Social engagement and the spread of infectious diseases

Lucie Lebeau*
University of California, Irvine

This version: December 6, 2020

Abstract

This paper endogenizes the spread of an infectious disease in a random matching model with pairwise meetings, where economic and social gains arise explicitly from person-to-person contacts. When agents can decide whether to engage in interactions, complementarities in the participation decisions of individuals susceptible to contracting the disease generate a large multiplicity of equilibria through adverse selection. The lower the participation of susceptible agents, the higher the prevalence of infection in the pool of participants, further discouraging the participation of susceptible agents. I document a variety of infection dynamics, including plateaus and multiple waves. Adverse selection leads to too much isolation from susceptible agents, and in the calibrated version of the model, the cost of forgone interactions offsets the welfare gains of flattening the curve and mitigating the human toll. When agents cannot opt out of the market but can instead choose whether to wear a mask, the equilibrium is unique. In the calibrated model the human toll is lower than when considering the participation margin, yet at a significantly smaller cost.

Keywords: disease transmission, social distancing, masks, COVID-19

*I am very grateful to Guillaume Rocheteau for his extensive feedback and support in writing this paper. I also thank Pierre-Olivier Weill for his insightful comments and suggestions. In addition, I thank seminar participants at Virginia Commonwealth University, University of Basel, and the UC Irvine Macro PhD workshop. I would also like to thank the UCI Department of Economics for financial research support. All errors are my own. Email address: llebeau@uci.edu.

1 Introduction

This paper investigates the transmission of an infectious disease in a random matching model where economic and social gains from trade directly stem from person-to-person contacts. In typical epidemiological models of disease transmission, infection dynamics are mechanically driven by exogenous, reduced-form parameters meant to encompass both the fundamentals of the disease and individuals' behaviors. Modeling individuals as rational and forward-looking agents, who face a trade-off between the benefits from engaging in social interactions and the infection risk it carries, allows us to endogenize the infection rate by letting agents' behavior react over time as the epidemic develops.

I consider two response margins, both of which have been at the forefront of public health recommendations since the onset of the COVID-19 pandemic: social engagement and mask-wearing.[1] Figure 1 shows that since the beginning of 2020, and throughout the world, a significant share of individuals have taken action along both of those margins. While policy mandates have certainly largely contributed to those reactions, they do not account for all of it. For example, although Finland never required mask-wearing in public, the right panel shows that by the end of September 2020, more than 20% of the Finns surveyed reported taking that precaution.[2]

The paper aims to answer the following questions. First, when and to which extent do individuals modify their behaviors in response to the outbreak of a contagious disease? Second, how are epidemic dynamics and economic outcomes impacted by these rational changes in behavior? In particular, can the equilibrium path still be precisely predicted, as is the case in epidemiological models that do not account for endogenous changes in behaviors? Third, how do these dynamics and outcomes differ when agents modulate the frequency of their social engagement relative to when they modulate the precautions they take during social interactions?

To answer these questions, I build a model that retains a structure similar to that of compartmental models of disease transmission, developed in the wake of Kermack and McKendrick (1927). The population is divided into three groups respectively labeled S, I, and R: susceptible individuals, who can contract the disease; infectious

[1] For example, in an interview with ABC News, Dr. Fauci, director of the National Institute of Allergy and Infectious Diseases in the United States, stated that "[the] best way that you can avoid — either acquiring or transmitting infection — is to avoid crowded places, to wear a mask whenever you're outside." The Centers for Disease Control and Prevention (CDC) makes similar recommendations.

[2] Additional evidence can be found in Farboodi et al. (2020), who use data based on cellphone tracking to show that everywhere across the United States individuals started to reduce their social activity before any policy measures were enacted.

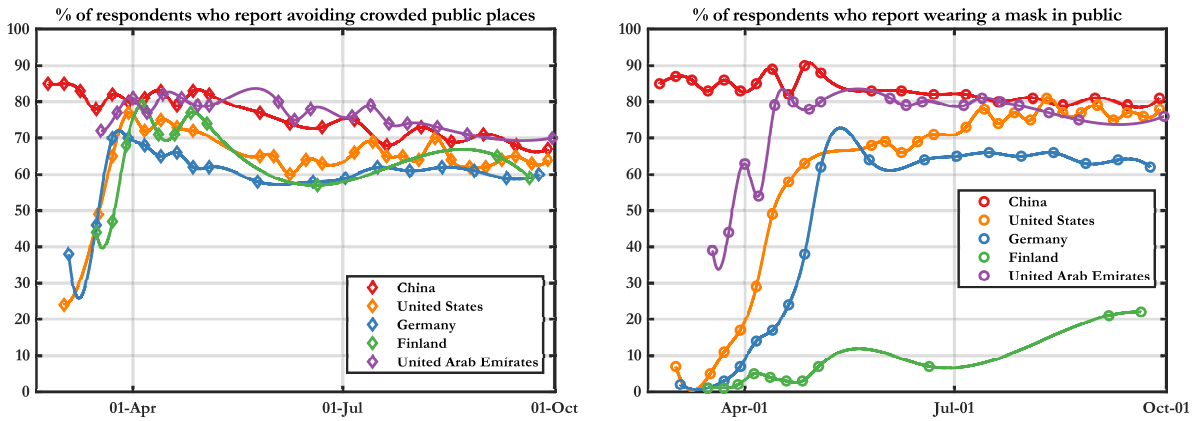

Figure 1: Protective behaviors in a sample of countries during the COVID-19 pandemic. Left Panel - Percentage of survey respondents who report avoiding crowded places. Right panel - Percentage of survey respondents who report wearing a mask in public. Source: YouGov.

individuals, who carry the disease and can transmit it; resistant individuals, who can neither contract nor transmit the disease.

Agents get to engage in bilateral interactions with randomly-chosen partners. Each interaction may or may not generate utility, reflecting the fact that some contacts may not be desired. When a susceptible agent enters in contact with an infectious agent, the former contracts the virus with some probability. The measure of infectious individuals grows from the flow of previously-susceptible individuals newly infected, and shrinks from the flow of individuals who recover from the disease. Infection dynamics are then driven by the effective reproduction number: the number of susceptible individuals expected to contract the disease from the same infectious agent while she is a carrier.

I first analyze equilibrium outcomes when agents are given the opportunity to stay home instead of engaging in social and economic interactions. Second, I assume that agents cannot opt out of the market but can wear a mask, which diminishes the probability of contracting (or transmitting) the virus during social interactions. The two decisions significantly differ in the trade-offs they imply: self-isolation offers absolute safety, but comes at the relatively large opportunity cost of forgoing social engagement; mask-wearing only offers a partial protection and carries an inconvenience cost, but still allows the wearer to engage in social and economic activity.

Throughout the paper, I study two limiting cases regarding the status of agents who recover from the disease. In a first specification, labeled SIS, it is assumed that agents

do not gain immunity, so that they are again susceptible after they recover. Under this assumption, the equilibrium reduces to a system of two ordinary differential equations that can be analyzed with phase diagrams. A second specification where recovered agents become resistant is also studied. This adds a third differential equation to the model, which is then calibrated to US data and solved numerically.[3]

In terms of calibration, while I mostly follow the method developed by Farboodi et al. (2020), I depart from them by making use of micro-founded data to calibrate two important parameters: the meeting rate and the transmissibility of the virus during an interaction. The former is calibrated using survey studies that document the number of interpersonal contacts experienced by respondents on a daily basis. The latter is calibrated following contact tracing studies, which track the contacts of individuals who have tested positive to the virus, and record whether those contacts, who have been asked to isolate, contracted the virus.

A first important result is that in a world where utility is directly derived from contacts between individuals, there exist complementarities between the participation decisions of susceptible agents. Since matching among agents is random, the risk of infection in a given contact depends on the composition of the pool of participants. More specifically, as infectious and resistant agents always participate, the probability of a susceptible agent contracting the virus in a given contact decreases the more susceptible peers participate. Equilibrium participation then becomes the outcome of a game between susceptible agents, whereby multiple Nash equilibria may coexist. For example, it could be rational for a susceptible individual to participate if all other susceptible agents participate, and to stay home if all other susceptible agents stay home. This gives rise to adverse selection: the fewer agents not carrying the virus participate, the higher the prevalence of infection in the pool of participants (i.e., the lower the "quality" of the pool), and the lower the net benefit of participating, which further drives non-carriers out of the market and increases the prevalence of the disease among participants.

The complementarities between the decisions of susceptible agents translate to a multiplicity of equilibrium paths in both the SIS and SIR specifications as long as the cost suffered by infected agents is in a medium range, a condition satisfied in the

[3] As of October 2020, while there is still disagreement regarding a definitive immunity in individuals who have recovered from COVID-19, there exists a body of evidence pointing at temporary immunity for the majority of cases. See references in https://www.nytimes.com/2020/08/16/health/coronavirus-immunity-antibodies.html, retrived on October 5, 2020.

calibrated model. I restrict my attention to classes of equilibria that satisfy some specified coordination rules for each instant along the equilibrium path where multiple Nash equilibria coexist.[4] I first consider two extreme rules, where susceptible agents either always coordinate to participate or to stay home whenever both could hold in equilibrium. The paths obtained provide bounds to all other equilibrium paths in the (S, I) phase plane. Following either rule, the infection curve is considerably flatter than in the benchmark case with no behavioral response, where the infection curve reaches a peak with around 40% of the population infected at the height of the epidemic. When agents coordinate to go out and engage in interactions whenever possible, the measure of infected agents never surpasses 7% of the population. At the other extreme, when agents coordinate on staying home as much as possible, it never goes past 2%.

I then consider other coordination rules. One specifies that in the multiplicity region, susceptible agents coordinate on going out with probability $x \in (0, 1)$. The lower x, the flatter the infection curve and the more delayed the development of the epidemic. Another rule is based on the idea that individuals may be more likely to coordinate on staying home when the epidemic seems more severe, so that coordination is determined by comparing the number of active cases to a set threshold. Using this rule, infection curves feature plateaus. The last rule I impose has to do with "isolation fatigue." After coordinating on the safe behavior of staying home for a long time, susceptible agents may get fatigued and switch to coordinate on participating. This coordination rule allows the equilibrium path to feature multiple waves of infections. These results highlight that countries or regions with similar fundamentals can still experience significantly different infection dynamics, driven by equilibrium beliefs.

Interestingly, across all equilibrium paths explored, as time goes to infinity, a relatively similar measure of agents will have been infected—between 78% and 80% of the population. In comparison, in the benchmark model, 96% of the population would eventually have been infected. A policy implication is that the shape of the infection curve is not necessarily, in itself, a good measure of how well a population is faring in terms of long-term outcomes: widely different shapes could eventually lead to similar steady states. Coordination is nevertheless extremely relevant when it comes to welfare, as the economic and social costs of forgone social contacts vary largely across equilibria.

When agents cannot opt out of the market but can decide to wear a mask, equilib-

[4] These can be seen as restricting agents' beliefs.

rium analysis is much simplified. First, infectious agents have no incentives to take this costly precaution. Second, there are no complementarities (in a static sense) between the decisions of different susceptible agents: taking as given the future course of the epidemic, the net benefit of wearing a mask for a given susceptible agent is independent of the behavior of other susceptible peers. As a result, the equilibrium is unique. In the SIR simulations, the cost of wearing a mask is calibrated as a fixed percent of the utility received by agents when they engage in interactions. For reasonable calibrations, the equilibrium path is such that susceptible agents do wear a mask once the epidemic has gained enough ground, and they stop doing so once it has sufficiently subsided. As expected, the costlier the masks, the shorter the amount of time during which they are worn. The infection curve is again considerably flattened, and a bit delayed.

Across all specifications, the number of active cases never gets past 12%. The cumulative measure of agents that has been infected by the end of the epidemic remains between 72% and 75% for the different cost specifications. In terms of welfare, the mask-wearing margin yields better outcomes than the participation margin. Not only is the long-run cumulative number of cases even lower than for the participation model, the associated costs are extremely low. This translates to a total welfare loss between 5.7 and 6.1 trillion dollars for the model with mask-wearing, compared to 7.6 trillion dollars in the benchmark with no behavior response.

Relation to the literature A large body of economic literature aimed at endogenizing the dynamics predicted by epidemiological models quickly developed in the wake of the COVID-19 outbreak.[5] This paper is most-closely related to a subset of those papers, which endogenize individual-level participation in a SIR model, using forward-looking rational agents who maximize their lifetime utility: Bethune and Korinek (2020), Farboodi et al. (2020), Garibaldi et al. (2020), McAdams (2020a), and Toxvaerd (2020).[6]

In those five papers, and different from the present paper, an agent's utility is not directly derived from each social interaction but from her level of "social activity" (a continuous variable). Additionally, in all but McAdams (2020a), that utility is inde-

[5]There did already exist a small economic literature related to infectious disease, spurred by the HIV outbreak in the 1990s. McAdams (2020b) provides a comprehensive review of economic epidemiology, with a focus on recent developments but also going back to those seminal papers.

[6]Note that Bethune et al. (2020) also studies an SIS specification. Other papers that endogenize economic activity in a SIR model at a more aggregate level include for example Eichenbaum et al. (2020) and Krueger et al. (2020).

pendent of the "social activity" of other agents. The main implication is that there are no complementarities between the participation decisions of different agents in the economy, which are essential to generate the infection dynamics obtained in the present paper. Indeed, in my paper, because utility stems from each individual contact, it inherently requires meeting other agents, and thus directly depends on the participation of other agents. Similar complementarities are highlighted by McAdams (2020a), written concurrently. In that paper, utility is specified as depending on the participation of other agents in reduced form, with a utility function that increases in aggregate participation. In other words, complementarities in participation decisions are directly built in. Both models predict equilibrium multiplicity, and highlight the role of coordination. What differs is that while McAdams focuses on the theory, I calibrate the model and explore the form that multiplicity takes when applied to COVID-19 in the US. I simulate paths at the two extremes of participation, when agents participate as much and as least as possible, and quantify the corresponding range of human and economic costs. I also explore additional coordination rules and show how they can give rise to infection dynamics such as plateaus and multiples waves, which can be observed in the data but are absent from Bethune and Korinek (2020) and Farboodi et al. (2020), two models that are also calibrated to the COVID-19 epidemic in the US.

The information structure also differs across the aforementioned papers. In the first two, it is assumed that agents do not know whether they are susceptible or infectious. In Garibaldi et al. (2020) and Toxvaerd (2020), like in my paper, agents know their status. In McAdams (2020a), agents who contract the virus originally do not know it, but may eventually learn it. When applied to the COVID-19 pandemic, the latter specification certainly seems the most appropriate, as 40% of infections are asymptomatic (Oran et al., 2020). Due to the structure of my model, however, uncertainty would require keeping track of the distribution of beliefs over time, depending on the exact history of matches encountered, a challenging problem left for future work.

The second part of the paper focuses on mask-wearing, a decision absent from the five papers mentioned above. This decision can be seen as a specific example of behavioral reaction along a "vigilance" margin, which is for example present in Keppo et al. (2020). Vigilance is costly, but decreases one's risk of infection (as well as others'). A major difference is that in Keppo et al. (2020), agents are partially myopic: they maximize an objective function that only depends on the state of the epidemic at that time. In the present paper, agents are perfectly forward-looking. Salanié and Treich

(2020) combine the two margins studied in this paper, mask-wearing and isolation, and focus on the impact of a mandatory policy related to mask-wearing on agents' isolation behavior, however in a static setting.

The adverse selection that occurs in the model with participation is similar to a mechanism highlighted by Kremer et al. (1996) in the context of HIV, where participation decisions impact not only the number of matches but also the composition of the pool: as more individuals choose abstinence, the prevalence of infection may increase.

The rest of the paper is organized as follows. Section 2 presents the model environment. Section 3 describes the infection dynamics predicted by standard SIS and SIR models absent the participation and mask-wearing margins, thus serving as a benchmark against which to compare subsequent results: outcomes when adding the participation margin, in Section 4, and outcomes when adding the mask-wearing margin, in Section 5.

2 Environment

Time t is continuous and goes on forever. The economy is populated by a measure P of infinitely-lived agents who discount the future at rate $r > 0$. At all points in time, agents can choose to engage in a meeting process. Meetings, or "social contacts," are bilateral and occur at random with a Poisson arrival rate $\alpha(N)$, where $N \leq P$ represents the measure of participating agents. When an agent enters in contact with another agent, she enjoys $y > 0$ utils with probability p and 0 otherwise—not all meetings may be desirable.[7]

We consider the existence of a virus that can spread in the population. Agents can be in one of three states: susceptible, infected, or resistant.[8] The measures of agents in each state are respectively denoted $S(t)$, $I(t)$, and $R(t)$, with $S(0)$, $I(0)$ and $R(0)$ taken as given. An agent's state is private information, exclusively known by that agent.

Agents who participate to the meeting process can take a protective measure, interpreted as wearing a mask, at a flow cost $k > 0$. The decision is made before any contact is realized. Upon contact with an infectious agent j, a susceptible agent i becomes infectious with probability $\tau^{ij} \in (0, 1)$, with $\{i, j\} \in \{m, n\}^2$. The superscripts

[7] Garibaldi et al. (2020) highlight the existence of those "unintended contacts" and explain their relation to the matching technology.

[8] In this model, being infected and infectious are strictly equivalent.

denote whether each agent in the meeting is wearing a mask (m) or not (n). It is assumed that $\tau^{mm} < \tau^{nm} < \tau^{mn} < \tau^{nn}$. While masks are most effective when both agents wear them, the second best occurs when only the infected agent wears one.

Infectious agents, who recover at Poisson rate $\gamma > 0$, suffer a flow cost $\psi > 0$ as long as they are infected. The flow cost of infection represents both the direct, physical cost of being sick, and the indirect cost associated with the prospect of dying from the disease.[9]

Throughout the paper, I study and compare two assumptions regarding the status of agents who recover from the virus. In one case, it is assumed that they gain immunity and become resistant, an absorbing state. In the other case, it is assumed that they transition back to being susceptible. The first case will be referred to as the SIR model and the second as the SIS model.

Preliminary results: equilibrium in a virus-free economy When $I(t) = 0$, the population is free from the virus. In this case, there are no incentives to refrain from participating in the meeting process, and no incentives to wear a mask during those meetings. The lifetime discounted utility of an agent in state j, where $j \in \{S, I, R\}$, is denoted V_j and is given by the Hamilton-Jacobi-Bellman (HJB) equation,

$$rV_j(t) = \alpha(P)\tilde{y} + \dot{V}_j(t), \tag{1}$$

where the dot represents a time derivative. Because the whole population participates to the meeting process, $N = P$, and agents meet at Poisson arrival rate $\alpha(P)$. Each meeting provides an expected benefit of $\tilde{y} \equiv py$ utils. The equilibrium path is such that V^j is constant and equal to the present value of all future meetings, $\alpha(P)\tilde{y}/r$.

3 A quick primer on epidemiological models

It will be helpful to first review the dynamics of the model when neither the participation nor the mask-wearing decision are taken into account, which will later serve as a benchmark. As such, the model looks like an off-the-shelves SIS/SIR model, where the number of participants is equal to the population size, $N = P$, and the transmission probability corresponds to that when neither agent wears a mask, τ^{nn}. At any point

[9] Note that death is not formally modeled otherwise. Agents always recover at time goes to infinity, so that the population remains constant.

in time, agents are either susceptible, infectious, or resistant, so that the following identity must hold,

$$P = S + I + R, \tag{2}$$

where, to simplify the exposition, I suppressed the explicit dependence of S, I and R on time. The measure of infected agents evolves according to

$$\dot{I} = \alpha(P)\tau^{nn}S\frac{I}{P} - \gamma I. \tag{3}$$

The first term on the right-hand side corresponds to the inflow of susceptible agents newly infected. S agents enters in contact with $\alpha(P)$ other agents, a proportion I/P of which is infected. For each of these contacts with infectious agents, the probability for the susceptible agent to catch the virus is τ^{nn}. The second term on the right-hand side corresponds to the outflow of previously-infected agents that recover. In a SIS model, $\dot{R} = 0$, so that R is constant. Assuming that the original stock of resistant individuals is null, in the SIS specification, $R = R(0) = 0$. The law of motion for S is equal to the negative of (3). In a SIR model, $\dot{R} = \gamma I$ and $\dot{S} = -\alpha(P)\tau^{nn}SI/P$.

We denote $\sigma \equiv \alpha(P)\tau^{nn}/\gamma$ the basic reproduction number. This number is often referred to as $\mathcal{R}_0$, but the notation σ is preferred here to avoid any confusion with the initial measure of resistant agents, $R(0)$. It corresponds to the number of people that an infectious individual would be expected to infect before recovering, *assuming that the whole population is susceptible*. It has to be distinguished from the effective reproduction number, $\sigma_e(t) \equiv \sigma S(t)/P$, a time-dependent variable which measures the number of people an infected individual would be expected to infect *given the actual measure of susceptible agents in the population at time t*.

Steady states In the SIS and SIR models, the system is at steady state when $\dot{I} = \dot{S} = \dot{R} = 0$. Plugging $\dot{I} = I = 0$ into (3), we directly obtain that in the SIS model, there always exists a virus-free steady state, with $I^* = 0$, $S^* = P$ and $R^* = 0$. There also exists an endemic steady state, with $I^* > 0$, as long as $\sigma > 1$. It is such that $I^* = P\left[1 - 1/\sigma\right] > 0$, $S^* = P/\sigma$ and $R^* = 0$. In this steady state, the flow of new infections exactly offsets the flow of recovered agents at each instant. In other words, in an endemic steady state, the effective reproduction number must be exactly equal to one. Because the effective reproduction number is never greater than the basic

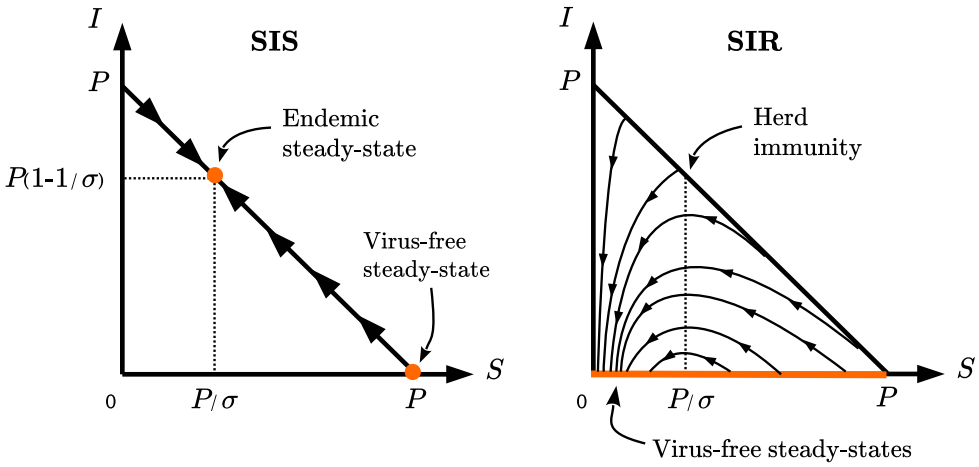

Figure 2: Dynamics of the SIS and SIR model with no participation nor mask-wearing decision, where $\sigma \equiv \alpha(P)\tau^{nn}/\gamma > 1$ is the basic reproduction number.

reproduction number, this requires the basic reproduction to be greater than one.[10]

In contrast, in the SIR model, there exists no endemic steady state. Intuitively, the SIR system cannot be at steady state as long as $I > 0$, as there would be a strictly positive increase in the measure of recovered individuals at each instant. As a result, in the SIR model, steady state is only achieved when the virus has entirely been eradicated, $I^* = 0$.[11] There exists a continuum of such virus-free steady states, indexed by $R^* \in [0, P]$, with $S^* = P - R^*$. In sum, any combination of S^* and I^* adding up to the total population can be sustained as a steady state.

Dynamics The dynamics of the SIS and SIR systems can be represented in phase diagrams. They are drawn in Figure 2 with S on the x-axis and I on the y-axis.

The left panel represents the dynamics under the SIS specification when $\sigma > 1$. Because we assumed $R = 0$, the pair (I, S) must always be located on the hypotenuse of the triangle. The number of infections increases (decreases) as long as the effective reproduction rate number is greater (smaller) than one, or equivalently, $S > (<) P/\sigma$. Once S reaches the threshold P/σ, the system has reached the endemic steady state. As a result, when the endemic steady state exists, it is globally stable. Starting from any arbitrarily low $I(0)$, the infection grows until it reaches the steady state. When

[10] Mathematically, for $\sigma_e = \sigma S/P = 1$ to hold given $S/P \in [0, 1)$, we need $\sigma > 1$.

[11] This result can be overturned by introducing birth and death dynamics to the system, in which case an endemic steady state can be sustained.

the endemic steady state does not exist, i.e., when $\sigma < 1$, the effective reproduction number is mechanically smaller than one, and the virus-free steady state is globally stable. Starting from any $I(0)$ the number of infections decreases until the virus is eradicated.

The dynamics under the SIR specification are represented in the right panel. Whether the measure of infected agents increases or decreases is still determined by the effective reproduction number: I increases (decreases) when $S > (<) P/\sigma$. However, S can only decrease, as contrary to the SIS model, there is no inflow of previously-infected susceptible agents. As a result, the effective reproduction number is strictly decreasing over time and must eventually drop below one regardless of its initial value. At this point, we say that the population has reached "herd immunity," i.e., there are not enough susceptible individuals left for each infected individual to infect more than one susceptible agent before recovering. Starting from an arbitrarily low $I(0)$ and along the hypotenuse ($R = 0$), the initial effective reproduction number must be greater than one, so that the stock of infected first increases. Once herd immunity has been reached, the measure of infected agents then steadily decreases until the system reaches the (virus-free) steady state. Steady states with $S > P/\sigma$ are not locally stable, and therefore could never be reached starting from any $I(0) > 0$, because the effective reproduction number as we approach $I = 0$ would be greater than one. Equivalently, only steady states such that $R^* \in [P(1 - 1/\sigma), P]$ can be reached. Because R^* indicates the aggregate measure of agents that have been infected during the epidemic, it means that at the minimum, a share $(1 - 1/\sigma)$ of the population would be infected over the course of the epidemic. Finally, it can be shown that R^* and S^* are uniquely determined by $I(0)$. The higher $I(0)$, the higher R^* and the lower S^*.

In the remaining of the paper, we relax the assumption that the whole population seeks to engage in social interactions at all times and allow participants to wear masks, so as to study the implications of those endogenous responses on both the dynamics of the epidemic and long-run outcomes. We will focus on the parameter region such that the basic reproduction number, σ, is greater than one: absent any behavioral response, the virus would become endemic in the SIS model, and starting from low enough $I(0)$, there would be an epidemic in the SIR model.

4 To go or not to go

In this section we focus on the participation decision of agents in the market. It is assumed that there is no protective measure, i.e., agents do not wear masks, and we let $\tau \equiv \tau^{nn}$ to simplify the notation. We first set up the SIS model and derive a few key analytical insights in Section 4.1, before moving on to the SIR model in Section 4.2. The SIR model is then calibrated and solved numerically in Sections 4.3 and 4.4.

4.1 SIS model

The HJB for an infectious agent is

$$rV_I = \alpha(N)\tilde{y} - \psi + \gamma(V_S - V_I) + \dot{V}_I, \tag{4}$$

where V_I is the expected lifetime discounted utility of being infectious and V_S that of being susceptible. The first term makes use of the equilibrium result that infectious agents would always choose to participate to the meeting process. While they could choose to stay home, it would never be rational for them to do so, as it carries an opportunity cost but no benefit.[12] The cost of being infected is captured by the second term, while the third term captures the potential upside of recovery. The HJB equation for susceptible agents is

$$rV_S = \max\left\{0, \alpha(N)\left[\tilde{y} + \tau\frac{I}{N}(V_I - V_S)\right]\right\} + \dot{V}_S. \tag{5}$$

The maximization represents the participation decision of the susceptible agent. She can decide to stay home, in which case her utility is normalized to zero, or she can decide to go out, in which case she gets utility from a fraction of the social contacts she will encounter, but also faces the downside of potentially getting infected. Due to random matching and full participation from infectious agents, the probability of a partner being infectious is I/N. Conditional on meeting an infectious agent, the probability of contracting the virus is τ.

[12] This would be different if preferences favored altruism, e.g., infected individuals could suffer a cost from endangering or infecting susceptible individuals, or if infectious individuals did not know their own state, in which case they may also fear getting infected. Additionally, for completeness, note that it could in fact be rational for an infectious agent to stay home if all other agents stay home as well. Going out would bring no benefit, so that the infectious agent would be indifferent. This is a pure coordination problem and we will ignore this type of equilibrium in the remaining of the paper.

Equations (5) and (4) can then be combined to obtain a single differential equation in $\omega \equiv V_S - V_I$,

$$\dot{\omega} = (r+\gamma)\,\omega + \alpha(S^p + I)\min\left\{\tilde{y}, \tau\frac{I}{N}\omega\right\} - \psi, \tag{6}$$

where S^p denotes the measure of susceptible agents who participate. A second differential equation comes from the law of motion for the measure of infected individuals,

$$\dot{I} = \alpha(N)\tau S^p\frac{I}{N} - \gamma I. \tag{7}$$

It is almost identical to (3), derived in the benchmark SIS model. The difference is that is now depends on the measure of participating susceptible agents, S^p, rather than the measure of susceptible agents, S (both directly and through N).

To close the model, we need to solve for the aggregate participation of susceptible agents, S^p. To do so, we first solve for the participation decision problem of a single susceptible agent, j. She engages in social interactions with probability a_j, given by

$$a_j \begin{cases} = 0 & < \\ \in [0,1] \text{ if } \tilde{y} & = \ \tau\frac{I}{N}\omega, \\ = 1 & > \end{cases} \tag{8}$$

where the aggregate participation of the whole population, N, is given by

$$N = S^p + I^p = \int_{i\in S} a_i di + I. \tag{9}$$

To decide whether to participate, the susceptible agent compares the expected utility from a social contact, $\tilde{y}$, to the expected chance of receiving the virus multiplied by the cost of getting infected, $\tau(I/N)\omega$. Because $N = S^p + I$, the chance of contracting the virus during a social contact depends on the participation decision of all other susceptible agents as well. Ceteris paribus, a higher number of susceptible agents shifts the composition of the pool of participants in such a way that I/N decreases, thereby reducing the riskiness of the pool and making any given interaction safer. In other words, an increase in the number of susceptible agents participating, keeping the number of infected participants the same, reduces the marginal cost of a contact without impacting its marginal benefit. This generates complementarities between

the participation decisions of susceptible agents: more participation from susceptible agents encourages the participation of other susceptible agents. Note that this feedback loop is independent of the assumptions made regarding the matching technology. In particular, it does not require increasing returns to scale.

Solving for S^p for a given pair (I, ω) is now akin to solving a Nash equilibrium, whereby the participation decision of each individual, driven by (8), must be the best response given the participation decisions of all other individuals. There are two dominance regions and one multiplicity region. When $\tilde{y} > \tau\omega$, there exists a unique Nash equilibrium, where all susceptible individuals participate, $S^p(I, \omega) = S$. Intuitively, in this region, the expected utility from any given social contact would be higher than the expected cost of infection even if all participants were infectious, so that the only possible equilibrium outcome is for everyone to participate. When $\tilde{y} < \tau(I/P)\omega$, there also exists a unique Nash equilibrium, although with no participation from susceptible individuals, $S^p(I, \omega) = 0$. In this region, the expected utility earned from a contact would be lower than the expected cost of infection, even if everyone were to participate, so that it can never be worth it for susceptible agents to participate. When $\tau(I/P)\omega \leq \tilde{y} \leq \tau\omega$, there exist multiple Nash equilibria: the two previous corner equilibria, as well as an equilibrium with partial participation, $S^p(I, \omega) = (\tau\omega/y - 1)\, I$. By construction, in the latter type of equilibrium, susceptible agents are indifferent between staying in or going out.

Equilibrium definition An equilibrium consists in a list of time paths for the two state variables $\{S(t), I(t)\}$ and the two control variables $\{S^p(t), \omega(t)\}$, such that (6) and (7) are satisfied, where S^p is a Nash equilibrium consistent with the individual participation decision rule given by (8) and (9), $S(t) = P - I(t)$ and $I(0)$ is given.

Steady states There always exists a virus-free steady state, where, by definition, $I^* = 0$ and $S^* = P$. It features full participation, $S^{p*} = N = S$, and is therefore identical to the virus-free steady state described in the benchmark SIS model. Because the virus is eradicated, there is no risk in participating. As long as other people are participating, it is therefore worthwhile to join.[13] In this steady state, $\omega^* = \psi/(r + \gamma)$.

There may also exist endemic steady states, which require $\alpha(P)\tau S^{p*}/(S^{p*} + I^*) =$

[13] As mentioned earlier, there could also be an equilibrium with no participation at all, $S^{p*} = N^* = 0$, but that would purely due to a coordination problem, not to the virus. In this scenario, susceptible agents would be indifferent between participating or not, so we assume that they do participate.

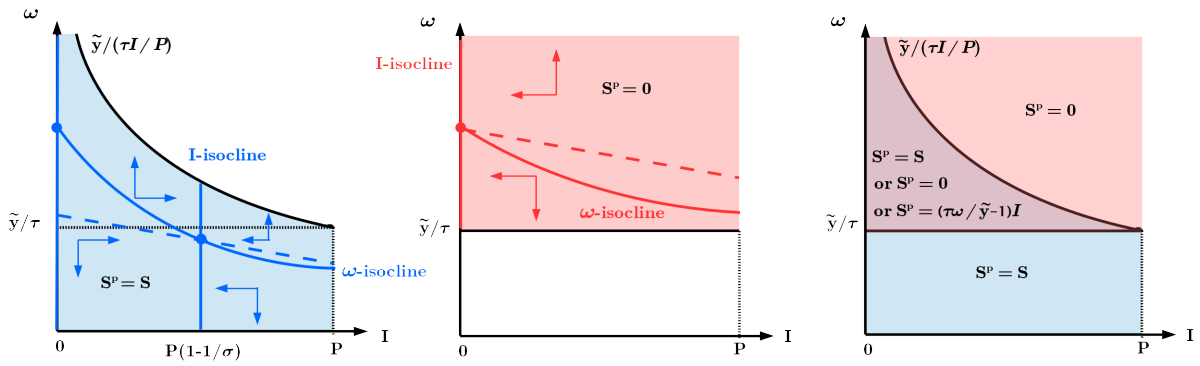

Figure 3: Construction of the phase diagram in the SIS model with participation.

γ. Now denoting the basic reproduction number $\sigma^p(S^p) \equiv \alpha(S^P + I)\tau/\gamma$ and the effective reproduction number $\sigma_e^p(S^p) \equiv \sigma^p(S^P)S^p/(S^p + I)$, this condition can be rewritten as $\sigma_e^p(S^p) = 1$. As was the case for the standard SIS model, the effective reproduction number must be equal to one for the system to be in a endemic steady state. What differs is that the effective reproduction number is now a jump variable and an increasing function of the participation of susceptible agents.

We obtain that $S^p = 0$ cannot hold in an endemic steady state, since the effective reproduction number would be equal to zero. An endemic steady state with full participation, $S^{p*} = S$, implies $S^* = P/\sigma$, $I^* = P(1 - 1/\sigma)$ and $N^* = P$. This steady state is identical to the endemic steady state in the standard SIS model. Note that $\omega^* = \psi/[r + \alpha(P)\tau]$, so that this steady state exists as long as $\tilde{y} \geq \tau\{1 - \gamma/[\alpha(P)\tau]\}\psi/[r + \alpha(P)\tau]$, i.e., the cost of infection ψ and the transmission probability τ are not too high relative to the expected utility from a social contact, $\tilde{y}$. Finally, there may also exist an endemic steady state with partial participation, $S^{p*} = (\tau\omega^*/\tilde{y} - 1)I^*$, as long as the cost of infection, ψ, is not too high nor too low (although it requires $\alpha'(.) > 0$).

Dynamics We can now study the dynamics of the SIS model with participation in a phase diagram. Figure 3 displays its construction, with I on the x-axis and ω on the y-axis (the second state variable, S, can be obtained from the identity $S + I = P$).

As depicted in the right panel, three regions can be delineated. In the blue region, $\tilde{y} > \tau(I/P)\omega$, so that there exists a Nash equilibrium with full participation. The blue arrows in the left panel depict the direction of motion in that region under full

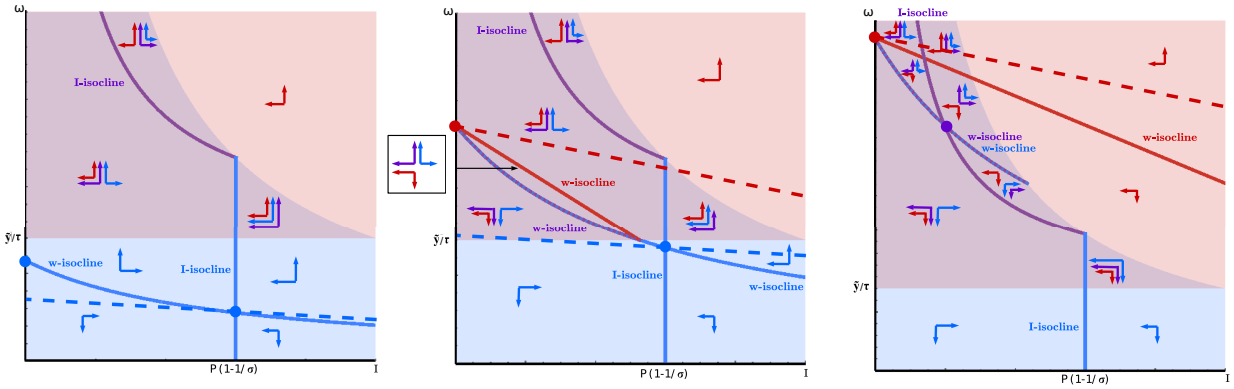

Figure 4: Dynamics of the SIS model with participation for three levels of the flow cost of infection, ψ. Left panel: low ψ. Middle panel: medium ψ. Right panel: high ψ.

participation, while the blue curves represent the isoclines, $\dot{I} = 0$ and $\dot{\omega} = 0$. The virus-free steady state, denoted by the blue dot on the y-axis, is unstable. This is intuitive: with full participation, the model is akin to the standard SIS model, and we saw that in a standard SIS model the virus-free steady state is unstable as long as the basic reproduction number is greater than one, which we assume. The endemic steady state, denoted by the blue circle at the intersection of the isoclines in the center of the region, is saddle-path stable. The saddle path is represented by the dashed blue curve.

In the red region, $\tilde{y} < \tau\omega$, so that there exists a Nash equilibrium with no participation. The red arrows in the middle panel represent the direction of motion in this region with no participation. Intuitively, without participation from susceptible agents, the measure of infectious agents always decreases. The virus-free steady state, represented by the red dot, is saddle-path stable. The saddle path is given by the dash red curve.

In the purple region, located exactly where the blue and red regions overlap, $\tau(I/P)\omega \leq \tilde{y} \leq \tau\omega$, so that the two corner Nash equilibria coexist with the partial participation equilibrium. In this region, the direction of motion therefore depends on whether susceptible agents coordinate on full participation, no participation, or partial participation.

To study equilibrium dynamics, we must combine the three graphs presented in Figure 3. The regions in which isoclines intersect, which depends on parameter values, will determine the set of equilibrium paths. Figure 4 displays the complete phase diagrams for three different values of the flow cost of infection, ψ.

As represented in the left panel, when the cost is low enough, the endemic steady state features full participation and is the only stable steady state.[14] The equilibrium path is unique and follows the saddle path. Dynamics are then exactly identical to those described under the benchmark model. Because getting sick only carries a small cost, agents do not react by adjusting their participation margin despite having the opportunity to do so.

When the cost of infection is in a medium range, as shown in the middle panel, both an endemic steady state with full participation and a virus-free steady state may be stable, depending on the coordination of susceptible agents. In this parameter region, there exists a large multiplicity of equilibria. For example, agents can still coordinate on participating throughout, in which case the dynamics are identical to those described in the left panel. At the other extreme, susceptible agents can coordinate on staying home as the number of infected asymptotically converges to zero. Agents can also switch coordination along the way, in which case the equilibrium paths for S and I may no longer be monotone.

Finally, when the cost of infection is very high, so that there does not exist an endemic steady state with full participation anymore, the only stable steady state is the virus-free steady state, as visible in the right panel.

Takeaways and discussion We can now highlight a few takeaways and discuss some of the key underlying assumptions.

1) Endogenous, policy-free equilibrium lock-down. When the cost of infection is high enough, the participation of susceptible individuals endogenously drops in equilibrium. This can be interpreted as an endogenous lock-down, which does not stem from any policy interaction but purely from individuals' fear of being infected. Ignoring this margin would then lead to overestimating the rate at which the virus is transmitted. This echoes results from other papers that also add a participation margin to SIR/SIS models, e.g. Bethune and Korinek (2020), Farboodi et al. (2020), Garibaldi et al. (2020), McAdams (2020a), and Toxvaerd (2020).

2) Impact of individual behaviors on the effective reproduction number. In the standard SIS model, the effective reproduction number is a state variable that evolves with the measure of susceptible agents. When allowing for the participation margin, it becomes

[14] More precisely, this occurs as long as the intersection of the blue saddle path with the y-axis occurs for $\omega \leq \tilde{y}/P$.

a jump variable that positively depends on the participation of susceptible individuals through two channels. First, the participation margin impacts the composition of the pool of agents an infectious individual is expected to meet: the lower the participation of susceptible agents, the higher the chance of a random encounter leading to virus transmission. Second, the participation of susceptible agents may also impact the number of meetings if the matching technology does not feature constant returns to scale, i.e., $\alpha'(.) \neq 0$. With increasing returns to scale, a drop in the participation of susceptible agents leads to a smaller arrival rate of meetings, and therefore fewer opportunities for the virus to be transmitted.

3) *Complementarities in participation and social utility.* The participation decisions of susceptible agents are complementary to each other, a phenomenon that is absent from Bethune and Korinek (2020), Farboodi et al. (2020), Garibaldi et al. (2020), and Toxvaerd (2020). In those papers, an increase in the participation of other susceptible agents is irrelevant for the participation decision of a susceptible agent. This difference is rooted in the assumptions made regarding the way individuals receive "social utility." In the three papers just mentioned, social utility is not directly derived from individual contacts, but more so derived from what could be interpreted as "time spent outside." Agents compare the utility from spending a given amount of time outside to the overall probability of getting infected during that time. The former is independent of the number of people outside. The latter increases in the participation of infectious agents, but is independent of the participation of susceptible agents. In the present paper, agents directly gain utility from each contact—contacts are exactly what provides utility. Therefore, when susceptible agents decide whether to engage in those contacts or to stay home, they compare the expected utility from each given contact to the chance of being infected in each of those contacts. It seems that both specifications may be appropriate for different real-world activities. The former seems to appropriately describe activities such as walking one's dog or going to the theater. The benefits from engaging in those activities do not directly depend on meeting other individuals, but the more people in the street or in the theater, the higher the chance of being infected. Agents would always be at least weakly better off with less participation from other agents. In contrast, the specification used in the present paper may be more appropriate to describe gregarious behaviors that have been at the heart of the COVID-19 spread: family and religious gatherings, parties, etc. In those events, utility directly comes from meeting with people, so that there would be no benefit from engaging in those

activities alone. When deciding whether to attend or not, one must compare the utility from being around people to the chance of being infected.[15]

4) *Equilibrium indeterminacy and variety of infection curve shapes.* When the cost of infection is not too high nor too low, the complementarities described above renders the equilibrium indeterminate. Coordination then plays an important role both for short-run dynamics and for long-run outcomes. In the short-run, because the coordination of susceptible agents can switch at any time when in the purple region, we can observe a variety of infection curves, for example featuring one or several peaks, as the number of infected may get up and down successively. This is very different than the dynamics obtained in the typical SIS model, where the infection curve is always monotone increasing or decreasing. In the long run, both the virus-free and the endemic steady states may be reached. This highlights the role of norms, or coordination rules, in the population.

5) *Adverse selection and its welfare cost.* As long as the cost of infection is not too low, it is easy to see that optimally, a planner would like to quarantine the small measure of infected individuals from the very beginning of the epidemic, and keep that rule in force forever. Isolated infectious agents would gradually recover, such that in the long run, only a very small measure of infected agents would remain infected and miss out on trade. In equilibrium, if anything happens, it is the exact opposite: infected agents always participate, and susceptible agents are forced to stay home. This is very costly, since susceptible agents initially form the wide majority of the population, and when they are confined, their measure only increases as infectious agents recover. This outcome is especially unfortunate when it occurs in the region where a Nash equilibrium with full participation could also occur, as it then resembles an extreme case of adverse selection. While it would be rational for a susceptible agent to participate if all other susceptible agents were also participating, as susceptible agents drop out, the risk of infection grows relative to the benefit of going out, which encourages participants to drop out further more. The equilibrium features a full unraveling, where only infectious agents are left in the market. It is worth noting that even though as currently described, this mechanism relies on the assumption that agents know their epidemiological status, the emergence of adverse selection would robust to different information specifications,

[15]Note that complementarities in economic activity are also present in McAdams (2020a), where aggregate activity is an argument of agents' utility function. This can be seen as a reduced-form representation of the meeting mechanism described in the present paper.

as long as agents infrequently get informed about their status (through the development of symptoms or through testing, for example).

4.2 SIR model

We now assume that infected agents who recover become permanently resistant. The HJB equation for susceptible agents remains identical to (5). The HJB for infected agents is similar to (4) with the third term replaced by $\gamma(V_R - V_I)$, where V_R is the expected lifetime discounted utility of a resistant agent. The HJB for resistant agents is

$$rV_R = \alpha(N)py + \dot{V}_R, \tag{10}$$

where we anticipate the result that in equilibrium, resistant agents have no incentives to stay out. We can easily show that $V_R - V_I = V_R^* - V_I^* = \psi/(r+\gamma)$, which is then used to again obtain a differential equation in ω,

$$\dot{\omega} = r\omega + \alpha(N)\min\left\{\tilde{y}, \frac{I}{N}\tau\omega\right\} - \frac{r}{r+\gamma}\psi. \tag{11}$$

The laws of motion for I is identical to (7) and the law of motion for R is given by

$$\dot{R} = \gamma I. \tag{12}$$

Finally, the participation decision of a susceptible agent is still given by (8), where $N = S^p + I + R$.

Equilibrium definition An equilibrium consists in a list of time paths for the three state variables $\{I(t), R(t), S(t)\}$ and the two control variables $\left\{S^P(t), \omega(t)\right\}$ such that (11), (7) and (12) are satisfied, where S^p is a Nash equilibrium consistent with the individual participation decision rule given by (8) where $N(t) = S^p(t) + I(t) + R(t)$, $S(t) = P - I(t) - R(t)$, $R(0)$ and $I(0)$ are given.

Steady states As was the case for the benchmark SIR model, any steady state must be virus-free, since we would otherwise have $\dot{I} > 0$. Absent any infection risk, the steady state participation of susceptible agents must be full, $S^{p*} = S$. This implies $\omega^* = \psi/(r+\gamma)$. There is a continuum of such steady states, indexed by $S^* \in [0, P]$, with $R^* = P - S^*$, as in the standard SIR model. Contrary to the SIS specification,

adding a participation margin does not impact the set of steady states in the SIR model.
We now calibrate the model so as to study equilibrium dynamics.

4.3 Calibration

The SIR model is calibrated to the COVID-19 epidemic in the United States, at a daily frequency. There are three types of parameters to be calibrated: epidemiological parameters, that depend on the characteristics of the virus; parameters related to the meeting and matching of agents; and parameters related to costs and preferences. Calibrated values are summarized in Table 1.

Epidemiological parameters. The baseline probability of transmission of the virus in a given contact between an infected and a susceptible agent, $\tau \equiv \tau^{nn}$, is calibrated based on medical studies. A meta-analysis of 172 observational studies by Chu et al. (2020) predicts an infection probability of 3% after a contact between two individuals at a distance of three feet. Contact tracing studies in China by Bi et al. (2020) and Luo et al. (2020) track the secondary infections stemming from contacts with positive individuals, and respectively find infection probabilities of 6.6% and 3.7%. An intermediate value of 5% is chosen for the calibration, so that $\tau^{nn} = 0.05$. The recovery parameter is calibrated following Farboodi and al. (2020), who assume an expected recovery time of 7 days, implying that infectious agents recover at a Poisson arrival rate of 1/7. While many patients take more than 7 days to recover, this value is closer to the expected time during which infected agents can transmit the virus. Finally, to match the model's timeline to the real-world timeline, we need to calibrate time 0, when the first infection occurs (i.e., $I(0) = 1$). We follow a study by Worobey et al. (2020), who pin down the first case of coronavirus that eventually led to an outbreak in the US to mid-February, and set $t = 0$ to February 14, 2020.[16] For simplicity, it is assumed that at this time, no one was resistant to the virus, so that $R(0) = 0$, and $S(0) = P - I(0)$.

Matching parameters The matching rate is assumed to be $\alpha(N) = \alpha N$, implying that the flow number of meetings, αN^2, displays increasing returns in the number of

[16] Robustness checks with a later start date, similar to that used in Farboodi et al. (2020), are presented in Appendix B.

participants. Population P is calibrated to the 2019 estimate of the US population from the Census Bureau, around 328.24 million individuals. The matching parameter α is calibrated to match the average number of in-person social contacts experienced by individuals on a daily basis when there is no epidemic. A telephone survey of four counties in North Carolina by DeStefano et al. (2011) reports an average of 10 contacts per day. Feehan and Cobb (2019) survey Facebook users and find an average of 12 contacts per day. We pick the conservative value of 10, implying $\alpha P = 10$.[17]

Costs and preferences The daily discount rate is calibrated to match an annual discount rate of 5%, hence $r = 0.05/365$. The expected utility from a meeting, $\tilde{y} \equiv py$ is normalized to 1. The cost of infection is calibrated relative to this unit utility, following the method used in Farboodi et al. (2020). Because the expected cost of death for an infected individual trumps all other costs associated with the infection, they calibrate the cost of infection to reflect the cost of potentially losing one's life. Formally, $\psi/(r+\gamma) = \pi\nu$, where π is the infection fatality rate and ν the value of statistical life. The left-hand side represents the discounted expected cost of the disease: infected agents pay a flow cost ψ as long as they are sick. The infection fatality rate is calibrated to 0.0062. To calibrate the value of statistical life, Farboodi et al. (2020) follow Hall et al. (2020), who estimate that each remaining year of life is worth \$270,000, and that COVID-19 victims would on average expect 14.5 remaining years of life. This implies $\nu = \$3,915,000$. To convert this value to utils, first note that it is equivalent to say that individuals are willing to pay a lump sum of \$3,915, or a daily stream of $\$3{,}915r$, to avoid a 0.1% probability of death. With a median yearly consumption of \$45,000 (Hall et al., 2020), this means that individuals would be indifferent between a 0.1% probability of death and forgoing $3,915r \cdot 365 \cdot 100/45,000 = 0.435\%$ of their consumption, or

$$\frac{\alpha P \tilde{y}}{r} - 0.001\nu = \frac{(1 - 0.435/100)\alpha P \tilde{y}}{r}. \tag{13}$$

We can now solve for ν and obtain $\nu = 317,550$ utils, from which we get $\psi = (r+\gamma)\pi\nu \approx 272$.[18] Note that the previous analysis implies an exchange rate between utils and

[17] This is also in line with studies run in Europe, e.g. Mossong et al. (2008) find a lower bound of 7.8 daily contacts in Germany and 19.8 daily contacts in Italy. In a similar study run in Hong-Kong, Leung et al. (2017) find an average number of contacts of 8.

[18] This analysis implies that the entirety of an individual's consumption requires person-to-person contacts. Robustness checks where individuals still obtain some baseline utility when they stay home are presented in Appendix B.

Param.	Definition	Target/Sources	Value
P	Population	US population	328.24M
α	Matching parameter	Average number of daily contacts	10/P
τ	Baseline transmissibility	Medical studies	0.05
γ	Recovery rate	7 days of infection on average	0.14
ψ	Cost of infection	Expected cost of death	272
$\tilde{y}$	Individual benefit from interaction	Normalisation	1
r	Discount rate	5% yearly discount rate	0.05/365
$t = 0$	Initial infection date	Epidemiological studies	Feb. 14, 2020
σ	Basic reproduction number	Implied	3.5
\$ per util	Exchange rate	Implied	12.33

Table 1: Calibrated parameters for the SIR model with participation.

dollars of $3,915,000/317,550 = 12.33$ dollars per util.

It is worth noting that Farboodi et al. (2020) and Bethune and Korinek (2020) use different approaches than the present paper to calibrate the timeline and the infection rate $\alpha\tau$. I explain those differences and explore the robustness of my results to those different approaches in Appendix B.

4.4 Results and discussion

Due to the coordination problem among susceptible agents, there exists a large multiplicity of equilibria. I consider a subset of equilibria that obey some coordination rules—more precisely, paths that exogenously dictates whether the susceptible population coordinates in or out of the market for each instant where both can be Nash equilibria. Key equilibrium results are described below. More detailed results are available in Table 2 in Appendix A. A description of the algorithm used to solve the model can be found in Appendix C.

Randomized coordination rule I first focus on rules that specify that susceptible agents coordinate to participate with probability x, where $x \in \{0, 0.1, 0.2, \ldots, 1\}$. The two extreme cases, $x = 0$ and $x = 1$, correspond to equilibria where susceptible agents respectively always coordinate to stay home, and always coordinate to go out. Those two equilibrium paths are represented in a phase diagram in the left panel of Figure 5 in blue and in red, alongside the path that would be observed in a standard SIR model with no participation decision. Labeled "No behavioral response," it is represented by a dashed black line and will serve as a benchmark. Although uninformative about the time dimension, this phase diagram is helpful to understand the relation between the

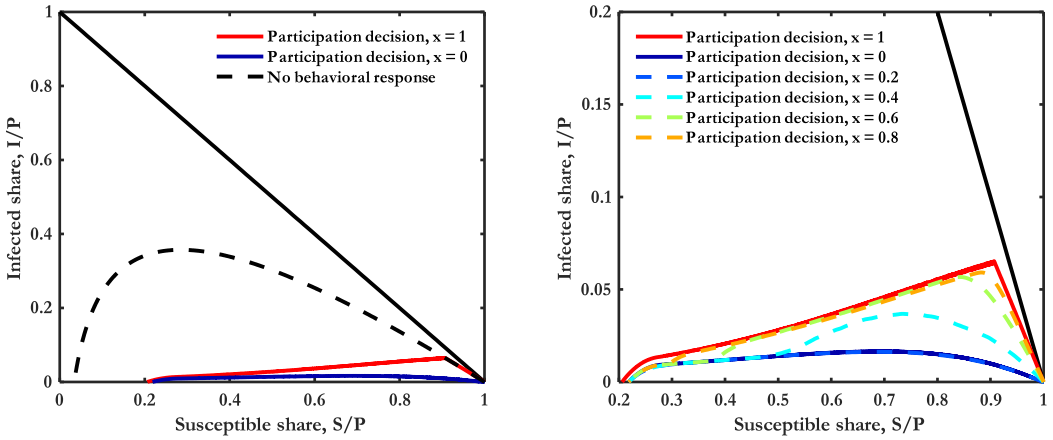

Figure 5: Equilibrium paths of the SIR model with participation decision, under the randomized coordination rule, represented in a phase plane. Left panel - Paths for the two extremes cases: in blue, agents always coordinate to stay home ($x = 0$), in red, agents always coordinate to go out ($x = 1$). The path with no participation decision, in dashed black, is provided for comparison. Right panel - Paths for the two extreme cases and intermediate cases. The graph is a magnified view of the bottom-right part of the full-size graph.

two state variables S and I. Starting from the bottom right corner, with $I(0) = 1$ and $S(0) = 0$, all three path originally feature an increase in the number of infected agents (and thus, mechanically, a decrease in the number of susceptible agents).

While the benchmark path and the path with maximum participation ($x = 1$) originally overlap, they diverge while I is still relatively low. In the benchmark case, the number of infected agents only starts to go down once the herd immunity threshold of $S/P = \gamma/(\alpha\tau P) = 0.2857$ has been reached. It then steadily declines and reaches the virus-free steady state with $S^*/P = 0.037$. Adding a participation margin allows for the number of infected agents to decrease much earlier, even when agents participate as much as possible. The virus-free steady state is reached with $S^*/P = 0.2061$. In other words, in this scenario, 79.39% of the population would have been infected by the end of the epidemic, compared to 96.63% in the benchmark case. Now looking at the other extreme case, where $x = 0$, we can observe slightly different dynamics. First, the relation between S and I is very flat, so that the measure of infected agents remains very low throughout the epidemic. However, it starts to declines for a lower S than was the case with $x = 1$, and eventually converges towards a virus-free steady state remarkably close to that resulting from the maximum equilibrium participation,

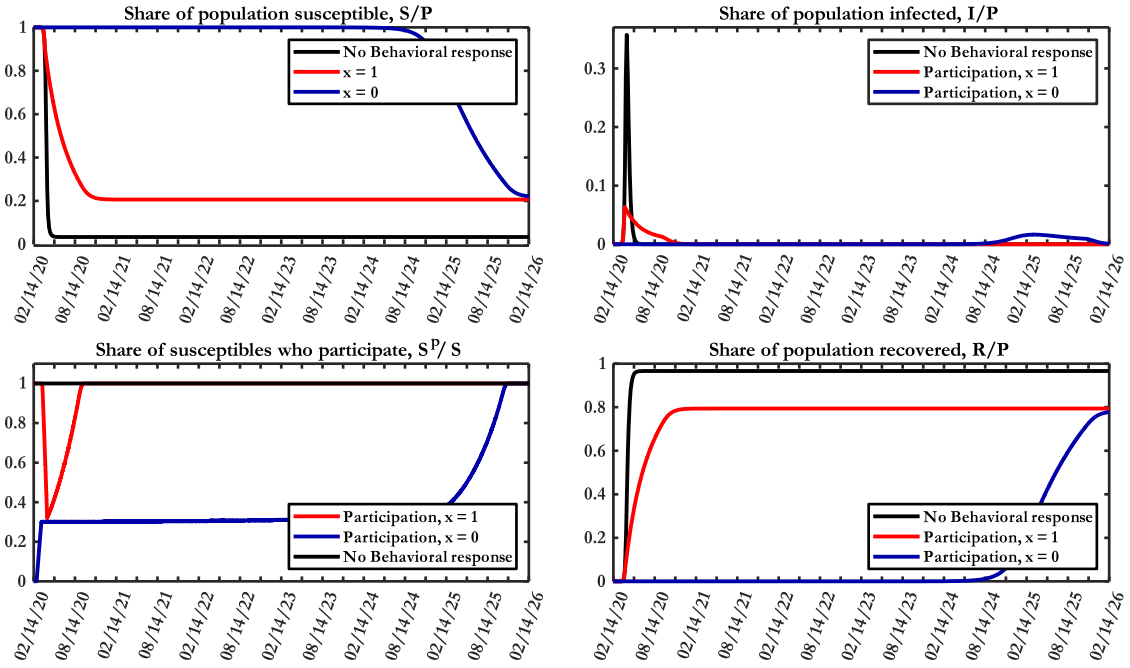

Figure 6: Time paths of epidemiological measures for the SIR model with participation, under the randomized coordination rule: susceptible agents coordinate to participate with probability x, for $x = 0$ and $x = 1$. Note that 3 months lapse between each tick on the x-axis, and that labels are plotted every 6 months.

with $S^*/P = 0.2185$, i.e., 78.15% of the population has been infected by the end of the epidemic.

To understand how those dynamics play out over time and how they relate to the participation decision, we can look at the four panels of time series in Figure 6. They run from February 14, 2020, when the first case is assumed to occur, to February 14, 2026. From the top-right panel, which displays the share of population infected over time, we can confirm that the addition of a participation margin allows to considerably flatten the curve, even when $x = 1$. In this case, the peak of infections only slightly goes over 6% over the population, against a peak at almost 40% of the population for the benchmark case. When $x = 0$, the curve is even flatter, and the measure of infectious agents remains below 2% of the population at all times.

The flattening of the curve is what allows for a much smaller number of individuals to be infected over the course of the epidemic. This can be seen directly by looking at the bottom-right panel, which shows the share of population recovered as a function of time, since the share of population recovered as time goes to infinity must be equal

to the share of population that has been infected in total. We can confirm that it is considerably smaller when the participation margin is introduced.

The evolution of participation over time is displayed in the bottom-left panel. Because of the discrete nature of the algorithm used to solve the model, values for S^p often jump between 0 and S when in the multiplicity region. As a result, participation is plotted as a rolling average with ten-day windows.[19] In the equilibrium path with the highest participation, $x = 1$, we can see that originally, all susceptible agents participate. By mid-March, around one month after the first infection, participation sharply drops. It reaches a low point at 30% of its regular level by mid-April. During this time, the peak of the epidemic is reached, and the number of infections starts declining. Participation then gets back up, and the economy reaches full participation again by September 2020. In the equilibrium path with the lowest participation, $x = 0$, participation drops to zero from day one. It then climbs back to reach a bit less than 30% of the baseline participation level by the beginning of March, and remains stable at this level for many months. Only around May of 2024 does it start to slowly increase again, the economy getting back to full participation by November 2025. As a result, when the coordination rules is such that susceptible individuals stay home, the epidemic is not only much flatter, it is considerably delayed. This is visible in the top-right panel of Figure 6.

We now turn our attention to welfare measures. The expected human toll of the epidemic is displayed in the left panel of Figure 7.[20] As expected, because the coordination rule bears little impact on the total number of agents infected by the end of the epidemic, its impact on the expected number of fatalities is also limited. The model predicts around 1.6 million fatalities for both $x = 0$ and $x = 1$, compared with almost 2 million without the participation margin. However, whether agents coordinate in or out does have a very large impact on the welfare costs of the epidemic, displayed in the right panel of Figure 7. The total costs, represented by full lines, can be broken down between the direct costs borne by infected agents and the opportunity costs due to susceptible agents staying home. When agents coordinate on going out, $x = 1$, total

[19] This explains why interior values can be obtained even though, in each period, there is either full or no participation.

[20] Even though the model does not feature deaths per se, in that the population always remains constant, the cost of infection ψ takes into account the expected cost of death for infectious agents. The expected number of fatalities is computed by multiplying the infection fatality rate, π, by the number of recovered agents.

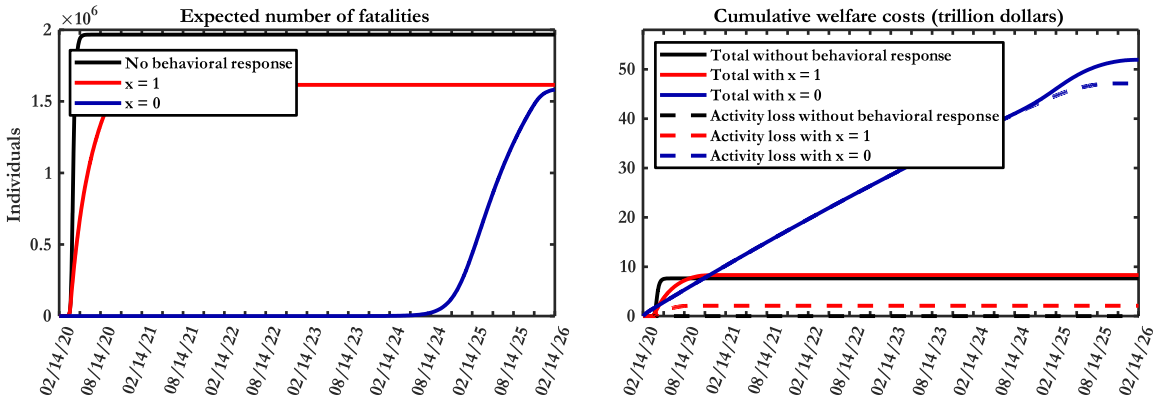

Figure 7: Time paths of expected number of fatalities and cumulative welfare losses for the SIR model with participation, under the randomized coordination rule: agents coordinate to participate with probability x, for $x = 0$ and $x = 1$. Total welfare losses correspond to the discounted cumulative sum of sickness costs borne by infected agents and activity losses due to susceptible agents staying home. Note that 3 months lapse between each tick on the x-axis, and that labels are plotted every 6 months.

costs amount to 8.3 trillions dollars, roughly 40% of US GDP.[21] Three quarters of this cost are due to sickness costs. When agents coordinate on staying in, $x = 0$, total costs skyrocket to around 51 trillion dollars. While sickness costs are 1.5 trillion dollars lower than when $x = 1$, 90% of of the total losses stem from the loss in activity. Recall that in that scenario, the participation of susceptible agents drops to 30% of its usual level for months on end. While this delays the epidemic, it does little to reduce the total number of infections. As a result, the losses due to the reduction in activity trump the gains in sickness costs by far, and the equilibrium with $x = 0$ is much worse, from a welfare perspective, than the equilibrium with $x = 1$. Interestingly, the benchmark path is even slightly better in terms of welfare than the equilibrium path with $x = 1$. While sickness costs are only a bit larger, there is no activity loss whatsoever, leading to a better net outcome. This is due to the adverse selection problem described earlier, and further discussed in *takeaway 5* at the end of the section.

Until now, we focused on the two extremes cases of the coordination rule, $x = 0$ and $x = 1$. We can now look at intermediate cases. Note that for those cases, the realized equilibrium path depends on the random draws made over time. For this reason, for each intermediate case, 50 simulations were run and then averaged. Going back to a phase diagram representation on the right panel of Figure 5, we can see

[21] All welfare numbers are computed as the discounted sums of welfare losses from time 0.

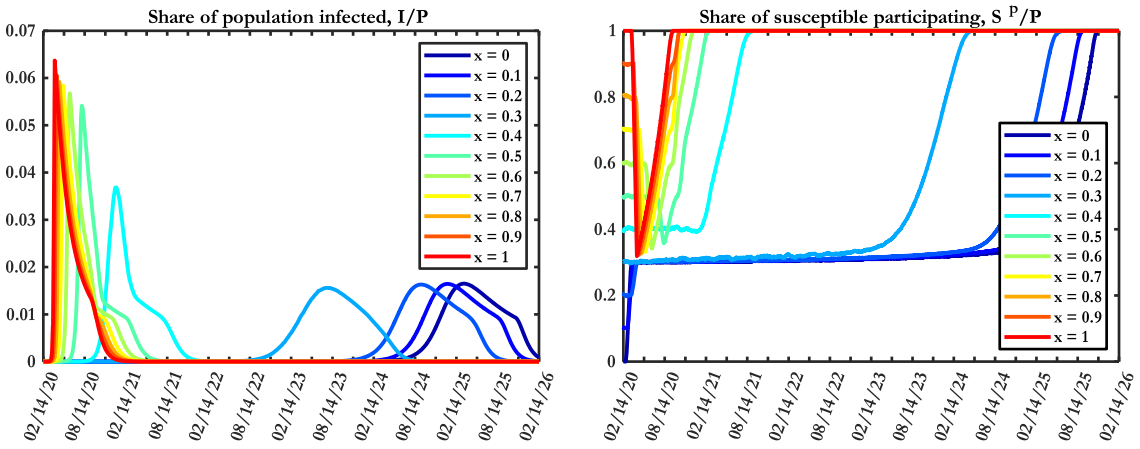

Figure 8: Time paths of share of infected population and share of susceptible population participating, under the randomized coordination rule: agents coordinate to participate with probability x, for $x = 0, 0.1, ..., 1$. Note that 3 months lapse between each tick on the x-axis, and that labels are plotted every 6 months.

that intermediate cases populate the phase plane in between the two cases studied earlier. Because all paths converge towards virus-free steady states very close to those described for $x = 0$ and $x = 1$, we can already see that the long term epidemiological outcomes will remain similar across those intermediates cases, with around 78% of the population infected. Mechanically, the expected number of fatalities remains around 1.6 million across cases. Similarly, because those paths are "bound" by the two extreme paths, welfare outcomes will lie in between the two welfare outcomes described earlier.

Figure 8 shows the particular shapes taken by the infection curves (in the left panel), as well as the associated participation decisions (on the right panel), for all x. As the probability of coordination to go out, x, decreases, the infection curve flattens and shifts to the right. In terms of participation, for $x \leq 0.3$, participation first starts at $100x\%$, then increases to 30%, and finally gradually goes back up to 100% after some time. The lower x, the longer the time during which participation remains at 30%. For $x \geq 0.4$, participation starts at $100x\%$, drops below 40%, then gets back up to 100% relatively quickly. The higher x, the earlier and the deeper the drop, and the faster the recovery.

Alternative coordination rules While the probability-based coordination rule was picked as a straightforward rule allowing to populate the phase diagram between the two extreme rules of always or never coordinating to go out, an infinite number of

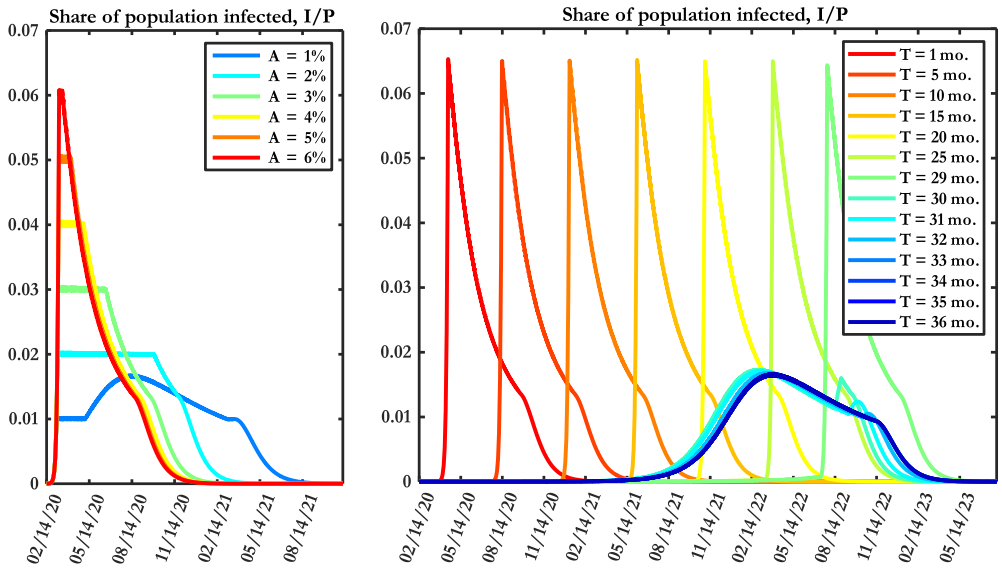

Figure 9: Time paths of expected number of fatalities and cumulative welfare losses for the SIR model with participation, under alternative coordination rules. Left panel - Susceptible agents coordinate to participate as long as $I < A$, and coordinate on staying in otherwise. Right panel - Susceptible agents coordinate to stay home as long as $t < T$, and coordinate to participate afterwards (fatigue specification).

alternative rules could be used. I present two such rules in this section.

The first alternative rule assumes that agents coordinate on staying in or going out depending on the measure of active infections cases, I. When this number is lower than some threshold A, agents feel safe and coordinate out. When it is higher, agents coordinate on staying home. Infections curves for $A \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6\}\%$ of the population are displayed in the left panel of Figure 9. As expected, when A is high, the infection curve resembles the one obtained when agents always coordinate to go out. Indeed, in this case, the measure of infected agents never reaches the threshold, so that the rule is never binding. In contrast, for low enough A, while the infection curve initially follows that of the $x = 1$ scenario, the number of infections then perfectly stabilizes once A active infections have been reached. Note that this does not mean that susceptible agents entirely stop going out. In that case, the number of infections would decline. Instead, susceptible agents go out exactly enough to maintain the effective reproduction number at one. It then eventually decreases, except for very low A. When the threshold is equal to 0.1% of the population, the number of infections initially stabilizes, but infections start growing again after a few months.

The second alternative rule assumes that individuals initially coordinate on staying home, which could be seen as the "good behavior" due to social norms, but eventually get fatigued and coordinate on going out from time T onward Infection curves for $T \in \{1, 5, 10, 15, 20, 25, 29, 30, 31, 32, 33, 34, 35, 36\}$ months are displayed in the right panel of Figure 9. Up to $T = 29$ months, fatigue arising later and later only shifts the infection curve to the right. For $T = 30$, that is, assuming that fatigue develops two and half year into the epidemic, an interesting phenomenon occurs. The start of the epidemic shifts back, in between the starts of the path with $T = 15$ and $T = 20$. Instead of a sharp increase in the number of active cases, the curve is much flatter. While it then resembles the curves observed earlier with low x, there are two big differences: first, the epidemic is not as delayed; second, there can be a second wave of infections. For example, for $T = 30$, while the number of infected individuals starts to go down around January 2022, it goes up again, this time more steeply, roughly six months later. While left for further exploration, it is easy to see that the existence of more than two infection waves can easily be obtained as equilibrium outcomes, conditional on picking the coordination rule appropriately.

Takeaways and discussion Key insights from the calibrated SIR model with participation are summarized below.

1) Equilibrium indeterminacy and multiplicity of infection dynamics. The calibrated model features an infinite number of equilibrium paths. The infection curve can be single-peaked like in the standard SIR model, but can also feature several peaks or even remain flat for some time. It can develop quickly, and disappear before the end of 2020, or be delayed for months and even years. Those dynamics are driven by the coordination of susceptible agents in the economy. While not explored in the paper at the moment, it could be interesting to consider factors that could impact the coordination of agents: country-wide experience with previous infectious disease and associated norms, policy messages and recommendations, etc.

2) Endogenous flattening of the infection curve. The idea of "flattening the curve" has consistently been put forward by policy makers during the COVID-19 pandemic, often with two main justifications: first, it would help lessen the load on hospitals, thereby increasing chances of recovery for infected individuals; second, it would help minimize the number of victims until a vaccine is found. Neither of these two incentives is present in the current model, since the recovery rate γ is independent of the number of

infections, and the possibility of a vaccine is not modeled. Nevertheless, we can observe a significant flattening of the curve across all of the examples presented, relative to the infection curve in a world with no participation margin. Individuals' expected cost of infection provides them with incentives to stay home, which ends up diminishing the death toll of the epidemic even in a world where herd immunity is the only way out. It is interesting to note that due to the complementarities between the participation decisions of susceptible agents, the endogenous "lock-down" can happen from the very beginning of the outbreak (e.g., in the case with $x = 0$), when the number of infected agents is still very low. This differs from Bethune and Korinek (2020), Farboodi et al. (2020), Garibaldi et al. (2020), and Toxvaerd (2020), where behaviors only start to change after the epidemic has gained more ground.

3) Participation and herd immunity. In a standard SIR model, the epidemic can only start declining once the number of susceptible individuals is low enough, which can be directly mapped to a minimum number of people having been infected and having gained immunity. This herd immunity threshold is very high for COVID-19: using the calibration presented in Section 4.3, we saw that it would require more than 71% of the population to have been infected for the outbreak to start declining. The simulations presented above show that the epidemic can start to wane much before that many people have been infected. For example, when $x = 1$, the number of active cases starts to decline after only 2.8% of the population has been infected, much before herd immunity has been achieved. It is sufficient for the *participating population* to have reached herd immunity, a sort of qualified herd immunity. As long as $\sigma_e^p = \alpha(S^p + I + R)\tau S^p/(S^p + I + R) < 1$, the measure of infected individuals decreases. Because this is a jump variable, it cannot be mapped into a corresponding measure of people having recovered, and it does not necessarily last forever (as visible in Figure 9, where some equilibrium paths go down then up again). Eventually, the entire population does gain herd immunity, and participation can get back to full participation.[22] Hence, while herd immunity is eventually reached by the population in both the standard SIR model and the SIR model with participation, the important difference is that in the latter model, the epidemic can start declining much before

[22] As I gets close to 0, all susceptible agents participate because the risk of infection is very small compared to the benefits from interacting with resistant agents. As a result, the effective reproduction number is greater than one when I is close to zero and $S > P/\sigma$. Thus, as was the case in the benchmark SIR model, only steady states where $S^* \in [0, P/\sigma]$ can be reached, which implies that there will always be at least a share $(1 - 1/\sigma)$ of the population eventually infected.

the threshold has been reached, eventually leading to a much smaller number of total infections.

4) Invariance of long-run epidemiological outcomes. One striking result is that despite the large multiplicity of equilibria, long-run epidemiological outcomes do not differ very much across those equilibrium paths. In all specifications reported in the analysis, the steady-state number of resistant agents lies between 78% and 81%. This implies that two regions with widely different infection curves may not necessarily be headed towards widely different long-run outcomes, and that infection curves by themselves may not be sufficient indicators of how well a region is doing.

5) Negative welfare outcomes. The issue of adverse selection was highlighted in section 4.1. Simulations with the SIR model confirm that it comes at a great cost. From the point of view of society, diminishing participation is beneficial: it allows the number of total infections to drop. However, in equilibrium, the drop of participation comes from susceptible agents exclusively. It would be much better both for susceptible agents and for society if infectious agents, who are the large minority, were the ones staying home. Even in the case where agents participate as much as possible, the losses due to missed interactions remain extremely large, and the gains due to fewer infections are not sufficient to offset those losses. Equilibria that simply delay the epidemic, in such a way that the infection curve is shifted to the right, are even worse. As discussed in *takeaway 4*, while they feature a delayed epidemic, those equilibrium paths do not come with significantly fewer infections, so that the diminished participation is a pure loss for society. It is worth noting that this negative welfare result largely depends on the way infections (and by proxy, deaths), are valued both privately and by society. Robustness checks show that either increasing the private cost of infection by 50%, or the social cost of infection by 45%, would make the equilibrium path with $x = 1$ reach lower welfare losses than the benchmark path.

5 To mask or not to mask

We now consider a model where agents cannot opt out from social interactions, so that $N(t) = P$, but can choose to wear a mask as a precaution. After solving for the equilibrium analytically in a SIS specification, I set up, calibrate and solve numerically for the SIR equilibrium outcomes. In a last part, the participation decision is reintroduced

in order to shed light on the interaction between the two reaction margins.

5.1 SIS model

The HJB equation for an infected agent is similar to (4), derived in the participation model, where we saw that infected agents have no incentives to restrict their participation. Likewise, in the SIR model with mask-wearing, infected agents prefer not to wear masks. As a result, the only relevant transmission rates are τ^{nn} and τ^{mn}. The second superscript being superfluous, notations for those two parameters are simplified to τ^n and τ^m. The HJB equation for a susceptible agent is

$$rV_S = \alpha(P)\tilde{y} + \max\left\{-k + \tau^m \frac{I}{P}(V_I - V_S), \tau^n \frac{I}{P}(V_I - V_S)\right\} + \dot{V}_S. \tag{14}$$

As in the benchmark model, the agent matches with other agents with a Poisson arrival rate $\alpha(P)$, in which case she can expect to earn $\tilde{y}$ utils. The second term of the maximization represents the expected cost of social engagement when not wearing a mask. In this case, neither she nor the infected trade partners she may meet wear masks, and the probability of infection when in contact with an infectious agent is τ^n. The first term in brackets represents the cost of social engagement when wearing a mask. The agent would suffers a flow disutility k, but would contract the disease with a smaller probability, $\tau^m < \tau^n$, in case of contact with an infectious agent.

Again, we can combine the two HJB equations to obtain one differential equation in $\omega \equiv V_S - V_I$,

$$\dot{\omega}(t) = (r+\gamma)\omega(t) - \psi + \min\left\{k + \alpha(P)\frac{I(t)}{P}\tau^m\omega(t), \alpha(P)\frac{I(t)}{P}\tau^n\omega(t)\right\}. \tag{15}$$

The measure of infected agents evolves according to the law of motion

$$\dot{I} = S^m\alpha(P)\frac{I}{P}\tau^m + (S - S^m)\alpha(P)\frac{I}{P}\tau^n - \gamma I, \tag{16}$$

where S^m denotes the measure of susceptible agents who wear a mask. The first term corresponds to the inflow of previously-susceptible agents who wore masks but still contracted the virus. The second term corresponds to susceptible agents who did not wear masks and got infected. The third term corresponds to previously-infected agents who recovered (and become susceptible again).

To close the model, we now need to solve for the measure of susceptible agents wearing a mask. To decide whether to wear a mask, a susceptible agent weighs the cost, k, against the benefit, $\alpha(P)\frac{I}{P}\omega(\tau^n - \tau^m)$. We directly get that in aggregate, the measure of susceptible agents who wear masks, S^m, is given by

$$S^m \begin{cases} = S & < \\ \in [0, S] \text{ if } k & = \quad \alpha(P)(\tau^n - \tau^m)I\omega/P \\ = 0. & > \end{cases} \tag{17}$$

Equilibrium definition An equilibrium consists in a list of time paths for the two state variables $\{S(t), I(t)\}$ and the two control variables $\{S^m(t), \omega(t)\}$, such that (15) and (16) are satisfied, where S^m is given by (17), $S(t) = P - I(t)$, and $I(0)$ is given.

Steady-states As in the benchmark and participation-based SIS models, an endemic and a virus-free steady state coexist. The virus-free steady state is such that $I^* = 0$, $S^* = P$, $S^m = 0$ and $\omega = \psi/(r+\gamma)$. Absent any risk of infection, there is no incentive for susceptible agents to be wearing masks.

In the endemic steady state, the distribution of infectious and susceptible agents as well as the prevalence of masks depend on the magnitude of k relative to $(\tau^n - \tau^m)$. Let $\underline{k} \equiv [\alpha(P)\tau^m - \gamma](\tau^n - \tau^m)\psi/\left\{\tau^m(r+\gamma) + \tau^n[\alpha(P)\tau^m - \gamma]\right\}$ and $\bar{k} \equiv [\alpha(P)\tau^n - \gamma](\tau^n - \tau^m)\psi/\left\{\tau^n(r+\gamma) + \tau^n[\alpha(P)\tau^n - \gamma]\right\}$. There are three cases.

When $k \geq \bar{k}$, $I^* = P(1 - 1/\sigma) \equiv I^{*n}$, $S^* = P/\sigma$, $S^{m^*} = 0$ and $w^* = \psi/[r + \alpha(P)\tau^n]$. Intuitively, if the cost of wearing a mask is high enough compared to the benefit, even the endemic steady state features no mask-wearing. In that case, it is identical to the endemic steady state from the benchmark SIS model.

When $k < \underline{k}$, the endemic steady state is such that $I^* = P\{1 - \gamma/[\alpha(P)\tau^m]\} \equiv I^{*m} < I^{*n}$, $S^* = P\gamma/[\alpha(P)\tau^m]$, $S^{m*} = S$, and $\omega^* = \psi/[r + \alpha(P)\tau^m]$. For a low enough k, all susceptible agents wear a mask in the endemic steady state. Recall that in an endemic steady state, the effective reproduction number must be equal to one. For the two steady states just described, it is given by $\sigma_e^m = \alpha(P)\tau^j S/(\gamma P)$, with j respectively equal to n and m. Given $\tau^m < \tau^n$, masks wearing allows the endemic steady state to occur with a lower proportion of infectious agents in the population.

Finally, when $k \in (\underline{k}, \bar{k})$, the endemic steady state is given by $I^* = P(r+\gamma)k/[\alpha(P)(\tau^n - \tau^m)\psi - \tau^n k] \in (I^{*m}, I^{*n})$, $\omega^* = \psi/(r+\gamma) - k\tau^n/[(r+\gamma)(\tau^n - \tau^m)]$. We can show that I^* is

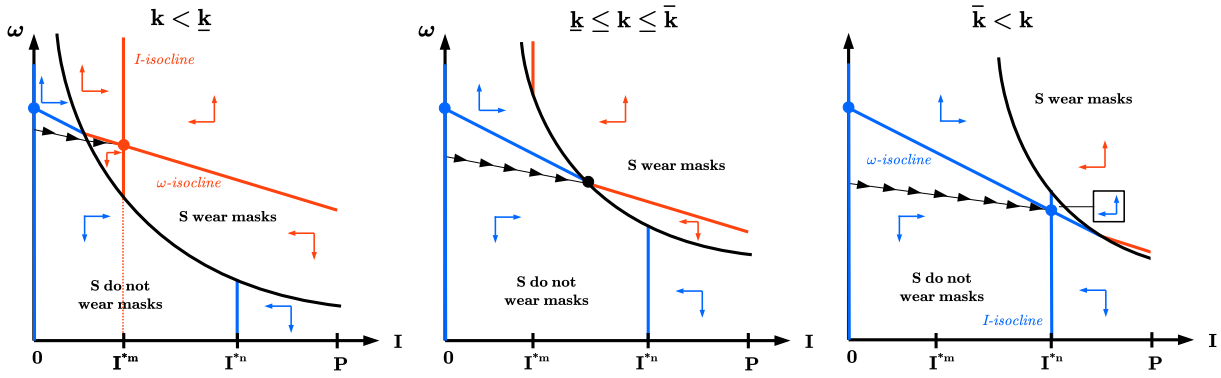

Figure 10: Dynamics of the SIS model with mask-wearing decision, as a function of the flow cost of wearing a mask, k.

increasing in k, so that $S^* = P - I^*$ is decreasing in k, and $S^{m^*} = (\alpha\tau^n S^* - \gamma)/(\tau^n - \tau^m)$ is decreasing in k as well. The costlier masks, the smaller the prevalence of masks at steady state, so that a higher proportion of the population must be infected.

Dynamics System dynamics can be analyzed with the help of the phase diagrams displayed in Figure 10. The left panel represents the case when $k < \underline{k}$, the middle panel the case when $k \in [\underline{k}, \bar{k}]$, and the right panel the case where $k > \bar{k}$. Each diagram is split into two zones by the mask-wearing indifference curve, plotted in black and given by $\omega = Pk/[(\tau^n - \tau^m)I]$. Above this curve, all susceptible agents wear masks. Below, none do. Along the curve, when I goes to zero, ω goes to infinity. As a result, at the limit, agents never wear masks because the cost is too high relative to the infrequent benefit. Then, the reasoning explained in Section 3 applies: as long as the basic reproduction number is greater than one, the measure of infected agents would always increase when I is close to zero, making the virus-free steady state unstable. This result is different than what had been obtained in the participation model. With the participation margin, it was possible for susceptible agents to coordinate on staying home even when I got close to zero because of the complementarities between their individual decisions. Here, there is no such complementarities. A corollary is that agents do not mask at the very onset of the epidemic either, when I is still low. Instead, dynamics can be described as follows.

When k is high, starting from a low $I(0)$, the unique equilibrium path is such that $\omega(0)$ jumps onto the saddle path, represented by the black line with arrows, and the

system remains on that path, with I growing and S going down, until it converges to the endemic steady state represented by the blue dot. There is no mask-wearing at any point in time. Thus, equilibrium dynamics are exactly identical to the equilibrium dynamics from the standard SIS model. When k is in the middle range, agents initially do not wear masks and the epidemic develops identically to the previous case, up until reaching the indifference curve. At this point, a measure S^{m*} of the susceptible agents start wearing a mask and the system is at steady state, represented by the black dot. When k is low, there is still no mask wearing at the onset of the epidemic, which initially develops no differently that when k is higher. However, agents eventually become indifferent between wearing a mask or not and at this point, the system reaches the saddle path that leads to the mask-wearing steady state (represented by the red dot). On this saddle path, all agents wear masks. In all three cases, dynamics are monotone in I, S and S^m.

5.2 SIR model

We now switch to the SIR specification, where infectious agents who recover become permanently resistant. Very few modifications are required. The HJB for susceptible agents remains identical to (14). The HJB for infectious agents also remains similar to that from the SIS model, but the term $(V_S - V_I)$ becomes $(V_R - V_I)$. Finally, the HJB for recovered agents is identical to the one derived for the SIR participation model, (10), with $N = P$ since resistant agents have no incentives to wear masks. We can then obtain a differential equation in ω,

$$\dot{\omega} = r\omega - \min\left\{k + \alpha(P)\tau^m \frac{I}{P}\omega, \alpha(P)\tau^n \frac{I}{P}\omega\right\} - r\frac{\psi}{r+\gamma}. \tag{18}$$

The law of motion is for I is given by (16), the low of motion for R by (12), and the aggregate measure of mask wearers by (17).

Equilibrium definition An equilibrium consists in a list of time paths for the three states variables $\{I(t), R(t), S(t)\}$ and the two control variables $\{S^m(t), \omega(t)\}$ such that (18), (16) and (12) are satisfied, where S^m is given by (17), $S(t) = P - I(t) - R(t)$, $R(0)$ and $I(0)$ are given.

Steady-states Again, no different than standard SIR models, steady state requires the economy to be virus-free, $I = 0$, and any $S^* \in [0, P]$ and $R^* = P - S^*$ constitute a steady state of the system. We showed earlier that there is no mask-wearing when I is low, so $S^{m*} = 0$.

We now calibrate the model in order to study its dynamics.

5.3 Calibration

Most parameters were already present in the model with participation, calibrated in Section 4.3, and are kept identical. There are two additional parameters to calibrate: the transmissibility of the virus in a contact where the susceptible agent wears a mask while the infectious agent does not, τ^{mn}, and the cost of wearing a mask, k.

There is no general consensus regarding the efficacy of masks, which significantly depends on the type of mask as well as the fit of the mask to the wearer. While it is generally understood that masks are most efficient when they are worn by the source of the virus (the infectious agents), several studies do suggest that masks also confer a benefit to the wearer. For example, Li et al. (2020) claim that masks reduce the risk of infection for the wearer between 40% to 70%. We pick an efficacy of 50%, so that $\tau^{mn} = 0.5\tau^{nn} = 0.025$.

The cost of mask-wearing is calibrated using a heuristic approach driven by the idea that mask-wearing is costly because of the physical inconvenience it imposes when going out and engaging in social interactions—it may be harder to breathe or to communicate, for example. We then assume that k is proportional to the utility received from meetings. More specifically, $k(N) = x\alpha N\tilde{y}/100$, i.e., agents suffer a disutility cost equal to $x\%$ of their utility from having to wear a mask. We then pick $x \in \{1, 5, 10\}$, which corresponds to daily costs of around \$1.23, \$6.16 and \$12.33 when participation is full (averaging ten meetings per person per day). Those three values will be referred to as low, medium, and high k thereafter.

5.4 Results and discussion

I now describe some key equilibrium results. More detailed results are available in Table 3 in Appendix A. Contrary to the model with participation, there are no complementarities between the mask-wearing decisions of agents, and the equilibrium is unique.

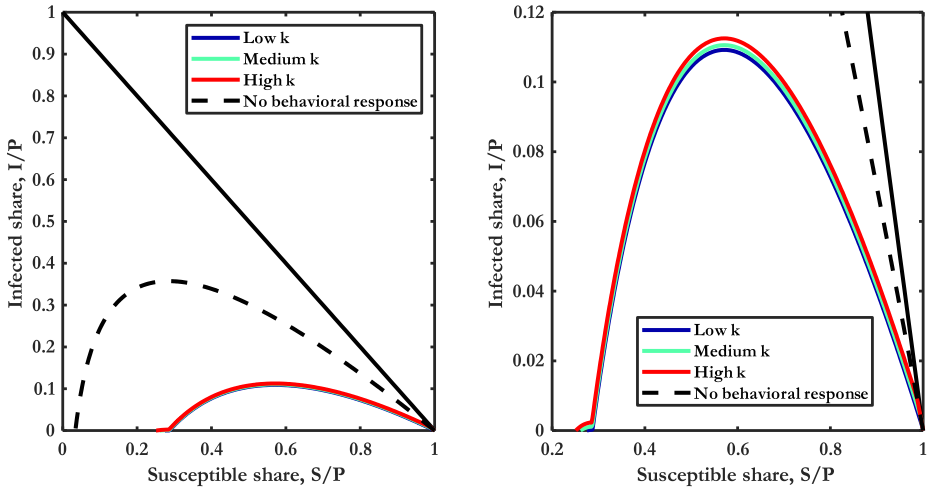

Figure 11: Equilibrium paths of the SIR model with mask-wearing decision, represented in a phase plane. Left panel - Paths for the three levels of the mask-wearing cost: low k in blue, medium k in green, and high k in red. The path with no mask-wearing decision, in dashed black, is provided for comparison. Right panel - Magnified view of the bottom-right of the graph displayed in the left panel.

The relationship between I and S is plotted in the phase planes displayed in Figure 11, for the three levels of k considered. The left panel shows the full graph, so that the equilibrium paths can easily be compared with the equilibrium path obtained absent the mask-wearing decision. We first notice that the paths do not differ much across the three levels of k, and are difficult to distinguish one from another. While the shape of the curve is roughly similar to that followed by the benchmark curve—it first goes up then goes down with a slightly flatter slope, it is much flatter, from the very beginning. As a result, the number of infected individuals is always lower (never breaking past 12%), the peak of the epidemic is reached for a lower S, and in the long-run, a considerably smaller number of individuals are infected.

One interesting feature to notice is that as I gets close to 0, the curve becomes significantly flatter until it reaches the steady state. This certainly corresponds to a shift back to no mask-wearing, which can be confirmed with the time plots. The right panel is a magnified view of the left panel, zooming in on the bottom-right corner. It is now easier to distinguish the three equilibrium paths. As expected, the lower the cost, the lower the peak of infections, and the lower the number of cumulative infections in the long run. More precisely, once at steady state, 72.51% of the population has been infected when k is low, 73.87% when k is in the middle, and 74.90% when k is high.

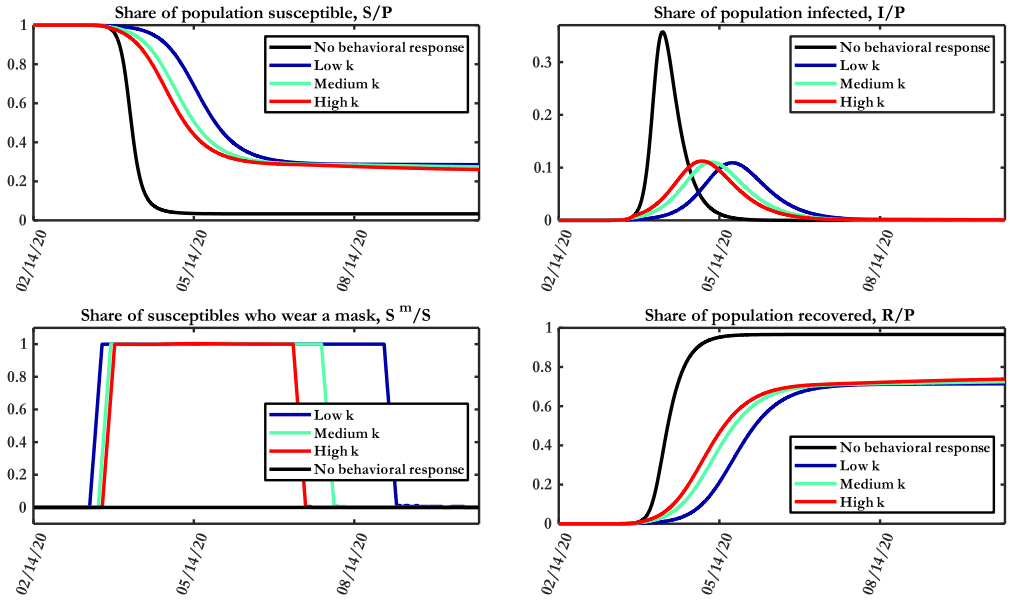

Figure 12: Time paths of epidemiological measures for the SIR model with mask-wearing, for the three levels of the mask-wearing cost: low k in blue, medium k in green, and high k in red. Note that 3 months lapse between each tick and label on the x-axis.

Recall that absent any behavioral response, 96.63% of the population gets infected, while in the equilibrium with participation, between 78.15% and 79.30% were.

We can now look at the dynamics of the epidemic on the four panels of Figure 12. From the top-right panel, we can see that not only is the infection curve flattened when we allow for mask-wearing, but it is also delayed. The lower the cost of wearing a mask, the greater the delay. In the benchmark scenario, the epidemic peaks around mid-April. With high k, it peaks around the beginning of May, while with low k it peaks around the end of May.

All of those differences can be explained by the time paths for the share of susceptible individuals who wear a mask, displayed in the bottom left panel. Mask-wearing does not occur in equilibrium either at the very beginning or the very end of the epidemic, when I is low and the chance of infection is too low to warrant bearing the cost of mask-wearing. As I goes up, it eventually becomes rational to wear a mask and as expected, this occurs first when k is low. At this point, susceptible individuals sharply go from no mask-wearing to full mask-wearing, in the matter of less than two weeks. This occurs in the second half of March 2020. All susceptible individuals then

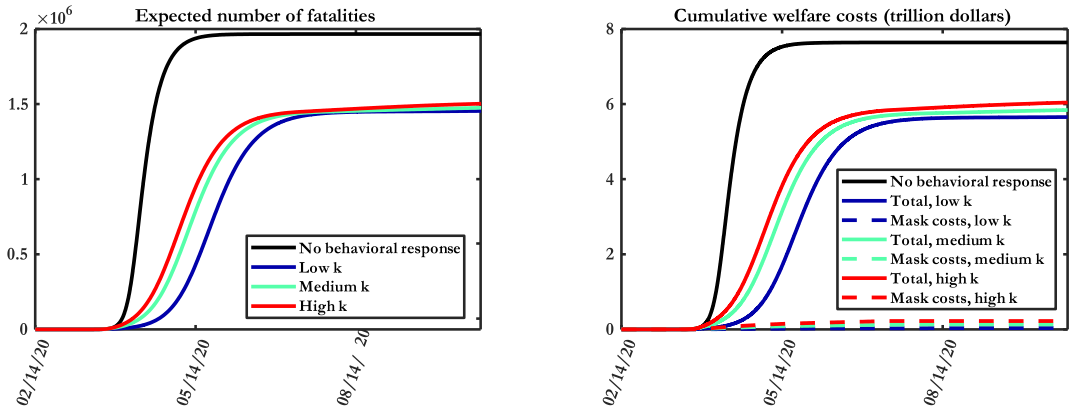

Figure 13: Time paths of expected number of fatalities and cumulative welfare losses for the SIR model with mask-wearing, for the three levels of the mask-wearing cost: low k in blue, medium k in green, and high k in red. Total welfare losses correspond to the discounted cumulative sum of sickness costs borne by infected agents and the mask-wearing costs borne by susceptible agents. Note that 3 months lapse between each tick and label.

keep wearing their mask until I is back to being low enough, again in a relatively quick fashion. This occurs at the end of July for the high k, during the first half of August for the medium k, and at the beginning of September for the low k.

We can now turn to the long-term impact for society and look at the expected number of fatalities as well as the welfare cost of the epidemic in the long-run. Those are represented in Figure 13. Because the number of total infections is relatively similar across all calibrations for k, the expected number of fatalities is relatively constant as well. We can see on the left panel that this number is around 1.5 million. Not only is it considerably smaller than the expected number of fatalities in the no-response scenario (approximately 2 million), it is also smaller than the expected number of fatalities in the model with the participation margin (approximately 1.6 million).

The right panel displays the discounted cumulative welfare losses. Those include the sickness costs borne by infected agents, and the mask-wearing costs borne by susceptible agents. While in the participation decision model, the welfare gains obtained thanks to the drop in infections and fatalities due to susceptible agents staying home were more than offset by the cost of foregone economic and social activity, the conclusion is much rosier here. Compared to the benchmark model, the economy goes from a cumulative loss of 7.6 trillion dollars in the benchmark case to between 5.7 and 6.1 trillion dollars when masks are introduced. Notably, the lower the cost of masks,

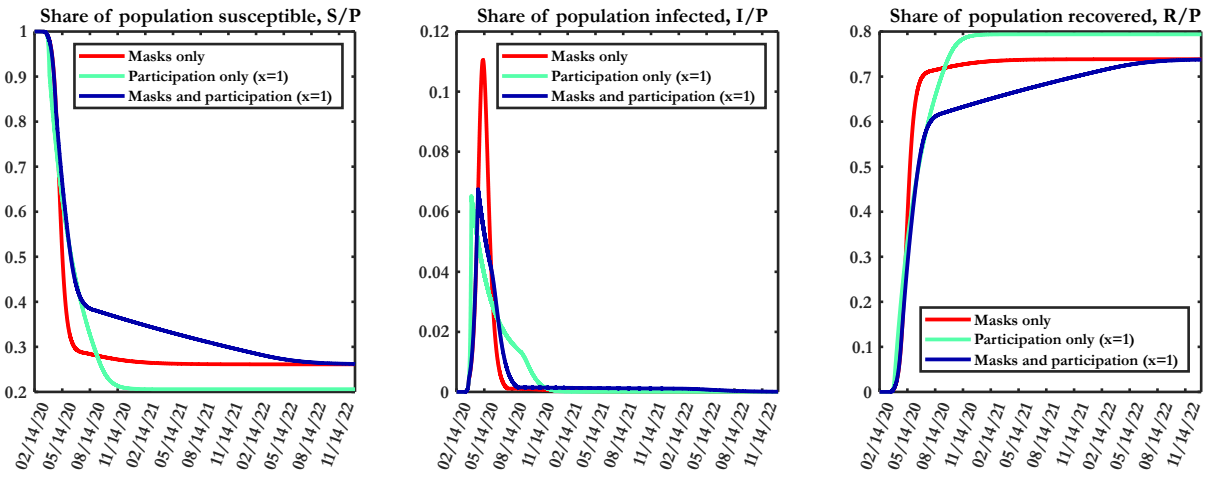

Figure 14: Time paths of epidemiological measures for the SIR model with mask-wearing and participation, assuming a medium cost for mask-wearing and coordination on participating (x=1). Note that 3 months lapse between each tick and label on the x-axis.

the smaller the welfare loss, even though the losses due to the costs of wearing masks are negligible compared to the losses due to the costs of sickness (between 0.62% to 3.65%).

5.5 The relation between masks and participation

The possibility for agents to opt out of the market and stay home is now reintroduced in addition to the option to wear a mask, which revives the multiplicity of equilibria. As an illustration, I focus on equilibria where susceptible agents always coordinate on participating when multiple Nash equilibria exist ($x = 1$). The cost of masks is calibrated to the intermediate value.

The time paths for epidemiological outcomes are represented in Figure 14. The middle panel displays the infection curve. The equilibrium path with both the mask and the participation margins is displayed in blue, and can be compared with the paths obtained when considering only masks (in red) or only participation (in green). We first notice that the infection curve is in between the two other curves. At its peak, active cases represent around 6.5% of the population, which is significantly lower than for the mask-only case, but a bit higher than the participation-only case. From the right panel, we can see that in the long run, the cumulative measure of agents that has

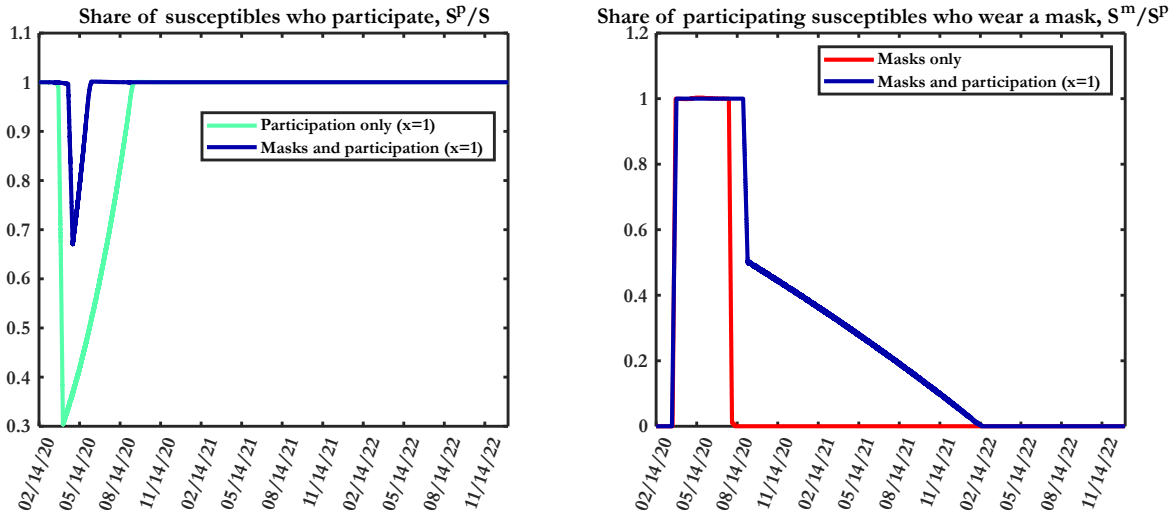

Figure 15: Time paths of participation and mask-wearing responses for the SIR model with mask-wearing and participation, assuming a medium cost for mask-wearing and coordination on participating (x=1). Note that 3 months lapse between each tick and label on the x-axis.

been infected is very similar when considering the two margins as when considering masks only, even though in the short run, the curve is steeper when only considering masks.

Figure 15 provides some additional insights. The left panel represents the share of susceptible agents who participate. As expected, adding the mask margin dampens the reaction along the participation margin. Susceptible agents start to reduce their participation approximately two to three weeks later, and go back to full participation more than a month earlier. Additionally, the level of participation does not drop as much. At its lowest, it reaches around 67% of full participation, compared to 30% without the mask margin.

The right panel represents the share of participating susceptible agents who wear a mask, contrasting the equilibrium path when only the mask margin is active to that when the two margins are active. While the two paths originally overlap, they start to diverge just before August. At this point, susceptible agents entirely stop wearing masks when this is the only margin considered. On the other hand, when participation is also considered, mask-wearing continues at 100% for a few additional weeks, then declines steadily, but slowly, only reaching 0% around February of 2022. A key takeaway is that, maybe surprisingly, adding the participation margin reinforces the

mask response. The intuition is that the mask margin is only effective for participating agents. If fewer agents participate, for a given share of mask wearers, infections are not as frequent, and immunity builds more slowly. In response, it becomes rational to keep wearing masks for a long time, as the number of active cases slowly vanishes.

6 Conclusion

This paper contrasted the impact of two response margins usually absent from epidemiological models of virus transmission—participation and mask-wearing—on infection dynamics and long-run outcomes in an economy where individuals gain utility from person-to-person social contacts. Rational agents and their decision making are embedded into a typical model of disease transmission in a micro-founded fashion, making use of search-and-matching methods to model interactions at a granular level.

When considering whether to participate in the matching process, individuals must weigh the risk of infection in each interaction with the benefits from that interaction. The more susceptible peers participate, the less risky the pool of participants, the smaller the risk of infection in a given interaction, and thus the higher the incentives for other susceptible agents to participate. These complementarities generate a continuum of equilibria and a rich set of dynamics, e.g. there can be multiple waves of infections, even absent any policy intervention. Another implication of the complementarities is that equilibria feature adverse selection, whereby susceptible agents leave the market while infectious agents stay, which comes at a very large welfare cost in the calibrated version of the model.

Predictions are markedly different when allowing agents to take precautions rather than to withdraw from the market. The equilibrium is unique, and infection dynamics are single-peaked. Mask-wearing being very efficient relative to its cost, both the human toll and welfare outcomes are considerably better under this specification than under the participation specification.

When both margins are considered simultaneously, a notable interplay arises. On one hand, the option to wear a mask reduces incentives to stay home, and dampens the drop in participation compared to a model with only a participation margin. On the other hand, the option to stay home makes the adoption of masks last longer, as a reduction in participation means that fewer people get exposed to the virus, and herd immunity builds more slowly.

There are many interesting avenues for further research. One could explore the impact of different policies on behaviors, in particular when taking the two response margins into consideration. For example, could a mask mandate lead to undesirable outcomes by increasing the participation of agents? Another path would be to expand the model to multiple regions, e.g., with different densities, so as to investigate cross-regional contagion dynamics and policies such as border closures.

References

Bethune, Z.A. and Korinek, A., 2020. Covid-19 infection externalities: Trading off lives vs. livelihoods (No. w27009). National Bureau of Economic Research.

Bi, Q., Wu, Y., Mei, S., Ye, C., Zou, X., Zhang, Z., Liu, X., Wei, L., Truelove, S.A., Zhang, T. and Gao, W., 2020. Epidemiology and transmission of COVID-19 in 391 cases and 1286 of their close contacts in Shenzhen, China: a retrospective cohort study. The Lancet Infectious Diseases.

Chu, D.K., Akl, E.A., Duda, S., Solo, K., Yaacoub, S., Schünemann, H.J., El-harakeh, A., Bognanni, A., Lotfi, T., Loeb, M. and Hajizadeh, A., 2020. Physical distancing, face masks, and eye protection to prevent person-to-person transmission of SARS-CoV-2 and COVID-19: a systematic review and meta-analysis. The Lancet.

DeStefano, F., Haber, M., Currivan, D., Farris, T., Burrus, B., Stone-Wiggins, B., McCalla, A., Guled, H., Shih, H., Edelson, P. and Wetterhall, S., 2011. Factors associated with social contacts in four communities during the 2007–2008 influenza season. Epidemiology & Infection, 139(8), 1181-1190.

Eichenbaum, M.S., Rebelo, S. and Trabandt, M., 2020. The macroeconomics of epidemics (No. w26882). National Bureau of Economic Research.

Farboodi, M., Jarosch, G. and Shimer, R., 2020. Internal and external effects of social distancing in a pandemic (No. w27059). National Bureau of Economic Research.

Feehan, D.M. and Cobb, C., 2019. Using an Online Sample to Estimate the Size of an Offline Population. Demography, 56(6), 2377-2392.

Garibaldi, P., Moen, E.R. and Pissarides, C.A., 2020. Modelling contacts and transitions in the SIR epidemics model. Covid Economics Vetted and Real-Time Papers, CEPR.

Hall, R.E., Jones, C.I. and Klenow, P.J., 2020. Trading off consumption and covid-19 deaths (No. w27340). National Bureau of Economic Research.

Keppo, J., Quercioli, E., Kudlyak, M., Smith, L. and Wilson, A., 2020. The behavioral SIR model, with applications to the Swine Flu and COVID-19 pandemics. Working paper.

Kermack, W.O. and McKendrick, A.G., 1927. A contribution to the mathematical theory of epidemics. Proceedings of the royal society of london.

Series A, Containing papers of a mathematical and physical character, 115(772), pp.700-721.

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

Krueger, D., Uhlig, H. and Xie, T., 2020. Macroeconomic dynamics and reallocation in an epidemic (No. w27047). National Bureau of Economic Research.

Leung, K., Jit, M., Lau, E.H. and Wu, J.T., 2017. Social contact patterns relevant to the spread of respiratory infectious diseases in Hong Kong. Scientific reports, 7(1), 1-12.

Li, T., Liu, Y., Li, M., Qian, X. and Dai, S.Y., 2020. Mask or no mask for COVID-19: A public health and market study. PloS one, 15(8), p.e0237691.

Liu, Y., Gayle, A.A., Wilder-Smith, A. and Rocklöv, J., 2020. The reproductive number of COVID-19 is higher compared to SARS coronavirus. Journal of travel medicine.

Luo, L., Liu, D., Liao, X., Wu, X., Jing, Q., Zheng, J., Liu, F., Yang, S., Bi, H., Li, Z. and Liu, J., 2020. Contact Settings and Risk for Transmission in 3410 Close Contacts of Patients With COVID-19 in Guangzhou, China: A Prospective Cohort Study. Annals of internal medicine.

McAdams, D., 2020a. Nash SIR: An economic-epidemiological model of strategic behavior during a viral epidemic. Covid Economics, 16, 115-134.

McAdams, D., 2020b. Economic Epidemiology in the wake of Covid-19. Covid Economics, 48, 1-45.

Mossong, J., Hens, N., Jit, M., Beutels, P., Auranen, K., Mikolajczyk, R., Massari, M., Salmaso, S., Tomba, G.S., Wallinga, J. and Heijne, J., 2008. Social contacts and mixing patterns relevant to the spread of infectious diseases. PLoS Med, 5(3), 74.

Oran, D.P. and Topol, E.J., 2020. Prevalence of Asymptomatic SARS-CoV-2 Infection: A Narrative Review. Annals of Internal Medicine.

Salanié, F. and Treich, N., 2020. Public and private incentives for self-protection. The Geneva Risk and Insurance Review, 1-10.

Toxvaerd, F.M.O., 2020. Equilibrium social distancing. Working paper.

Worobey, M., Pekar, J., Larsen, B.B., Nelson, M.I., Hill, V., Joy, J.B., Rambaut, A., Suchard, M.A., Wertheim, J.O. and Lemey, P., 2020. The emergence of SARS-CoV-2 in Europe and the US. Working paper.

Appendix A - Results tables

	R^*/P (%)	Deaths (M)	Max I/P (%)	Peak date	Min S^p/S (%)	End date
			Benchmark			
	96.63	1.9665	35.73	12-Apr-20	-	11-Sep-20
x			**Probability-based rule**			
0	78.15	1.5903	1.66	21-Feb-25	0.00	23-Jan-27
0.1	78.15	1.5904	1.65	8-Dec-24	9.83	13-Nov-26
0.2	78.15	1.5904	1.64	18-Aug-24	19.24	24-Jul-26
0.3	78.15	1.5903	1.57	7-Jul-23	29.34	19-Jun-25
0.4	78.15	1.5903	3.70	27-Dec-20	38.77	18-Nov-22
0.5	78.14	1.5903	5.44	31-Jul-20	34.87	2-Jun-22
0.6	78.14	1.5903	5.70	7-Jun-20	32.91	30-Mar-22
0.7	78.14	1.5903	5.88	9-May-20	31.77	21-Feb-22
0.8	78.15	1.5903	5.97	22-Apr-20	31.37	29-Jan-22
0.9	78.37	1.5948	6.14	10-Apr-20	30.86	28-Dec-21
1	79.39	1.6157	6.53	2-Apr-20	30.40	22-Oct-21
T			**Fatigue rule (months)**			
1	79.39	1.6157	6.53	16-Apr-20	0	5-Nov-21
5	79.39	1.6157	6.50	11-Aug-20	0	2-Mar-22
10	79.39	1.6157	6.51	4-Jan-21	0	26-Jul-22
15	79.39	1.6157	6.52	30-May-21	0	19-Dec-22
20	79.4	1.6158	6.50	24-Oct-21	0	15-May-23
25	79.39	1.6156	6.50	20-Mar-22	0	9-Oct-23
29	79.39	1.6158	6.44	16-Jul-22	0	3-Feb-24
30	79.40	1.6158	1.73	16-Feb-22	0	22-Oct-23
31	79.04	1.6085	1.72	22-Feb-22	0	26-Nov-23
32	78.49	1.5974	1.68	11-Mar-22	0	17-Jan-24
33	78.15	1.5904	1.66	19-Mar-22	0	20-Feb-24
34	78.15	1.5904	1.66	19-Mar-22	0	20-Feb-24
35	78.15	1.5904	1.66	19-Mar-22	0	20-Feb-24
36	78.15	1.5904	1.66	19-Mar-22	0	20-Feb-24
A			**Active-cases rule (% of P)**			
0.1	78.33	1.5942	1.67	-	17.10	30-Jun-22
0.2	79.39	1.6156	2.03	-	13.97	2-Jan-22
0.3	79.39	1.6157	3.04	-	9.22	3-Nov-21
0.4	79.39	1.6157	4.05	-	6.88	12-Oct-21
0.5	79.39	1.6156	5.06	-	5.53	3-Oct-21
0.6	79.39	1.6157	6.08	-	4.60	30-Sep-21

Table 2: Full set of results for the calibrated SIR model with participation. Peak date corresponds to the date when the highest I/P is reached. End date corresponds to the date when $I < 1$.

	R^*/P (%)	Deaths (M)	Max I/P (%)	Peak date	Max S^m/S (%)	End date
Low k	72.51	1.4757	10.92	21-May	100	9-Feb-27
Mid k	73.87	1.5033	11.06	9-May	100	5-Oct-23
High k	74.90	1.5243	11.25	4-May	100	15-Nov-22

Table 3: Full set of results for the calibrated SIR model with mask-wearing. Peak date corresponds to the date when the highest I/P is reached. End date corresponds to the date when $I < 1$.

Appendix B - Alternative calibrations

Farboodi et al. (2020) and Bethune and Korinek (2020) Farboodi et al. (2020), Bethune and Korinek (2020), and the present paper use different approaches to calibrate the timeline of the model as well as $\alpha\tau$.

To match the timeline of the model to the real world, Farboodi et al. (2020) set time $t = 0$ to March 13th, 2020. This corresponds to the date when the authors first observe social distancing in the US, from which they infer that it corresponds to the date when the population became aware of the virus. They then make use of the fact that by then, 51 fatalities had been reported, in order to derive the implied distribution of S, I and R at time 0. In comparison, the present paper's timeline would imply 51 fatalities by March 10th absent any reaction along the participation margin. Therefore, under the assumption that the population did not react until March 13th, the two calibrations are very close.

However, because being aware of the virus before reacting through social distancing turns out to be consistent with equilibrium behavior in my model, the approach followed by Farboodi et al. (2020) may not be the most appropriate. Additionally, there is some evidence supporting that the US population may have gained awareness of the virus earlier than March 13th. On January 20th, the Center for Disease Control and Prevention (CDC) announced that three airports would start screening for COVID-19. The next day, it confirmed the first case in Washington state. On January 31st, the World Health Organization (WHO) declared a global public health emergency, and the US declared a public health emergency on February 3rd. Correspondingly, the Google Trend tools show that searches for the term "coronavirus" experienced a first significant uptick during the last week of January.

As for Bethune et al. (2020), they calibrate their timeline by fitting time 0 to mid-May, assuming that by that time, 0.3% of the population was infected. This is

considerably different from the calibrations in this paper and in Farboodi et al. For comparison, with my calibration and when agents coordinate on participating ($x = 1$), 0.3% of the population would be infected by March 24, implying a discrepancy of more than one month at the minimum.

Table 4 shows the results obtained when using the timing from Farboodi et al. (2020), denoted FSJ, compared to the timing used in this paper, denoted L. By construction, outcomes with the benchmark model are exactly identical. Surprisingly, outcomes for the participation model with $x = 1$ and for the mask model are also identical (up to numerical errors). This is because in those two cases, there is actually no reaction from individuals, on either margin, before March 13th. Thus, even if my calibration allows for agents to react, while Farboodi et al. (2020) do not, the equilibrium paths remain identical. Results are a bit different for the case with $x = 0$, where agents react from the very beginning (February 15th), if allowed to do so. Then, the timeline of the epidemic differs noticeably. In particular, it is much shorter in the Farboodi et al. setup, which is intuitive since a higher number of people is infected earlier on.

	Benchmark		**Participation**				**Masks**	
			x = 1		**x = 0**		**Mid k**	
	L	**FJS**	**L**	**FJS**	**L**	**FJS**	**L**	**FJS**
R^*/P **(%)**	96.63	96.63	79.39	79.39	78.15	78.15	73.87	73.87
Deaths (M)	1.9665	1.9665	1.6157	1.6157	1.5903	1.5904	1.5033	1.5033
Max I/P **(%)**	35.73	35.73	6.53	6.52	1.66	1.66	11.06	11.05
Peak date	12-Apr-20	12-Apr-20	2-Apr-20	2-Apr-20	21-Feb-25	22-Sep-21	9-May-20	9-May-20
Min S^p/S **(%)**	-	-	30.40	30.40	0	0	100	100
End date	11-Sep-20	11-Sep-20	22-Oct-21	22-Oct-21	23-Jan-27	28-Aug-23	5-Oct-23	5-Oct-23

Table 4: Robustness of results to alternative timeline, used in Farboodi et al. (2020). Peak date corresponds to the date when the highest I/P is reached. End date corresponds to the date when $I < 1$.

Second, the present paper is, to the best of my knowledge, the first paper in the economic literature related to COVID-19 to use a micro-founded approach to calibrating the matching rate of individuals (αP) and the transmissibility of the virus (τ). Most of the literature estimates those two parameters jointly to match the contagion dynamics observed at the beginning of the epidemic (under the assumption that at that time, behavioral responses had not yet kicked in). Farboodi et al. (2020) target the growth rate of infections by the beginning of the epidemic, which they estimate to 30%. Noting that $\dot{I}(0)/I(0) = \alpha\tau S(0) - \gamma$, we can then easily solve for $\alpha\tau = [\dot{I}(0)/I(0) + \gamma]/S(0)$.

In comparison, this paper's calibration yields a growth rate of the measure of infected agents of 35.7% by March 13th, assuming no behavioral response until then. Bethune and Korinek (2020) target a basic reproduction number, σ, of 2.5. Since $\sigma = \alpha P \tau / \gamma$, this again allows to easily solve for $\alpha\tau = \sigma\gamma/P$. The calibration used in the present paper implies a basic reproduction number of 3.5, while that used in Farboodi et al. (2020) implies a basic reproduction number of 3.1. All of those numbers are consistent with the range of reproduction number estimated by epidemiological studies, from 1.5 to 7 (see Liu et al., 2020, for example).

Baseline utility In the calibration presented in 4.3, it is assumed that social contacts are the only source of income/utility, since $\alpha P \tilde{y}$ is calibrated to match the median yearly consumption in the US. This assumption may result in overstating the cost of staying home and forgoing social activities if some amount of income/utility can be gained without requiring to go out or to come into contact with other people. We now relax this assumption.

Denote y^h the baseline flow utility from consumption enjoyed by individuals regardless or whether they are at home or engaging in social interactions, and $Y = \alpha P \tilde{y} + y^h$ the total flow utility of an individual in a world with no virus (ensuring full participation from all individuals). We can now normalize Y to 1 and vary the share of consumption coming from social engagement, $\alpha P \tilde{y}/Y$, by varying $\tilde{y}$.

The table below presents the results obtained when solving for the equilibrium path with the participation margin active, imposing $x = 1$ (susceptible agents coordinate on participating whenever possible), for $\tilde{y} \in \{0.7, 0.8, 0.9, 1\}$. The specification with $\tilde{y} = 1$ corresponds to the specification assumed in the main text.

	$\tilde{y}$			
	0.7	**0.8**	**0.9**	**1**
R^*/P (%)	78.04	78.51	78.97	79.39
Deaths (M)	1.5882	1.5978	1.6070	1.6157
Max I/P (%)	4.28	5.01	5.75	6.52
Peak date	31-Mar-20	01-Apr-20	01-Apr-20	02-Apr-20
Min S^p/S (%)	29.99	29.99	30.40	30.40
End date	25-Mar-22	21-Jan-22	02-Dec-21	22-Oct-21
Welfare cost ($T)	8.3471	8.3404	8.3325	8.3241

Table 5: Robustness of results to varying the share of total utility requiring social contacts. Peak date corresponds to the date when the highest I/P is reached. End date corresponds to the date when $I < 1$.

Varying the share of total consumption that requires interpersonal has little impact on the cumulative measure of agents infected by the virus in the long run.

In the short run, it does impact participation: a lower $\tilde{y}$ makes the benefit of going out smaller relative to the risk of infection. As a result, the lower $\tilde{y}$, the bigger the reaction of susceptible agents, the lower the peak of the infection curve, and the longer the epidemic lasts.

The impact on welfare of varying the magnitude of $\tilde{y}$ results from two different forces. A lower $\tilde{y}$ directly implies that a drop in social activities is not as costly. But we saw that a lower $\tilde{y}$ encourages a comparatively larger reduction in activity, which indirectly could increase the welfare cost. The last line of table 5 shows that the indirect effect dominates—the welfare cost is slightly lower when $\tilde{y}$ is higher.

Appendix C - Numerical algorithm for the SIR model with participation

In this section, I describe the algorithm used to solve the SIR model with participation. First, the model was discretized. The laws of motion for the measure of infected and recovered agents are given by

$$I_{t+1} = \left[1 - \gamma + \alpha\tau(P - R_t - I_t)\right] I_t \tag{19}$$

and

$$R_{t+1} = R_t + \gamma I_t. \tag{20}$$

The difference between the lifetime discounted utility of a susceptible agent and that of an infectious agent is given by

$$\omega_t = -\alpha(S_t^p + I_t + R_t) \min\left\{\tilde{y}, \tau \frac{I_t}{S_t^p + I_t + R_t}\beta\omega_{t+1}\right\} + (1-\beta)\omega^* + \beta\omega_{t+1}, \tag{21}$$

where $\beta \equiv 1/(1+r)$ and $\omega^* \equiv \psi/\left(1-\beta+\beta\gamma\right)$. Finally, the decision for a susceptible agent to participate or to stay home is governed by

$$a_{j,t} \begin{cases} = 0 & < \\ \in [0,1] \text{ if } \tilde{y} & = \tau \frac{I_t}{S_t^p + I_t + R_t} \beta \omega_{t+1}. \\ = 1 & > \end{cases} \tag{22}$$

The model is then solved forward, following these steps:
(1) Set I_0, R_0 and S_0 to their calibrated values, and pick a guess for ω_0.
(2) Making use of (21), compute $\omega_1(S_0^p)$ for $S_0^p = S_0$ and $S_0^p = 0$, so as to determine whether we are in the multiplicity region or one of the two dominance regions. If $\tilde{y} < \tau(I_0/P)\beta\omega_1(S_0^p = S_0)$, the unique Nash equilibrium is such that $S_0^p = 0$. If $\tilde{y} > \tau[I_0/(I_0+R_0)]\beta\omega_1(S_0^p = 0)$, the unique Nash equilibrium is such that $S_0^p = S_0$. Otherwise, we are in the multiplicity region, and S_0^p is determined by one of the following coordination rules, chosen ex-ante and used for the whole algorithm:

- *Always participate:* $S_0^p = S_0$
- *Never participate:* $S_0^p = 0$
- *Participation with probability x:* draw d from a uniform distribution bounded by 0 and 1. If $d \leq x$, $S_0^p = S_0$, otherwise $S_0^p = 0$
- *Fatigue rule:* for a given T, if time 0 is less than T, $S_0^p = 0$, otherwise $S_0^p = S_0$
- *Active cases rule:* for a given A, if $I_0 < A$, then $S_0^p = S_0$, otherwise $S_0^p = 0$

(3) Record the corresponding I_1, R_1, S_1 and ω_1.
(4) Iterate over steps (2) and (3) to obtain ω_2, ω_3, ... up to ω_M, where M is set to a very large number.
(5) Let ϵ be an arbitrarily small number. If $|\omega_M - \omega^*| < \epsilon$, the algorithm has converged. If $\omega_M - \omega^* > \epsilon > 0$, go back to step (1), with a lower initial guess for ω_0. Otherwise, go back to step (1), with a higher initial guess for ω_0.