and $f' \neq 0$ on I. Dividing both sides by $m > 0$, we obtain the desired result. $\square$

Proof of Theorem 3.9. Fix $z \in Z$ and let $\mu = \mu(z)$ and $a_U^* = a_U^*(z)$. Define the functions $f, g : A \to \mathbb{R}$ by

$$\begin{aligned}
f(x) &= (1-\mathrm{e}^{-r\Delta})u(x) + \mathrm{e}^{-r\Delta}\,\mathrm{E}_z((1-\sigma\mu p(a_U^*)x)\mathrm{e}^{-\nu\Delta}V_U(z') + \sigma\mu p(a_U^*)xV_{I_k}(z')),\\
g(x) &= (1-\mathrm{e}^{-r\Delta})u(x) + \mathrm{e}^{-r\Delta}\,\mathrm{E}_z((1-\sigma\mu p(x)x)\mathrm{e}^{-\nu\Delta}V_U(z') + \sigma\mu p(x)xV_{I_k}(z')),
\end{aligned}$$

where (with a slight abuse of notation) $p(x) := \beta\Delta(a_I I_k + xI_u)$. Taking the derivatives of f, g, we obtain

$$\begin{aligned}
f'(x) &= (1-\mathrm{e}^{-r\Delta})u'(x) - \mathrm{e}^{-r\Delta}\sigma\beta\Delta\mu(a_I I_k + a_U^* I_u)\,\mathrm{E}_z(\mathrm{e}^{-\nu\Delta}V_U(z') - V_{I_k}),\\
g'(x) &= (1-\mathrm{e}^{-r\Delta})u'(x) - \mathrm{e}^{-r\Delta}\sigma\beta\Delta\mu(a_I I_k + 2xI_u)\,\mathrm{E}_z(\mathrm{e}^{-\nu\Delta}V_U(z') - V_{I_k}),\\
f''(x) &= (1-\mathrm{e}^{-r\Delta})u''(x) < 0,\\
g''(x) &= (1-\mathrm{e}^{-r\Delta})u''(x) - 2\mathrm{e}^{-r\Delta}\sigma\beta\Delta\mu I_u\,\mathrm{E}_z(\mathrm{e}^{-\nu\Delta}V_U(z') - V_{I_k}) < 0,
\end{aligned}$$

where the last two inequalities use $u'' < 0$ and $\mathrm{e}^{-\nu\Delta}V_U(z') \geq V_U(z') > V_{I_k}$ by Proposition 3.1. Since $A = [\underline{a}, 1]$ is nonempty, compact, convex, and f, g are continuous and strictly concave, they achieve unique maxima. By Definition 3.5, the maximum of f equals a_U^*. Since the last term of (3.14) does not depend on a_U^*, by Definition 3.7 the maximum of g equals $a_U^\dagger = a_U^\dagger(z)$.

Let us first show $a_U^* - a_U^\dagger \geq 0$. Note that at $x = a_U^*$, we have

$$\begin{aligned}
f'(a_U^*) &= (1-\mathrm{e}^{-r\Delta})u'(a_U^*) - \mathrm{e}^{-r\Delta}\sigma\beta\Delta\mu(a_I I_k + a_U^* I_u)\,\mathrm{E}_z(\mathrm{e}^{-\nu\Delta}V_U(z') - V_{I_k})\\
&\geq (1-\mathrm{e}^{-r\Delta})u'(a_U^*) - \mathrm{e}^{-r\Delta}\sigma\beta\Delta\mu(a_I I_k + 2a_U^* I_u)\,\mathrm{E}_z(\mathrm{e}^{-\nu\Delta}V_U(z') - V_{I_k})\\
&= g'(a_U^*).
\end{aligned}$$

If $f'(a_U^*) \leq 0$, then $g'(a_U^*) \leq f'(a_U^*) \leq 0$. Since g is strictly concave and $a_U^\dagger$ achieves its maximum, it must be $a_U^\dagger \leq a_U^*$. If $f'(a_U^*) > 0$, since f is strictly concave, it must be $a_U^* = 1$. Therefore $a_U^\dagger \le 1 = a_U^*$. In either case, we have $a_U^* - a_U^\dagger \ge 0$.

Next, let us show the upper bound in (3.17). If $f'(a_U^*) < 0$, since f is strictly concave, it must be $a_U^* = \underline{a}$. Then $\underline{a} \le a_U^\dagger \le a_U^* = \underline{a}$, so $a_U^* - a_U^\dagger = 0$ and the bound (3.17) is trivial. Similarly, if $g'(a_U^\dagger) > 0$, since g is strictly concave, it must be $a_U^\dagger = 1$. Then $1 = a_U^\dagger \le a_U^* \le 1$, so $a_U^*(z) - a_U^\dagger(z) = 0$ and the bound (3.17) is trivial. Therefore without loss of generality we may assume $f'(a_U^*) \ge 0 \ge g'(a_U^\dagger)$. If $f'(a_U^*) > 0 > g'(a_U^\dagger)$, then $1 = a_U^* \ge a_U^\dagger = \underline{a}$ and the bound (3.17) is trivial. Therefore we may assume either $f'(a_U^*) \ge 0 = g'(a_U^\dagger)$ or $f'(a_U^*) = 0 \ge g'(a_U^\dagger)$.

Case 1: $\boldsymbol{f'(a_U^*) \ge 0 = g'(a_U^\dagger)}$. Extrapolate f, g for $a > 1$ as

$$f_1(x) = \begin{cases} f(x), & (\underline{a} \le x \le 1) \\ f(1) + f'(1)(x-1) + \frac{1}{2} f''(1)(x-1)^2, & (x > 1) \end{cases}$$
$$g_1(x) = \begin{cases} g(x), & (\underline{a} \le x \le 1) \\ g(1) + g'(1)(x-1) + \frac{1}{2} f''(1)(x-1)^2. & (x > 1) \end{cases}$$

Clearly f_1, g_1 agree with f, g on $A = [\underline{a}, 1]$. Furthermore, f_1 is twice continuously differentiable and strictly concave on $[\underline{a}, \infty)$; g_1 is continuously differentiable and strictly concave. Since $f_1'(a_U^*) \ge 0 > -\infty = f_1'(\infty)$, f_1 achieves a unique maximum at some $a_1 \in [a_U^*, \infty)$ satisfying $f_1'(a_1) = 0$. Since $g'(a_U^\dagger) = 0$, g_1 achieves a unique maximum at $b_1 := a_U^\dagger$ and $g_1'(b_1) = 0$. Therefore

$$b_1 = a_U^\dagger \le a_U^* \le a_1. \tag{A.13}$$

Noting that $f_1''(x) = f''(1)$ for $x > 1$ and $m := \min_{a \in A} |u''(a)| > 0$, we obtain

$$\min_{x \ge \underline{a}} |f_1''(x)| = \min_{\underline{a} \le x \le 1} |f''(x)| = \min_{\underline{a} \le x \le 1} (1 - \mathrm{e}^{-r\Delta}) |u''(x)| = (1 - \mathrm{e}^{-r\Delta}) m. \tag{A.14}$$

By the definitions of f_1, g_1, we have

$$f_1'(x) - g_1'(x) = \begin{cases} \mathrm{e}^{-r\Delta} \sigma \beta \Delta \mu (2x - a_U^*) I_u \,\mathrm{E}_z(\mathrm{e}^{-\nu\Delta} V_U(z') - V_{I_k}), & (\underline{a} \le x \le 1) \\ f'(1) - g'(1). & (x > 1) \end{cases}$$

Taking the absolute value of both sides and setting $x = b_1 = a_U^\dagger \in [\underline{a}, 1]$, we obtain

$$\begin{aligned} |f_1'(b_1) - g_1'(b_1)| &= \mathrm{e}^{-r\Delta} \sigma \beta \Delta \mu \left| 2a_U^\dagger - a_U^* \right| I_u \,\mathrm{E}_z(\mathrm{e}^{-\nu\Delta} V_U(z') - V_{I_k}) \\ &\le \mathrm{e}^{-r\Delta} \sigma \beta \Delta \mu a_U^* I_u \,\mathrm{E}_z(\mathrm{e}^{-\nu\Delta} V_U(z') - V_{I_k}), \end{aligned} \tag{A.15}$$

where the last inequality uses $-a_U^* \le 2a_U^\dagger - a_U^* \le a_U^*$ because $0 \le a_U^\dagger \le a_U^*$. Noting that $f_1'(a_1) = g_1'(b_1) = 0$, combining (A.13), (A.14), (A.15), and applying Lemma A.5 to $(f, g) = (f_1', g_1')$ and $(a, b) = (a_1, b_1)$, we obtain

$$a_U^* - a_U^\dagger \le a_1 - b_1 \le \frac{|f_1'(b_1) - g_1'(b_1)|}{\min_{x \ge \underline{a}} |f_1''(x)|} \le \frac{\sigma\beta\Delta}{m(\mathrm{e}^{r\Delta} - 1)} \mu a_U^* I_u \, \mathrm{E}_z(\mathrm{e}^{-\nu\Delta} V_U(z') - V_{I_k}),$$

which implies the upper bound in (3.17).

Case 2: $f'(a_U^*) = 0 \ge g'(a_U^\dagger)$. Extrapolate f, g for $a < \underline{a}$ as

$$f_2(x) = \begin{cases} f(x), & (\underline{a} \le x \le 1) \\ f(\underline{a}) + f'(\underline{a})(x - \underline{a}) + \frac{1}{2} g''(\underline{a})(x - \underline{a})^2, & (x < \underline{a}) \end{cases}$$
$$g_2(x) = \begin{cases} g(x), & (\underline{a} \le x \le 1) \\ g(\underline{a}) + g'(\underline{a})(x - \underline{a}) + \frac{1}{2} g''(\underline{a})(x - \underline{a})^2. & (x < \underline{a}) \end{cases}$$

Clearly f_2, g_2 agree with f, g on $A = [\underline{a}, 1]$. Furthermore, f_2 is continuously differentiable and strictly concave on $(-\infty, 1]$; g_2 is twice continuously differentiable and strictly concave. Since $g_2'(-\infty) = \infty > 0 \ge g_2'(a_U^\dagger)$, g_2 achieves a unique maximum at some $a_2 \in (-\infty, a_U^\dagger]$ satisfying $g_2'(a_2) = 0$. Since $f'(a_U^*) = 0$, f_2 achieves a unique maximum at $b_2 := a_U^*$ and $f_2'(b_2) = 0$. Therefore

$$a_2 \le a_U^\dagger \le a_U^* = b_2. \tag{A.16}$$

Noting that $g_2''(x) = g''(\underline{a})$ for $x < \underline{a}$, $g''(x) \le f''(x) = (1 - \mathrm{e}^{-r\Delta}) u''(x) < 0$, and $m := \min_{a \in A} |u''(a)| > 0$, we obtain

$$\min_{x \le 1} |g_2''(x)| = \min_{\underline{a} \le x \le 1} |g''(x)| \ge \min_{\underline{a} \le x \le 1} (1 - \mathrm{e}^{-r\Delta}) \, |u''(x)| = (1 - \mathrm{e}^{-r\Delta}) m. \tag{A.17}$$

By the same argument as in the previous case, we can derive

$$|f_2'(b_2) - g_2'(b_2)| \le \mathrm{e}^{-r\Delta} \sigma\beta\Delta\mu a_U^* I_u \, \mathrm{E}_z(\mathrm{e}^{-\nu\Delta} V_U(z') - V_{I_k}). \tag{A.18}$$

Noting that $g_2'(a_2) = f_2'(b_2) = 0$, combining (A.16), (A.17), (A.18), and applying Lemma A.5 to $(f, g) = (g_2', f_2')$ and $(a, b) = (a_2, b_2)$, we obtain

$$a_U^* - a_U^\dagger \le b_2 - a_2 \le \frac{|g_2'(b_2) - f_2'(b_2)|}{\min_{x \le 1} |g_2''(x)|} \le \frac{\sigma\beta\Delta}{m(\mathrm{e}^{r\Delta} - 1)} \mu a_U^* I_u \, \mathrm{E}_z(V_U(z') - V_{I_k}),$$

which implies the upper bound in (3.17).

If $\sigma = 1$, since $I_u = 0$, we have $a_U^\dagger(z) = a_U^*(z)$ by (3.17). $\square$

B Observational equivalence between single and multiple signals

In Section 2 we supposed that an infected agent is informed of the infection with probability σ and that a known (unknown) infected agent dies with probability δ (0). In this Appendix we elaborate upon this assumption and show that it can model the case in which there are multiple signals.

Suppose now that there are multiple signals that may be received upon infection, indexed $j = 1, \dots, J$. Let $\sigma_j > 0$ be the probability of receiving signal j upon infection, with $\sigma := \sum_{j=1}^J \sigma_j \in (0, 1]$. Let $\delta_j \in [0, 1]$ be the fatality rate conditional on receiving signal j and $\delta := (\sum_{j=1}^J \sigma_j \delta_j)/\sigma$ be the expected fatality conditional on receiving any signal. For instance, the signal could encode the result of a laboratory test as well as the type and severity of symptoms. Suppose a type j known infected agent takes action a_j with associated flow utility u_j. Then by (3.2), the value function of a type j agent is

$$V_j := \frac{(1 - \mathrm{e}^{-r\Delta})u_j + \mathrm{e}^{-r\Delta}\gamma\Delta\delta_j u_D}{1 - \mathrm{e}^{-r\Delta}(1 - \gamma\Delta)}.$$

Therefore the expected continuation value of being known infected is

$$V_{I_k} := \frac{1}{\sigma}\sum_{j=1}^J \sigma_j V_j = \frac{(1 - \mathrm{e}^{-r\Delta})u_I + \mathrm{e}^{-r\Delta}\gamma\Delta\delta u_D}{1 - \mathrm{e}^{-r\Delta}(1 - \gamma\Delta)}, \tag{B.1}$$

where $u_I := (\sum_{j=1}^J \sigma_j u_j)/\sigma$ is the expected flow utility of being known infected. Because (B.1) is identical to (3.2) and the continuation values $\{V_j\}$ affect the behavior of unknown agents only through V_{I_k} due to expected utility maximization, a model with multiple signals and heterogeneous fatality is observationally equivalent to a model with a single signal and uniform fatality. The flow utility u_I and action a_I in Section 2 can be interpreted as the average across those of known infected agents.

C Perfect recall Markov equilibrium

We now consider an extension in which we allow the beliefs of the agents to enter as a separate state variable, the evolution of which depends directly upon their actions. The evolution of the aggregate state variables and the value functions of the known infected, recovered, and dead agents are unchanged relative to the perfect Bayesian equilibrium, and we must therefore only alter the problem of the unknown agents. Now (2.6) will remain their belief *in equilibrium*, but not if they deviated from the equilibrium action in the past.

If the average action of unknown agents is $\tilde{a}$, an agent with belief μ who chooses activity a will believe that if he is susceptible, he becomes infected at rate $\beta a((1-\sigma)\tilde{a}+\sigma a_I)I$. His belief that he is susceptible is then the probability that he was susceptible in the previous period multiplied by the probability that he was not infected without diagnosis during the last period, or

$$\mu_{+\Delta} = \mu(1-\Delta(1-\sigma)\beta a((1-\sigma)\tilde{a}+\sigma a_I)I). \tag{C.1}$$

We then define a *perfect recall Markov equilibrium* as an allocation in which the unknown agents optimize given the law of motion of the population shares *and* their beliefs, and these beliefs are consistent with equilibrium behavior and the Bayes rule. Formally, we adopt the continuous-time formulation used in the numerical analysis and proceed as follows. Rearranging (C.1) and sending $\Delta \to 0$ gives the evolution of beliefs

$$\dot{\mu} = -\mu(1-\sigma)\beta a\big((1-\sigma)\tilde{a}(S,I)+\sigma a^*_{I_k}\big)I. \tag{C.2}$$

The Hamilton-Jacobi-Bellman for unknown agents is

$$\begin{aligned}(r+\nu)V_U(S,I,\mu) = \max_{a_U\in[\underline{a},1]} \; & ru(a_U) \\ & + \sigma\mu\beta a_U\big((1-\sigma)\tilde{a}(S,I)+\sigma a^*_{I_k}\big)I[V_{I_k}(S,I)-V_U(S,I,\mu)] \\ & - \beta S\tilde{a}(S,I)\big((1-\sigma)\tilde{a}(S,I)+\sigma a^*_{I_k}\big)I\partial_S V_U(S,I,\mu) \\ & + \big(\beta S\tilde{a}(S,I)\big((1-\sigma)\tilde{a}(S,I)+\sigma a^*_{I_k}\big)-\gamma\big)I\partial_I V_U(S,I,\mu) \\ & - \mu(1-\sigma)\beta a_U\big((1-\sigma)\tilde{a}(S,I)+\sigma a^*_{I_k}\big)I\partial_\mu V_U(S,I,\mu).\end{aligned}$$

Given an average action $\tilde{a}$ of unknown agents, there is an associated policy function $a(S,I,\mu;\tilde{a})$ solving the problem of a given unknown agent. Now define an operator $J(\tilde{a})(S,I) = a(S,I,S/[1-\sigma+\sigma S];\tilde{a})-\tilde{a}$. The equilibrium notion we adopt in this section is then that of a Markov perfect equilibrium, in which all agents solve their individual problems taking the aggregate law of motion as given, and the associated law of motion is consistent with individual behavior.

Definition C.1. A *perfect recall Markov equilibrium* consists of value functions $V_U(S, I, \mu)$ and $V_{I_k}(S, I)$ for unknown agents and known infected agents together with a policy function $a(S, I, \mu)$ for unknown agents such that:

- The functions $V_U(S, I, \mu)$, $a(S, I, \mu)$ and $V_{I_k}(S, I)$ solve the problems of the unknown agents and known infected agents, respectively.

- The law of motion of the aggregate state is consistent with the policy function of the unknown agents, or $J(\tilde{a}) = 0$.

As Figure 8 shows, in the perfect recall Markov equilibrium agents take higher actions everywhere in the state space and the length of the pandemic is reduced.

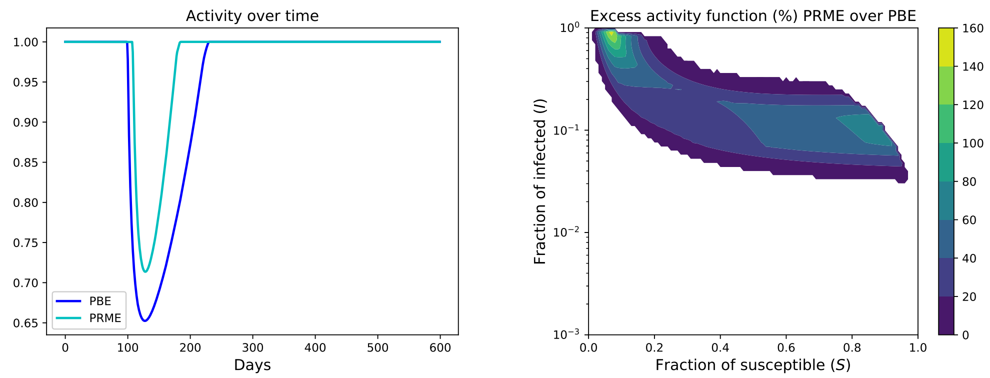

Figure 8: Comparison of perfect Bayesian and perfect recall Markov equilibria.

D Solution algorithm

Our model with (large but) finitely many competitive agents is not computationally tractable because the state space Z in (2.1) is very large.[13] To make the model computationally tractable, we make two approximations. First, we suppose that the number of agents N is large and apply the strong law of large numbers, which makes the transitions of the aggregate state deterministic as in

[13] To be specific, since there are N agents and 6 agent types, there are $\binom{N+5}{5}$ combinations of aggregate states. Even if we note that the distinction between R_k and D agents is payoff irrelevant and combine them into one type, and N is small such as $N = 100$, the number of possible states $\binom{104}{4} = 4{,}598{,}126$ is very large. With a more realistic N, which is of an order over a million, the number of possible combinations is astronomical.

(3.19). Under this approximation, we have $I_k = \sigma I$ and $I_u = (1-\sigma)I$. Since $R_k + R_u + D = 1 - S - I$ and an infected agent becomes symptomatic with probability σ, we have $R_k + D = \sigma(1-S-I)$ and $R_u = (1-\sigma)(1-S-I)$. Therefore the belief $\mu(z)$ in (2.6) becomes

$$\mu(z) = \frac{S}{S + I_u + R_u} = \frac{S}{S + (1-\sigma)I + (1-\sigma)(1-S-I)} = \frac{S}{1-\sigma+\sigma S}.$$

Because the behavior of R_k, R_u, D agents does not affect state transitions, with a slight abuse of notation we can define the minimal state space by

$$Z = \{(S, I) \mid S \geq 0, I \geq 0, S + I \leq 1\}. \tag{D.1}$$

Second, to aid in numerical accuracy, we produce the figures in the main text by considering the continuous-time limit of our model and applying the finite-state Markov chain method of Kushner and Dupuis (1992). The existence of an approximate equilibrium can be established by arguments that are essentially identical to those in Theorem 3.6.

D.1 Exponential grid

When solving for an SIR model numerically, because the fraction of infected agents I varies by many orders of magnitude (say between 10^{-6} and 10^{-1}), it is important to use a grid that properly covers the relevant range of the state space, such as an exponential grid. In general, suppose we would like to construct an N-point exponential grid on a given interval (a, b), where the lower bound a is not necessarily positive. A natural idea to deal with such a case is as follows.

Constructing exponential grid.

(i) Choose a shift parameter $s > -a$.

(ii) Construct an N-point evenly-spaced grid on $(\log(a+s), \log(b+s))$.

(iii) Take the exponential and subtract s.

The remaining question is how to choose the shift parameter s. Suppose we would like to specify the median grid point as $c \in (a, b)$. Since the median of the evenly-spaced grid on $(\log(a+s), \log(b+s))$ is $\frac{1}{2}(\log(a+s) + \log(b+s))$, we need to take $s > -a$ such that

$$c = \exp\left(\frac{1}{2}(\log(a+s) + \log(b+s))\right) - s \iff s = \frac{c^2 - ab}{a+b-2c}.$$

Note that in this case $s + a = \frac{c^2-ab}{a+b-2c} + a = \frac{(c-a)^2}{a+b-2c}$, so $s + a$ is positive if and only if $c < \frac{a+b}{2}$. Therefore, for any $c \in \left(a, \frac{a+b}{2}\right)$, it is possible to construct an exponentially-spaced grid with end points (a, b) and median point c.

To solve the model in Section 4 numerically, we construct finite grids for S, I and define the minimal state space Z in (D.1) by the Cartesian product of the S, I grids. For the S-space (fraction of susceptible agents), we use a 100-point uniform grid on $[10^{-8}, 1]$. For the I-space we use a 400-point exponential grid on $[10^{-8}, 1]$ with a median point of 10^{-4}. We now write these grids as $\Sigma_S = \{0, 1/N_S, \ldots, 1 - 1/N_S, 1\}$ and $\Sigma_I = \{I_0, I_1, \ldots, I_{N_I}\}$, and define

$$\begin{aligned}
\Delta_{Ii}^- &= I_i - I_{i-1} & i &= 1, \ldots, N_I \\
\Delta_{Ii}^+ &= I_{i+1} - I_i & i &= 0, \ldots, N_I - 1
\end{aligned}$$

and declare $\Delta_{I0}^- = \Delta_{IN_I}^+ = 0$ and $\Delta_S = 1/N_S$. We then write $\Sigma \coloneqq \Sigma_S \times \Sigma_I$.

D.2 Perfect Bayesian Markov competitive equilibrium

We consider the numerical method for computing the perfect Bayesian equilibrium. Let $V_h(S, I)$ be the value function of type $h = U, I_k$ agents, $\tilde{a}(S, I)$ be the policy function of unknown agents, and let partial derivatives be denoted by ∂_S etc. In this case the Hamilton-Jacobi-Bellman equation for the unknown agent becomes

$$\begin{aligned}
(r+\nu)V_U(S,I) = \max_{a_U \in [\underline{a}, 1]} & \, ru(a_U) \\
& + \sigma\mu\beta a_U\big((1-\sigma)\tilde{a}(S,I) + \sigma a_{I_k}^*\big) I [V_{I_k}(S,I) - V_U(S,I)] \\
& - \beta S\tilde{a}(S,I)\big((1-\sigma)\tilde{a}(S,I) + \sigma a_{I_k}^*\big) I \partial_S V_U(S,I) \\
& + \big(\beta S\tilde{a}(S,I)\big((1-\sigma)\tilde{a}(S,I) + \sigma a_{I_k}^*\big) - \gamma\big) I \partial_I V_U(S,I). \quad \text{(D.2)}
\end{aligned}$$

We suppose that at an arbitrary $(S, I) \in \Sigma$ there are three possible transitions, to $(S - \Delta_S, I)$, $(S, I - \Delta_I^-)$ and $(S, I + \Delta_I^+)$, with associated probabilities p^{-S}, p^{-I}

and p^{+I} given by

$$p^{-S} = \frac{\Delta_t}{\Delta_S}\beta S\tilde{a}(S,I)\big(\big((1-\sigma)\tilde{a}(S,I)+\sigma a^*_{I_k}\big)I, \tag{D.3a}$$

$$p^{\pm I} = \frac{\Delta_t}{\Delta_I^\pm}\max\left\{\pm\big[\beta S\tilde{a}(S,I)\big((1-\sigma)\tilde{a}(S,I)+\sigma a^*_{I_k}\big)-\gamma\big]I,0\right\}. \tag{D.3b}$$

The probability with which an unknown agent becomes known infected is then

$$p^{uk} = \Delta_t\sigma\mu\beta a_U\big((1-\sigma)\tilde{a}(S,I)+\sigma a^*_{I_k}\big)I, \tag{D.4}$$

where the superscript indicates this is the probability of transitioning from the unknown u state to the known infected state k. The Bellman equation for unknown agents is then

$$\begin{aligned}
V_U(S,I) &= \max_{a_U\in[\underline{a},1]} \Delta_t r u(a_U) + \mathrm{e}^{-(r+\nu)\Delta_t} p^{uk}[V_{I_k}(S,I)-V_U(S,I)] \\
&+ \mathrm{e}^{-(r+\nu)\Delta_t}\big(p^{-S}V_U(S-\Delta_S, I+\Delta_S) + p^{-I}V_U(S, I-\Delta_I^-) + p^{+I}V_U(S, I+\Delta_I^+)\big) \\
&+ \mathrm{e}^{-(r+\nu)\Delta_t}\big(1-p^{-S}-p^{-I}-p^{+I}-p^{uk}\big)V_U(S,I).
\end{aligned}$$

Omitting terms independent of a_U, dividing by $r\Delta_t$ and sending $\Delta_t \to 0$, the maximization problem becomes

$$\max_{a_U\in[\underline{a},1]} u(a_U) + \bar{p}^{uk} r^{-1}[V_{I_k}(S,I)-V_U(S,I)]$$

where $\bar{p}^{uk} = p^{uk}/\Delta_t$. This problem is equivalent to

$$\max_{a_U\in[\underline{a},1]} u(a_U) + c a_U\big((1-\sigma)\tilde{a}(S,I)+\sigma a^*_{I_k}\big),$$

where

$$c = \sigma\mu\beta r^{-1}[V_{I_k}(S,I)-V_U(S,I)]I. \tag{D.5}$$

When the utility function takes the form (4.1), the first-order condition rearranges to $a_U = a_U^{\text{FOC}} := [-c]^{-1/\alpha}$ and the optimal choice of a_U is therefore

$$a_U = 1_{c\geq 0} + (1-1_{c\geq 0})\max\left\{\underline{a}, \min\left\{1, a_U^{\text{FOC}}\right\}\right\}. \tag{D.6}$$

Note that the linear system characterizing the value function of the unknown agents is

$$
\begin{aligned}
0 =& ru(a_U) + \bar{p}^{uk} V_{I_k}(S, I) \\
&+ \bar{p}^{-I} V_U(S, I - \Delta_I^-) + \bar{p}^{+I} V_U(S, I + \Delta_I^+) + \bar{p}^{-S} V_U(S - \Delta_S, I + \Delta_S) \\
&+ \big(r + \nu + \bar{p}^{-S} + \bar{p}^{-I} + \bar{p}^{+I} + \bar{p}^{uk}\big) V_U(S, I). \qquad \text{(D.7)}
\end{aligned}
$$

We then iterate upon the policy function a_U. Beginning with an arbitrary guess a_U, we compute the value function of the unknown agent by solving the linear system (D.7), replace a_U with the implied policy function a_U in (D.6), and repeat until convergence.

D.3 Efficient action

We consider the numerical method for computing the efficient action. Taking the limit as $\Delta \to 0$ in the main text, one can show that the value function of the planner is of the form $W(S, I, D) = Du_D - C(S, I)$, where C solves the Hamilton-Jacobi-Bellman equation

$$
\begin{aligned}
(r + \nu) C(S, I) = \min_{a_U \in [\underline{a}, 1]} & r\big((1 - \sigma + \sigma S)[-u(a_U)] + \sigma I[-u(a^*_{I_k})]\big) + \nu \bar{C}_{\text{vac}} I \\
&+ \gamma \delta \sigma I[-u_D] - \beta S I a_U \big((1 - \sigma) a_U + \sigma a^*_{I_k}\big) \partial_S C(S, I) \\
&+ [\beta S a_U \big((1 - \sigma) a_U + \sigma a^*_{I_k}\big) - \gamma] I \partial_I C(S, I),
\end{aligned}
$$

where

$$
\bar{C}_{\text{vac}} := \frac{\sigma}{r + \gamma}(-\gamma \delta u_D - r u(a^*_{I_k})) \qquad \text{(D.8)}
$$

is interpreted as the cost of the pandemic (in terms of utility) when the vaccine has arrived. We first construct a locally consistent Markov chain for the law of motion of the state variables. The local consistency requirements are given by

$$
\mathrm{E}[\Delta X_S] = -\Delta_t \beta S I a_U \big((1 - \sigma) a_U + \sigma a^*_{I_k}\big) + o(\Delta_t) \qquad \text{(D.9a)}
$$

$$
\mathrm{E}[\Delta X_I] = \Delta_t \big(\beta S I a_U \big((1 - \sigma) a_U + \sigma a^*_{I_k}\big) - \gamma I\big) + o(\Delta_t). \qquad \text{(D.9b)}
$$

For an arbitrary $(S, I) \in \Sigma$ there are three possible transitions, to $(S - \Delta_S, I)$, $(S, I - \Delta_I^-)$ and $(S, I + \Delta_I^+)$, with associated probabilities p^{-S}, p^{-I} and p^{+I}. The local consistency requirements (D.9) become

$$
\begin{aligned}
-\Delta_S p^{-S} &= -\Delta_t \beta S I a_U\big((1-\sigma)a_U + \sigma a^*_{I_k}\big) + o(\Delta_t)\\
-\Delta_I^- p^{-I} + \Delta_I^+ p^{+I} &= \Delta_t\big(\beta S a_U\big((1-\sigma)a_U + \sigma a^*_{I_k}\big) - \gamma\big) I + o(\Delta_t).
\end{aligned}
$$

Inspection reveals it will suffice to set

$$
\begin{aligned}
p^{-S} &= \frac{\Delta_t}{\Delta_S}\beta S I a_U\big((1-\sigma)a_U + \sigma a^*_{I_k}\big) && \text{(D.10a)}\\
p^{+I} &= \frac{\Delta_t}{\Delta_I^+}\beta S I a_U\big((1-\sigma)a_U + \sigma a^*_{I_k}\big) && \text{(D.10b)}\\
p^{-I} &= \frac{\Delta_t}{\Delta_I^-}\gamma I && \text{(D.10c)}
\end{aligned}
$$

To ensure the chain remains on the grid we declare that at $I = 1$ we have $a_U \leq \hat{a}_1(S)$, the point at which $0 = \beta S a_U((1-\sigma)a_U + \sigma a^*_{I_k}) - \gamma$. For $S > 0$ we have

$$
\hat{a}_1(S) = \frac{1}{2\beta S(1-\sigma)}\Big(-\beta S\sigma a^*_{I_k} + \sqrt{[\beta S\sigma a^*_{I_k}]^2 + 4\gamma\beta S(1-\sigma)}\Big).
$$

We then have the Bellman equation

$$
\begin{aligned}
C(S,I) =& \nu\Delta_t \bar{C}_{\text{vac}} I + \min_{a_U\in[\underline{a},1]} -\Delta_t r\big((1-\sigma+\sigma S)u(a_U) + \sigma I u(a^*_{I_k})\big) - \Delta_t\gamma\delta\sigma I u_D\\
&+ \mathrm{e}^{-(r+\nu)\Delta_t}\big(p^{-S}C(S-\Delta_S, I) + p^{+I}C(S, I+\Delta_I^+) + p^{-I}C(S, I-\Delta_I^-)\big)\\
&+ \mathrm{e}^{-(r+\nu)\Delta_t}(1 - p^{-S} - p^{+I} - p^{-I})C(S,I).
\end{aligned}
$$

Omitting terms independent of the controls, using (D.10) and sending $\Delta_t \to 0$ gives

$$
\begin{aligned}
\min_{a_U\in[\underline{a},1]} &- r(1-\sigma+\sigma S)u(a_U) + \beta S I a_U\big((1-\sigma)a_U + \sigma a^*_{I_k}\big)[-C^{BS}]\\
&+ \beta S I a_U\big((1-\sigma)a_U + \sigma a^*_{I_k}\big)C^{FI}.
\end{aligned}
$$

Dividing by $-r(1-\sigma+\sigma S)$ gives

$$
\max_{a_U\in[\underline{a},1]} u(a_U) + \beta S a_U\big((1-\sigma)a_U + \sigma a^*_{I_k}\big)\frac{I(C^{BS} - C^{FI})}{r(1-\sigma+\sigma S)}. \tag{D.11}
$$

Note that (D.11) is of the form

$$
G(a,b,c) \coloneqq \max_{a_U\in[a,b]} u(a_U) + c a_U\big((1-\sigma)a_U + \sigma a^*_{I_k}\big), \tag{D.12}
$$

where $c := \beta SI\big(C^{BS} - C^{FI}\big)/[r(1-\sigma+\sigma S)]$. If $c \geq 0$ then $a_U = 1$. Otherwise the objective function is concave, and we only need to evaluate the first-order conditions. The first-order condition condition for (D.12) when u is CRRA is

$$0 = a_U^{-\alpha} + c(\sigma a_{I_k}^* + 2(1-\sigma)a_U). \tag{D.13}$$

Lemma D.1. *For any $0 < a < b$, the solution to (D.12) is*

$$a_U(c) = b1_{c\geq 0} + (1 - 1_{c\geq 0}) \max\left\{a, \min\left\{a_U^{\text{FOC}}(c), b\right\}\right\},$$

where $a_U^{\text{FOC}}(c)$ solves (D.13).

To avoid overflow in the code we divide all quantities by Δ_t and consider the limit of the above as $\Delta_t \to 0$. The linear system we solve at each stage is

$$\begin{aligned}
0 = & - r\big((1-\sigma+\sigma S)u(a_U) + \sigma I u(a_{I_k}^*)\big) - \gamma\delta\sigma I u_D + \nu \bar{C}_{\text{vac}} I \\
& + \bar{p}^{-S} C(S - \Delta_S, I) + \bar{p}^{+I} C(S, I + \Delta_I^+) + \bar{p}^{-I} C(S, I - \Delta_I^-) \\
& - (r + \nu + \bar{p}^{-S} + \bar{p}^{+I} + \bar{p}^{-I})C,
\end{aligned} \tag{D.14}$$

where for each direction we have $\bar{p} = p/\Delta_t$. We then iterate upon the policy function a_U. Beginning with an arbitrary guess a_U, we compute the planner's objective function by solving the linear system (D.14), replace a_U with the implied policy function a_U in Lemma D.1, and repeat until convergence.

D.4 Perfect recall Markov equilibrium

We consider the numerical method for computing the perfect recall Markov equilibrium. We define a grid for beliefs $\Sigma_\mu = \{0, 1/N_\mu, \ldots, 1 - 1/N_\mu, 1\}$ for some integer $N_\mu \geq 1$. Let $\Sigma := \Sigma_S \times \Sigma_I \times \Sigma_\mu$.

For unknown agents, we must specify the transition probabilities for their beliefs and the probability with which they become known infected. The latter quantity is simply (D.4). We suppose that at an arbitrary $(S, I, \mu) \in \Sigma$, there are four possible transitions, to $(S - \Delta_S, I, \mu)$, $(S, I - \Delta_I^-, \mu)$, $(S, I + \Delta_I^+, \mu)$ and $(S, I, \mu - \Delta_\mu)$, with associated probabilities p^{-S}, p^{-I}, p^{+I} and $p^{-\mu}$. The local consistency requirements for the aggregate state are again satisfied if we choose p^{-S}, p^{-I} and p^{+I} according to (D.3). Using (C.2), the local consistency requirement for the belief variable is

$$-\Delta_\mu p^{-\mu} = -\Delta_t(1-\sigma)\mu\beta a_U\big((1-\sigma)\tilde{a}(S, I) + \sigma a_{I_k}^*\big)I + o(\Delta_t),$$

which will be satisfied if we choose

$$p^{-\mu} = \frac{\Delta_t}{\Delta_\mu}\mu(1-\sigma)\beta a_U\big((1-\sigma)\tilde{a}(S,I) + \sigma a^*_{I_k}\big)I = \frac{1}{\Delta_\mu}(1/\sigma - 1)p^{uk}, \tag{D.15}$$

where p^{uk} is given by (D.4). The Bellman equation for unknown agents is then

$$\begin{aligned} V_U(S,I,\mu) = \max_{a_U\in[\underline{a},1]} &\Delta_t r u(a_U) + \mathrm{e}^{-(r+\nu)\Delta_t}\frac{\Delta_\mu \sigma}{1-\sigma}p^{-\mu}[V_{I_k}(S,I) - V_U(S,I,\mu)] \\ &+ \mathrm{e}^{-(r+\nu)\Delta_t}\big(p^{-S}V_U(S-\Delta_S, I+\Delta_S,\mu) + p^{-I}U(S, I-\Delta_I^-,\mu;u)\big) \\ &+ \mathrm{e}^{-(r+\nu)\Delta_t}\big(p^{+I}V_U(S, I+\Delta_I^+,\mu) + p^{-\mu}V_U(S,I,\mu-\Delta_\mu)\big) \\ &+ \mathrm{e}^{-(r+\nu)\Delta_t}\left(1 - p^{-S} - p^{-I} - p^{+I} - p^{-\mu} - \frac{\Delta_\mu\sigma}{1-\sigma}p^{-\mu}\right)V_U(S,I,\mu). \end{aligned}$$

Omitting terms independent of a_U, dividing by $r\Delta_t$ and sending $\Delta_t \to 0$, the maximization problem becomes

$$\begin{aligned} &\max_{a_U\in[\underline{a},1]} u(a_U) \\ &\quad + \bar{p}^{-\mu} r^{-1}\left(\frac{\Delta_\mu\sigma}{1-\sigma}[V_{I_k}(S,I) - V_U(S,I,\mu)] + V_U(S,I,\mu-\Delta_\mu) - V_U(S,I,\mu)\right), \end{aligned}$$

where $\bar{p}^{-\mu} = p^{-\mu}/\Delta_t$. Using (D.15), this problem is equivalent to

$$\max_{a_U\in[\underline{a},1]} u(a_U) + c a_U\big((1-\sigma)\tilde{a}(S,I) + \sigma a^*_{I_k}\big),$$

where

$$c = \frac{\mu\beta}{r\Delta_\mu}(\Delta_\mu\sigma[V_{I_k}(S,I) - V_U(S,I,\mu)] + (1-\sigma)[V_U(S,I,\mu-\Delta_\mu) - V_U(S,I,\mu)])I. \tag{D.16}$$

When the utility function takes the form (4.1), the first-order condition rearranges to $a_U = a_U^{\mathrm{FOC}} := [-c]^{-1/\alpha}$ and the optimal choice of a_U is therefore

$$a_U = 1_{c\geq 0} + (1 - 1_{c\geq 0})\max\left\{\underline{a}, \min\left\{1, a_U^{\mathrm{FOC}}\right\}\right\}. \tag{D.17}$$

Note that the linear system characterizing the value function of the unknown agents is

$$
\begin{aligned}
0 =& ru(a_U) + \frac{\Delta_\mu \sigma}{1-\sigma} \bar{p}^{-\mu} V_{I_k}(S,I) + \bar{p}^{-\mu} V_U(S,I,\mu-\Delta_\mu) \\
&+ \bar{p}^{-I} V_U(S, I-\Delta_I^-, \mu) + \bar{p}^{+I} V_U(S, I+\Delta_I^+, \mu) + \bar{p}^{-S} V_U(S-\Delta_S, I+\Delta_S, \mu) \\
&+ \left(r + \nu + \bar{p}^{-S} + \bar{p}^{-I} + \bar{p}^{+I} + \bar{p}^{-\mu} + \frac{\Delta_\mu \sigma}{1-\sigma} \bar{p}^{-\mu} \right) V_U(S,I,\mu). \qquad \text{(D.18)}
\end{aligned}
$$

We then iterate upon the policy function a_U. Beginning with an arbitrary guess a_U, we compute the value function of the unknown agent by solving the linear system (D.18), replace a_U with the implied policy function a_U in (D.17), and repeat until convergence.

E Calibrating u_D

We calibrate the flow utility from death u_D based on the case study from Sweden, which did not introduce mandatory lockdowns. For this purpose, we first obtain the daily cumulative number of reported cases and deaths for Sweden from Johns Hopkins University CSSE (Footnote 12). Let $N_{C,t}, N_{D,t}$ be the cases and deaths up to date t. We divide these numbers by the population $N = 10.38 \times 10^6$ to obtain the population shares of reported cases $C_t := N_{C,t}/N$ and deaths $D_t := N_{D,t}/N$. We compute the case fatality rate as $\delta_{\mathrm{CFR},t} := D_t/C_t$. Figure 9a shows the evolution of the case fatality rate. The fact that the CFR peaked at around 12 percent in May 2020 and has settled down below 2 percent at the time of writing suggests that the reporting rate has significantly changed over time and that the reported cases are unreliable.

We thus estimate the fraction of the infected population I_t using the accounting equation from the SIR model

$$
D_{t+1} - D_t = \gamma \delta_{\mathrm{IFR}} I_t \iff I_t = \frac{D_{t+1} - D_t}{\gamma \delta_{\mathrm{IFR}}}. \qquad \text{(E.1)}
$$

To control for the day-of-week effect (deaths seem to be unreported over the weekend), we take the 7-day moving average of (E.1), which is plotted in Figure 9b. The estimated peak prevalence is thus $\max_t I_t = 0.0663$. Finally, we choose the value for u_D to match the peak prevalence in the model and obtain $u_D = -12.22$.

As a robustness check, we also compute u_D from the value of statistical life. Consider an individual consuming a constant flow normalized to 1 and facing a small probability d of death for one period. Letting $\beta \in (0,1)$ be the agent's

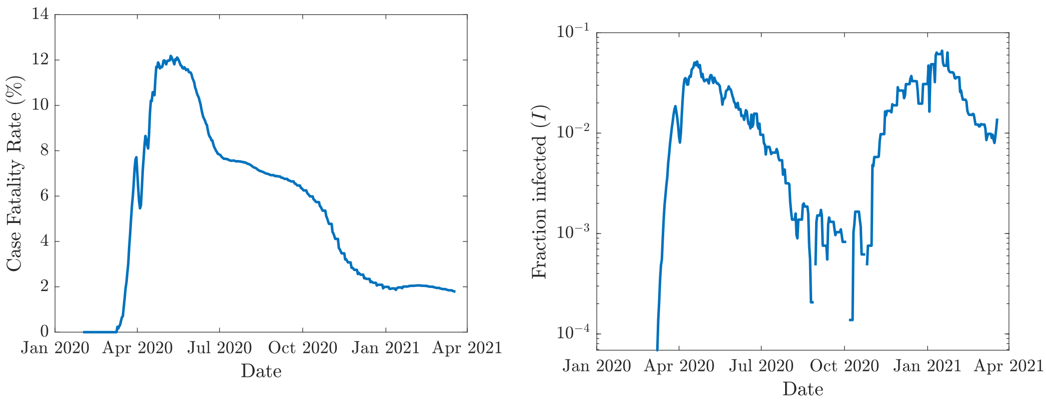

(a) Case Fatality Rate (percent) over time.

(b) Estimated prevalence over time.

Figure 9: Case Fatality Rate and estimated prevalence in Sweden.

discount factor, p be the willingness to pay to avoid the possibility of death, V be the continuation value of being alive, and V_D be the continuation value of being dead, by definition we have

$$(1-\beta)u(1-p)+\beta V=(1-\beta)u(1)+\beta((1-d)V+dV_D). \tag{E.2}$$

Using $V=u(1)$ and $V_D=u_D$, (E.2) simplifies to

$$(1-\beta)u(1-p)=(1-\beta-\beta d)u(1)+\beta d u_D. \tag{E.3}$$

The *value of statistical life* v is defined by the willingness to pay p scaled such that one life is saved on average, so $v=p/d$. Since

$$u(1-p)\approx u(1)-u'(1)p=u(1)-u'(1)vd$$

by the Taylor approximation, solving (E.3) for u_D, we obtain

$$u_D\approx u(1)-\frac{1-\beta}{\beta}u'(1)v. \tag{E.4}$$

Since $u(1)=0$ and $u'(1)=1$ for the CRRA utility (4.1), (E.4) further simplifies to

$$u_D\approx -\frac{1-\beta}{\beta}v. \tag{E.5}$$

Hall et al. (2020) note that the Environmental Protection Agency (EPA) recommends using $v =$ \$7.4 million in 2006,[14] when per capita consumption was about \$31,000. Thus the value of life in units of one period consumption is

$$v = \frac{7{,}400{,}000}{31{,}000} = 238.7,$$

and with 5 percent annual discounting, the flow utility from death in (E.5) is

$$u_D = -0.05 \times 238.7 = -11.94,$$

which is very similar to the value we obtained from the case study from Sweden.

[14] https://www.epa.gov/environmental-economics/mortality-risk-valuation