The Hammer and the Dance:

Equilibrium and Optimal Policy during a Pandemic Crisis*

Tiziana Assenza, Fabrice Collard, Martial Dupaigne, Patrick Fève, Christian Hellwig†, Sumudu Kankanamge, Nicolas Werquin

August 4, 2020

Abstract

How should governments balance controlling the COVID-19 pandemic with limiting its economic costs? We develop a theoretical framework embedding a stylized economic interaction game into a dynamic S-I-R model of epidemic propagation to shed light on the trade-offs between potentially competing economic and health policy objectives during a pandemic crisis characterized by: *(i)* fast propagation and *(ii)* asymptomatic transmission. We show that the epidemiological strategies implemented to contain the pandemic — aptly described as a "hammer and dance"— are also based on sound economic principles that go beyond the usual medical arguments. Health and economic objectives are not mutually exclusive.

JEL Codes: E6, H12, I1.

Keywords: COVID-19 pandemic, optimal policy, fast propagation, asymptomatic transmission, immunization externalities, infection externalities, shadow cost of infection risk, epidemiological and economic objectives.

At the outbreak of the recent COVID-19 pandemic, the option to implement a strict lockdown to contain the spread of the disease was largely and forcefully debated. On March 23, 2020, President Trump tweeted "*WE CANNOT LET THE CURE BE WORSE THAN THE PROBLEM ITSELF*" (original capital letters) therefore suggesting the existence of a strong trade-off

*We are grateful for comments from Aditya Goenka and Martin Hellwig as well as from seminar participants at the Toulouse School of Economics, Columbia University and Covid Search and Matching workshops (Chicago University). We acknowledge funding from the French National Research Agency (ANR) under the Investments for the Future program (Investissements d'Avenir, grant ANR-17-EURE-0010).

†All authors are affiliated to the Toulouse School of Economics. Corresponding author: christian.hellwig@tse-fr.eu, Toulouse School of Economics, 1 Esplanade de l'université, 31080 Toulouse Cedex 06 (France).

between health and economic prosperity. On July 2nd, 2020 Dr. Anthony Fauci, one of the lead members of the Trump Administration's White House Coronavirus Task Force, countered that "*You don't want to balance lives against the economy. So let's get public health to help us to get the economy open as opposed to two opposing forces*" (interview with the Journal of the American Medical Association), indicating that this trade-off may not be as binding. This paper develops a theoretical framework to shed light on the trade-offs between potentially competing economic and health policy objectives during a crisis akin to the COVID-19 pandemic. As our main policy conclusion, we argue that *health and economic objectives are not mutually exclusive*.

This COVID-19 pandemic has been characterized by two important and challenging features: *(i)* it propagates fast in the population (Silverman, Hupert, and Washburne, 2020; Moghadasa et al., 2020),[1] and *(ii)* many infections and transmissions are asymptomatic and go undetected (See Moghadasa et al., 2020; Oran and Topol, 2020; Li et al., 2020; Lavezzo et al., 2020).[2] To cope with such episodes, epidemiologists have suggested policies consisting of two phases (see Ferguson, Laydon, and Gemma Nedjati-Gilani et al., 2020; Pueyo, 2020): *(i)* a strong initial confinement phase that brings new infections under control, and *(ii)* a subsequent phase of gradual deconfinement during which the epidemic slowly progresses until the population reaches a state of herd immunity or a vaccine or cure is found. During this second phase, the effective reproduction rate of the virus is stabilized near 1, so that the pandemic is neither completely suppressed nor allowed to take off again. Pueyo (2020) dubbed these two phases respectively *the Hammer* and *the Dance*. In this paper, we show that such epidemiological strategies are also sound from an economic point of view: a full economic recovery is possible only once the pandemic is under control.

In an initial stage, we keep the model deliberately simple and disregard the factors that are usually used to motivate the epidemiological approach to policy such as medical congestion, potential vaccine and cure, testing and tracing, the use of face masks, transitory immunity… Doing so allows us to obtain a stark characterization of optimal confinement and deconfinement policies, how they perform relative to the *laissez-faire* equilibrium and of the role of externalities in shaping the results. In a second stage, we bring back all these ingredients and show that

[1] For instance, Moghadasa et al. (2020) show that the share of the population that contracts the disease in an at risk population during a specified time interval (1 day) —the attack rate— is above 25%, indicating a fast propagation of the disease in the population. Silverman, Hupert, and Washburne (2020) showed that the early epidemic in the US was unlikely to have been doubling slower than every 4 days.

[2] Moghadasa et al. (2020) report that 50% of the level of the attack rate is explained by the undocumented cases. Oran and Topol (2020) report that, in the Wuhan case, asymptomatic persons accounted for approximately 40% to 45% of SARS-CoV-2 infections. Li et al. (2020) use observations of reported infection within China and estimate that 86% of all infections were undocumented before the implementation of travel restrictions. Furthermore, undocumented infections were the source for about 80% of the documented cases. Lavezzo et al. (2020) report that, in the Italian city of Vo' , 42.5% of the confirmed SARS-CoV-2 infections detected were asymptomatic.

while they have a quantitative impact on the optimal and equilibrium allocations and the policy restrictions, they leave the main message and conclusions of the paper unaffected.

Structure of the model. We embed a stylized economic interaction game into a dynamic model of epidemic propagation. In our model, agents' decisions interact in two ways: on an economic level (*economic stage game*) they determine instantaneous utilities, and on a sanitary level (*confinement stage game*) they determine the likelihood with which economic interactions result in the transmission of an infectious and potentially lethal disease. We embed these two stage games into a *dynamic S-I-R model* of epidemic propagation (See Kermack and McKendrick, 1927): Agents are initially susceptible to infection by interacting with other infected agents. Once infected, they subsequently either recover or die from the disease. Recovery confers permanent immunity.[3] We assume that agents do not know their own health state: only death is observable. While restrictive at face value, this assumption captures the reality of asymptomatic infections in COVID-19. In addition, by making all agents symmetric (except the deceased), this assumption saves on state variables and simplifies the equilibrium characterization.

The *economic stage game* is kept abstract and only assumes that the equilibrium is efficient. The model is thus general enough to encompass many textbook economic models that satisfy sufficient conditions for the Second Welfare Theorem, or similar principles developed for other forms of economic interactions – centralized and decentralized market interactions, non-market interactions in hierarchies and organizations, search and assignment markets, ... Importantly, assuming that the pre-pandemic equilibrium satisfies a generalized efficiency condition ensures that any rationale for policy interventions comes solely from the trade-off between competing economic and health care objectives. The model specification is general enough to allow for multi-dimensional actions and can therefore address sectoral differences in equilibrium and policy responses to the pandemic crisis.

The *confinement stage game* determines individual infection probabilities from individual and aggregate actions in the economic stage game. We assume that the confinement game satisfies an analogue of the static efficiency condition: there exists an extreme confinement equilibrium that minimizes global infection risks.

Our model gives rise to a trade-off between current utility and instantaneous infection risks or future mortality that can be naturally summarized by a *shadow cost of infection risks*. Static decisions and policy trade-offs (*i.e.,* which sectors to open and which ones to close) then all revolve around aligning private and social marginal rates of substitution between instantaneous utility and infection risks to this shadow cost. These marginal rates of substitution have a natural interpretation as the marginal disutility of social distancing, so our equilibrium and

[3]This assumption is largely debated in the epidemiological literature (See World Health Organization, 2020; Kissler et al., 2020) and is relaxed in an extension of the baseline model.

optimality conditions equalize the private and social marginal costs and benefits of social distancing. The dynamics of equilibrium and optimal policy in turn depend on the dynamics of this shadow cost, which admit a simple recursive characterization from the S-I-R implied population dynamics.

Main results. Our main theoretical result fully characterizes and compares the *laissez-faire* equilibrium and optimal policy path when the epidemic spreads fast, *i.e.,* close to the limit with instantaneous propagation. In this limit, the equilibrium and the planner's problem can both be solved in closed form. We also extend the main insights away from this limit.

We establish theoretically that the equilibrium dynamics and the socially optimal policy are both characterized by the aforementioned *Hammer* and *Dance* phases. Key to this result is that the shadow cost of infection risks is determined by the product of the speed of propagation of the pandemic and the marginal cost of infection. At their optimal behavior, the agent and the central planner both equalize the marginal disutility of social distancing to this shadow cost. Therefore, in the fast propagation limit, the marginal disutility is high at first —corresponding to a strong *Hammer*— until new infections are brought under control. At this point, the economy enters a gradual recovery during which new infections are kept low enough to offset the high speed of propagation — corresponding to a slow *Dance*.

However, the equilibrium differs from the planner's solution in the timing and intensity of early lockdown measures, as well as in the speed of convergence towards a long-term recovery. The dynamic equilibrium and optimal policy coincide, if and only if static efficiency conditions (alignment of private and social marginal costs of social distancing across multiple sectors or activities) are augmented by a dynamic efficiency condition that requires offsetting static and dynamic spill-overs. Static spill-overs result from the presence of static externalities in instantaneous utility and infection risks. Dynamic spill-overs correspond to a gap between the private and social shadow costs of infection risks. They arise because in equilibrium, agents only consider the static trade-off between instantaneous utilities and concurrent infection risks, taking all future dynamics as given, while the planner, instead, fully internalizes the dynamic consequences of current actions.

The offsetting spill-overs condition is generically violated as the pandemic progresses, due to the interplay of two competing dynamic externalities. First, *infection externalities* imply that agents in equilibrium fail to internalize that, by risking an infection, they expose others to higher future infection risks. Second, *immunization externalities* imply that agents voluntarily opt for strong confinement to "wait out the storm": herd immunity is a collective good, and while everyone shares the benefits, no one is eager to contribute by catching an infection.

The balance of infection and immunization externalities determines how the Hammer and Dance dynamics play out at the equilibrium and the planner's solution. In our benchmark model,

without the prospect of a vaccine and congestion in hospitals, the planner's best long-term plan is to reach herd immunity quickly, but without infecting more agents than necessary. In that case, immunization externalities dominate early on, so that, in equilibrium, agents restrict too strongly their economic interactions: everyone holds out until the worst is over, which makes the pandemic progress much more slowly and last longer, thus significantly amplifying its economic costs. By contrast, the planner optimally lets infections and mortality reach an early peak. On the other hand, infection externalities dominate during the recovery phase: once the peak of the epidemic has passed, agents grow impatient to return to their prior activities, without internalizing that, by risking an infection, they expose others to higher future infection risks. The equilibrium recovery thus starts from too low a level and occurs too fast, which generates both unnecessary economic hardship and a high number of avoidable deaths in the long-run.

At the instantaneous propagation limit, optimal policy immediately brings the population to a state that optimally trades off between economic prosperity and mortality in the long-run, and then permanently stalls the pandemic at this long-run optimum without ever reaching herd immunity: the Hammer happens instantly and the Dance lasts forever. This makes the contrast between fast convergence to the long-run optimum at the planner's solution and very gradual propagation and strong hold-out incentives at the equilibrium especially stark.

All these results are obtained from a simple, though compelling, framework that neglects various potentially important aspects of a pandemic (medical congestion, hope for a cure or a vaccine, testing and tracing…). We then assess numerically the robustness of our main findings to the introduction of these elements by considering various extensions that fall into two main categories. The first comprises fundamental static and dynamic spill-overs that modify the dynamic first-order conditions, but keep the static equilibrium and social planner's trade-offs between utility and infection risks unchanged: static externalities, medical sector congestion, transitory immunization and the potential development of a vaccine. In particular, medical congestion and the prospect of a vaccine make the initial peak of infections very costly for the social planner: the immunization externality is offset or overturned by a congestion externality or by the option value of delaying infections. The planner now favors an earlier and stronger hammer than agents at the equilibrium, and opts for a longer and economically costlier dance with the hope of saving more lives in the long run. The second set of extensions considers policy interventions that aim to directly managing the static trade-off between utility and infection risks, such as the introduction of masks or testing and contact tracing. For example, the use of face masks reduces infection rates and allows the planner and agents to ease economic restrictions.

Related literature. Our contribution differs from the rapidly growing literature that studies the impact of SARS-CoV-2 and the trade-offs between economic activity and infection risks

that a social planner faces (*e.g* Atkeson, 2020; Alvarez, Argente, and Lippi, 2020; Gonzalez-Eiras and Niepelt, 2020; Piguillem and Shi, 2020) in two respects. First, we complement these predominantly quantitative studies by providing a theoretical framework and results that shed light on the different forces at play. Notably, our framework unifies a number of contrasting results which follow from specific assumptions about static and dynamic externalities. Second, we assume that the health status is unobservable. Most of the existing literature assumes that the health status is observable to the individual if not to the planner, then, however, focus on simple policies that do not condition on this information. Sophisticated policies would instead use this information, if necessary by eliciting it through direct revelation mechanisms that exploit differential responses to exposure risks. By assuming that health status is not observable, our model directly addresses the informational challenges posed by the COVID-19 pandemic.

Bethune and Korinek (2020), Eichenbaum, Rebelo, and Trabandt (2020a), Eichenbaum, Rebelo, and Trabandt (2020b), Farboodi, Jarosch, and Shimer (2020), Jones, Philippon, and Venkateswaran (2020a) and Jones, Philippon, and Venkateswaran (2020b) all highlight differences between the competitive equilibrium and the planner's solution because infection has a higher shadow cost for the planner than for an individual agent. We complement these quantitative studies with theoretical results on static and dynamic spill-overs. Among other things, we highlight that the optimality of early lockdowns in their planner's solution is the result of medical congestion externalities present in their models, that will be easily reversed when immunization spill-overs dominate the short-run policy trade-offs. Farboodi, Jarosch, and Shimer (2020) share our conclusion that optimal confinement policies may be long lasting and carefully balanced to keep new infections under control. Our model provides a general foundation for this prescription and the suggestion by Budish (2020) that optimal policy should maximize economic welfare, subject to keeping the pandemic under control. Indeed equilibrium and optimal policy during the *Dance* phase keep the effective reproduction rate of the pandemic very close to 1.

Toxvaerd (2020), Garibaldi, Moen, and Pissarides (2020), and Krueger, Uhlig, and Xie (2020) emphasize private incentives for flattening the infection curve. Garibaldi, Moen, and Pissarides (2020) develop a "bottom-up" approach to social distancing based on insights from search theory and contrast static and dynamic externalities similar to our setup, but stop short of a full comparison of the planner and the equilibrium solutions. We instead adopt a "top-down" approach that abstracts from the specifics of a given interaction to focus on general principles while characterizing the full dynamics of infection and immunization spill-overs. Krueger, Uhlig, and Xie (2020) emphasize the role of static substitutability across sectors and sorting by susceptible agents into low risk activities. As in our setup, a planner would subsidize (or tax) some sectors according to their specific infection risk externalities. In comparison, we show that private incentives for social distancing might go too far if the planner values early immunization.

Section 1 describes our model economy and details the different stages of our economic/confinement game. Section 2 offers a full characterization of the dynamics of the equilibrium and the central planner's allocations. Section 3 proposes the various extensions to the model. A last section offers some concluding remarks.

1 Environment

Time is discrete and infinite. In each period $t = 1, 2, \ldots$ a mass Λ_t of surviving agents interacts in an *economic stage game* which determines their instantaneous utilities. But these economic interactions also expose agents to the risk of being infected with and potentially dying from an infectious disease. Hence, we juxtapose this economic stage game with a *confinement stage game*, which summarizes how the agents' decisions determine their risk of infection.

1.1 Static Interaction Games

Let $\mathcal{X} \subset \mathbb{R}^K$ be a compact, convex set of feasible economic actions or choices, with non-empty interior $int(\mathcal{X})$. Let $x \in \mathcal{X}$ denote an individual action, and $X \in \mathcal{X}$ the aggregate choice of the other agents.[4] Individual and aggregate choices jointly determine the agents' instantaneous utility in the economic stage game and their risk of infection in the confinement stage game. Allowing for K-dimensional $\mathcal{X}$ allows us to address sectoral differences in economic spill-overs or infection risks. The public policy discussion is well aware of such differences when a distinction is made between essential and non-essential sectors, or restrictions are targeted towards sectors that increase infection risk, such as mass transport, travel or entertainment events.

Economic Stage Game. Let $\mathcal{U} : \mathcal{X} \times \mathcal{X} \to \left[0, \overline{V}\right]$ denote the static flow utility of choosing $x \in \mathcal{X}$, when all other agents choose $X \in \mathcal{X}$, where $\mathcal{U}(\cdot)$ is continuous, strictly concave and twice continuously differentiable over the interior of $\mathcal{X} \times \mathcal{X}$. We assume the following about $\mathcal{U}(\cdot)$:

Assumption 1. *There exists* $\widetilde{X} \in int(\mathcal{X})$*, such that* $\mathcal{U}\left(\widetilde{X}, \widetilde{X}\right) = \overline{V}$.

Assumption 1 states that the agents' utility in the static game is maximized at an interior optimum $\widetilde{X}$. Moreover, since $\mathcal{U}\left(x, \widetilde{X}\right) \leq \overline{V}$ for all $x \in \mathcal{X}$, the economic best-case scenario $\widetilde{X}$ also represents a symmetric Nash equilibrium of the economic stage game, *i.e.* the symmetric Nash equilibrium $\widetilde{X}$ decentralizes the social planner's solution. Assumption 1 says that our static economy admits a variant of the second welfare theorem or its analogue in frictional economies, which focuses our discussion on a benchmark in which the economy operates efficiently "in

[4] To simplify exposition and notation, we will focus throughout on symmetric pure strategy profiles and equilibria. We use x and X to distinguish between individual choices and aggregate variables.

normal times". Any rationale for active policy interventions then comes as a direct consequence of inefficient collective responses to the epidemic risk.

Confinement Stage Game. Let ρ denote the probability with which an agent is infected within a given period, conditional on being susceptible to infection. We assume that ρ varies with individual and aggregate choices, and in addition that it is proportional to the fraction of agents that are already infected, denoted by $\pi(i)$. Specifically, suppose that as a function of her choice $x \in \mathcal{X}$ and the aggregate action $X \in \mathcal{X}$, an agent is infected with probability $\rho = \mathcal{R}(x, X) \cdot \pi(i)$, where $\mathcal{R} : \mathcal{X} \times \mathcal{X} \to [\underline{R}, 1]$ is continuous, strictly convex and twice differentiable over the interior of $\mathcal{X} \times \mathcal{X}$. The aggregate infection rate at X is then given by $\mathcal{R}(X, X) \cdot \pi(i)$. We make the following additional assumption about $\mathcal{R}(\cdot)$:

Assumption 2. *There exists* $\hat{X} \in int(\mathcal{X})$*, such that* $\mathcal{R}\left(\hat{X}, \hat{X}\right) = \underline{R} \geq 0$*. Moreover* $\hat{X} \neq \widetilde{X}$*.*

Assumption 2 states that the agent's infection rate is minimized at an interior optimum $\hat{X}$, and since $\mathcal{R}\left(x, \hat{X}\right) \geq \underline{R}$, this action also aligns private and social returns from reducing infection risks. Assumption 2 is the direct analogue of assumption 1 for infection risk and implies that $\hat{X}$ represents a symmetric Nash equilibrium in the confinement game in which all agents aim to minimize infection risk $\mathcal{R}(x, X)$. This "extreme confinement equilibrium" $\hat{X}$ maximizes the long-term survival rate within the population and represents the best-case scenario from a health policy perspective.[5] Since $\hat{X} \neq \widetilde{X}$, there is a conflict between maximizing economic well-being $\mathcal{U}(\cdot)$ and minimizing infection risk $\mathcal{R}(\cdot)$ that is at the core of our analysis.

We let $\overline{R} = \mathcal{R}\left(\widetilde{X}, \widetilde{X}\right) > \underline{R}$ denote the infection risk at the economic optimum and $\underline{V} = \mathcal{U}\left(\hat{X}, \hat{X}\right) \in \left(0, \overline{V}\right)$ denote the instantaneous utility at the extreme confinement equilibrium. Any collective action with a strictly higher infection risk than $\overline{R}$ or lower welfare than $\underline{V}$ would be worse from the perspective of both economics and health care.

Remark. Both stage games implicitly assume scale invariance, *i.e.* instantaneous utilities and infection probabilities are independent of the mass of participating players. Scale invariance is common to many economic and epidemiological models in which interactions depend on the proportion of different types of agents in the population, rather than their absolute numbers. It is not critical for our analysis, but generates some useful simplifications along the way.

1.2 Dynamic Game

In the *dynamic game*, the economic stage game is infinitely repeated among the mass Λ_t of agents who remain alive in period t. The epidemic is summarized by a simple S-I-R structure:

[5] For example, $\hat{X}$ may be interpreted as a form of "extreme social distancing" to the point where there are no face-to-face interactions between any two agents at equilibrium: Suppose that an infection only occurs through physical contact between two individuals. Then if literally no one else is out on the streets, I will not be able to encounter anyone and hence not risk an infection, even if I am out on the street.

initially, a positive fraction is already infected with the disease, while the remainder is susceptible to infection. Susceptible agents become infected by interacting with other infected agents. After infection, an agent dies with constant probability δ and recovers with constant probability γ within each period; with probability $1-\gamma-\delta$ the agent remains infected the next period. Recovery confers immunity and is permanent. Only death is observable, so agents never know whether they are susceptible to infection, infected or have already recovered. Consistent with this assumption, their instantaneous utility function $\mathcal{U}(\cdot)$ is independent of their health status. Hence they are all ex ante identical, which simplifies the analytic characterization of the symmetric equilibrium and saves on the number of state variables. This assumption captures the reality of asymptomatic transmission and infections of COVID-19.

Conditional on surviving, each agent takes a sequence of decisions $x^\infty = \{x_t\}_{t=0}^\infty \in \mathcal{X}^\infty$ to maximize expected discounted utility flows, taking as given the choices $X^\infty = \{X_t\}_{t=0}^\infty \in \mathcal{X}^\infty$ of the other agents. We focus on a symmetric equilibrium, in which all agents take the same equilibrium action. An agent's expected discounted utility flow is

$$V_0 = (1-\beta)\sum_{t=0}^{\infty}\beta^t\Lambda_t\left(x^{t-1},X^{t-1}\right)\mathcal{U}(x_t,X_t)$$

where $\Lambda_t\left(x^{t-1},X^{t-1}\right)$ is the probability that the agent survives to period t, which is a function of the initial distribution π_0, individual and aggregate choices $\left(x^{t-1},X^{t-1}\right)$ up to period $t-1$, and $\beta \in (0,1)$ is the time discount factor. This welfare criterion summarizes the dynamic trade-off between instantaneous utilities $\mathcal{U}(x_t,X_t)$ and survival probabilities $\Lambda_t\left(x^{t-1},X^{t-1}\right)$. We assume perfect foresight, *i.e.* infection risk is idiosyncratic to each agent, but aggregate population shares of agents in different states are perfectly predictable.

The utilitarian social planner's solution consists of a sequence of symmetric choices $X^\infty \in \mathcal{X}^\infty$ that maximizes agents' expected discounted utility. A symmetric Nash equilibrium in the dynamic game is a sequence of symmetric choices $X^\infty \in \mathcal{X}^\infty$ that are optimal given that all agents also adhere to X^∞. Agents internalize the impact of their choices on their own infection and survival probabilities, but take aggregate transition rates as given.

We represent this dynamic game using the proportions $\pi_t(s)$ and $\pi_t(i)$ of susceptible and infected agents as state variables, taking the initial distribution as given with $\pi_0(i) > 0$ and $\pi_0(s) = 1-\pi_0(i) < 1$. We then characterize the planner's problem recursively as a function of the vector $\pi = (\pi_t(s),\pi_t(i))'$, and the equilibrium as a Markov-perfect equilibrium in π. The vector π admits the representation

$$\pi_{t+1} = \Lambda_t/\Lambda_{t+1}\cdot T(\rho_t)\,\pi_t, \text{ where } T(\rho) = \begin{pmatrix} 1-\rho & 0 \\ \rho & 1-\gamma-\delta \end{pmatrix}$$

and $\rho_t = \mathcal{R}(x_t,X_t)\cdot\pi_t(i)$ denotes the probability with which an agent is infected in period t,

as described above. The mass of surviving agents evolves according to

$$\Lambda_{t+1} = (1 - \delta\pi_t(i))\,\Lambda_t,$$

or $\Lambda(\pi) = \gamma/(\gamma + \delta(1 - \pi(i) - \pi(s)))$, as a function of the current population state π.[6]

In the dynamic model, current choices affect instantaneous utilities $\mathcal{U}(x, X)$ directly, and continuation values indirectly through their effect on the resulting infection rate $\mathcal{R}(x, X)$. We thus decompose the planner's solution and equilibrium characterization into implementation rules $X^*(\cdot)$ and $X^{eq}(\cdot)$ that determine optimal and equilibrium actions as functions of the implemented infection risk R, and a reduced-form interaction game that determines the choice of R recursively. We further characterize the planner's instantaneous utility $\mathcal{V}^*(R)$ as a function of infection risks R, and agents' instantaneous utility $\mathcal{V}^{eq}(r, R)$ as a function of their own choice r and the aggregate choice R of infection risks. We then study dynamic optimal policy and equilibrium through the lens of a reduced-form trade-off, treating infection risks R as the main policy variable, and provide conditions under which the two coincide.

Dynamic Planner's Problem. The utilitarian social planner's objective maximizes

$$V_0^* = \max_{X^\infty \in \mathcal{X}^\infty} (1-\beta) \sum_{t=0}^{\infty} \beta^t \Lambda_t\left(X^{t-1}, X^{t-1}\right) \mathcal{U}(X_t, X_t)$$

where $\Lambda_t\left(X^{t-1}, X^{t-1}\right)$ represents the fraction of agents alive in period t. Using the recursive characterization of $\Lambda_t\left(X^{t-1}, X^{t-1}\right)$, we represent the planner's value in period t as $V_t^* = \Lambda_t \cdot v^*(\pi_t)$, where $v^*(\pi)$ satisfies the following Bellman equation:

$$v^*(\pi) = \max_{X \in \mathcal{X}} \left\{(1-\beta)\mathcal{U}(X, X) + \beta(1 - \delta\pi(i))\, v^*(\pi_{+1})\right\}$$
$$\text{where } \pi_{+1} = (1 - \delta\pi(i))^{-1} \cdot T(\mathcal{R}(X, X)\,\pi(i)) \cdot \pi.$$

Let $X^*(\pi)$ denote the social planner's optimal decision rule. We decompose $X^*(\pi)$ into a target infection rate $R^*(\pi)$ and a static implementation rule $X^*(R)$. Let $\mathcal{V}^*(R)$ be given by the static constrained optimization problem

$$\mathcal{V}^*(R) \equiv \max_{X \in \mathcal{X}} \mathcal{U}(X, X) \text{ s.t. } \mathcal{R}(X, X) \leq R,$$

and $X^*(R) = \arg\max_{X \in \mathcal{X}, \mathcal{R}(X) \leq R} \mathcal{U}(X, X)$ as the planner's static implementation rule for a given target R. The function $\mathcal{V}^*(R)$ is strictly increasing, concave and satisfies the Inada conditions $\mathcal{V}^*\left(\overline{R}\right) = \overline{V}$, $\mathcal{V}^*(\underline{R}) = \underline{V}$, $\mathcal{V}^{*\prime}\left(\overline{R}\right) = 0$ and $\lim_{R \to \underline{R}} \mathcal{V}^{*\prime}(R) = \infty$. Using $\mathcal{V}^*(\cdot)$, we

[6] This characterization of $\Lambda(\pi)$ follows from $1 - \Lambda(\pi) = (\delta/\gamma)\,\Lambda(\pi)\,(1 - \pi(i) - \pi(s))$, *i.e.* the fraction of deceased $1 - \Lambda(\pi)$ must equal the fraction of agents who have recovered $\Lambda(\pi)(1 - \pi(i) - \pi(s))$, scaled by the odds ratio of dying vs. recovering δ/γ.

re-cast the planner's problem as a choice over R:

$$v^*(\pi) = \max_{R \in [\underline{R}, \overline{R}]} \{(1-\beta)\mathcal{V}^*(R) + \beta(1-\delta\pi(i))v^*(\pi_{+1})\}$$
$$\text{where } \pi_{+1} = (1-\delta\pi(i))^{-1} \cdot T(R\pi(i)) \cdot \pi.$$

Markov-Perfect Equilibrium. Consider now the dynamic decision problem of an individual agent. Let $X(\pi)$ denote the aggregate decision rule followed by the other agents, and let π^k denote agent k's private belief about her own infection state. The probability that the agent survives until next period is $1-\delta\pi^k(i)$.[7] Her decision problem is stated as follows

$$\hat{v}\left(\pi^k, \pi\right) = \max_{x \in \mathcal{X}} \left\{(1-\beta)\mathcal{U}(x, X(\pi)) + \beta\left(1-\delta\pi^k(i)\right)\hat{v}\left(\pi^k_{+1}, \pi_{+1}\right)\right\}$$
$$\text{where } \pi^k_{+1} = \left(1-\delta\pi^k(i)\right)^{-1} \cdot T(\mathcal{R}(x, X(\pi))\pi(i)) \cdot \pi^k$$
$$\pi_{+1} = (1-\delta\pi(i))^{-1} \cdot T(\mathcal{R}(X(\pi), X(\pi))\pi(i)) \cdot \pi$$

A symmetric Markov-perfect equilibrium consists of an aggregate choice function $X^{eq}(\cdot)$ and a value function $\hat{v}\left(\pi^k, \pi\right)$, such that $X^{eq}(\cdot)$ is a best response to itself given a private belief $\pi^k = \pi$ and continuation values $\hat{v}\left(\pi^k_{+1}, \pi_{+1}\right)$, and $\hat{v}\left(\pi^k, \pi\right)$ solves the above Bellman equation for an aggregate decision rule $X(\cdot) = X^{eq}(\cdot)$.

Like the planner's solution, we decompose the Markov-perfect equilibrium into a static implementation rule $X^{eq}(R)$ for a given target infection rate R, and a reduced form dynamic interaction game that determines the equilibrium target infection rate $R^{eq}(\pi)$. For given aggregate choice X, consider the static decision problem

$$\max_{x \in \mathcal{X}} \mathcal{U}(x, X) \text{ s.t. } \mathcal{R}(x, X) \leq r,$$

and let $\hat{x}(r, X)$ denote the best-response action that implements infection risk r, given an aggregate choice of X. A static implementation rule $X^{eq}(R)$ can be part of a symmetric Markov-perfect equilibrium, if and only if $X^{eq}(R)$ is a best response to itself, *i.e.* $X^{eq}(R) = \hat{x}(R, X^{eq}(R))$.

Now, define $\mathcal{V}^{eq}(r, R) = \mathcal{U}(\hat{x}(r, X^{eq}(R)), X^{eq}(R))$ as the indirect utility of choosing r when the aggregate action $X^{eq}(R)$ implements an equilibrium infection rate R. This reduced-form utility function is strictly increasing, concave in r, and satisfies the Inada conditions at $\overline{R}(R) = \max_{x \in \mathcal{X}} \mathcal{R}(x, X^{eq}(R)) \geq \mathcal{R}(X^{eq}(R), X^{eq}(R))$ and $\underline{R}(R) = \min_{x \in \mathcal{X}} \mathcal{R}(x, X^{eq}(R)) \leq \mathcal{R}(X^{eq}(R), X^{eq}(R))$ whenever $\arg\min_{x \in \mathcal{X}} \mathcal{R}(x, X^{eq}(R)) \in int(\mathcal{X})$.[8] Hence if infection rates

[7] Individual survival probabilities evolve recursively according to $\Lambda_{t+1}\left(x^t, X^t\right) = \Lambda_t\left(x^{t-1}, X^{t-1}\right) \cdot \left(1-\delta\pi^k_t(i)\right)$, or $\Lambda\left(\pi^k\right) = \gamma / \left(\gamma + \delta\left(1-\pi^k_t(i) - \pi^k_t(s)\right)\right)$.

[8] When $K = 1$, the solution to the constrained optimization problem is determined directly from the constraint: $\mathcal{V}^{eq}(r, R) = \mathcal{U}(x(r, R), X^{eq}(R))$, where $x(r, R)$ solves $r = \mathcal{R}(x, X^{eq}(R))$.

are minimized on the interior of $\mathcal{X}$, we recover the same Inada conditions for the private efficiency frontier as for the planner, but its support varies with the aggregate infection rate R.

The equilibrium infection rate is determined as the outcome of a reduced form interaction game in the choice of $r \in \left[\underline{R}, \overline{R}\right]$, using the indirected utility function $\mathcal{V}^{eq}(r, R)$: Given an aggregate infection rate function $R(\pi)$, we restate $\hat{v}\left(\pi^k, \pi\right)$ as follows:

$$\hat{v}\left(\pi^k, \pi\right) = \max_{r \in \left[\underline{R}, \overline{R}\right]} \left\{(1-\beta)\,\mathcal{V}^{eq}(r, R(\pi)) + \beta\left(1 - \delta\pi^k(i)\right)\hat{v}\left(\pi^k_{+1}, \pi_{+1}\right)\right\}$$

$$\text{where } \pi^k_{+1} = \left(1 - \delta\pi^k(i)\right)^{-1} \cdot T(r\pi(i)) \cdot \pi^k$$

$$\pi_{+1} = (1 - \delta\pi(i))^{-1} \cdot T(R(\pi)\,\pi(i)) \cdot \pi$$

A symmetric Markov perfect equilibrium is then characterized by a target infection rate function $R^{eq}(\cdot)$ and a value function $\hat{v}(\cdot,\cdot)$, such that $R^{eq}(\pi)$ is a best response to itself given an initial private belief $\pi^k = \pi$ and continuation values $\hat{v}\left(\pi^k_{+1}, \pi_{+1}\right)$, and $\hat{v}(\cdot,\cdot)$ solves the above Bellman equation for an aggregate infection rate function $R(\cdot) = R^{eq}(\cdot)$.

Figure 1: Equilibrium and Pareto Frontiers

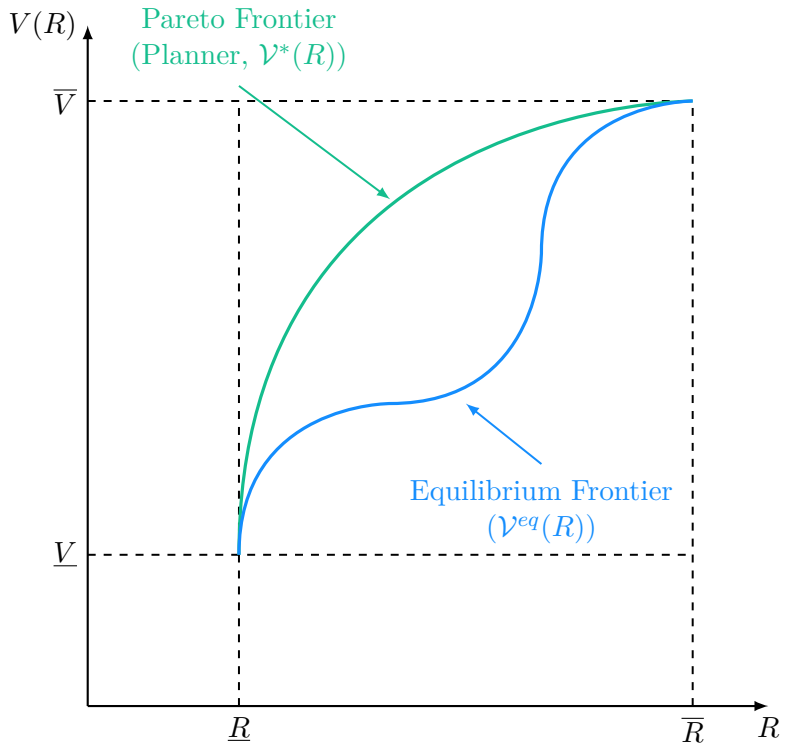

Figure 1 summarizes the characterization of the planner's Pareto frontier $\mathcal{V}^*(R)$ (in green) and the equilibrium frontier $\mathcal{V}^{eq}(R, R)$ (in blue). The general shape of $\mathcal{V}^{eq}(\cdot)$ depends on economic and infection risk spill-overs, but $\mathcal{V}^{eq}(\cdot)$ satisfies the same Inada conditions as $\mathcal{V}^*(R)$ at $\overline{R}$ and $\underline{R}$, since the equilibria and planner's solutions both converge to the same limit when the hybrid game converges to either the economic stage game (Assumption 1) or the confinement

game (Assumption 2). Without these assumptions, the Inada conditions no longer hold and $\mathcal{V}^{eq}(\cdot)$ can take any shape below $\mathcal{V}^{*}(\cdot)$ at its boundaries.

1.3 Static equilibrium and efficiency conditions

The planner's implementation rule $X^{*}(R)$ satisfies the static optimality condition

$$\nabla \mathcal{U}(X^{*}, X^{*}) = \mathcal{V}^{*\prime}(R) \cdot \nabla \mathcal{R}(X^{*}, X^{*}), \tag{1}$$

where ∇ denotes the total gradient with respect to the individual action $x = X^{*}$ and the aggregate action $X = X^{*}$. Equation (1) gives us the standard result that the planner's solution equates the social marginal rates of substitution in $\mathcal{U}(\cdot)$ to the social marginal rates of substitution in $\mathcal{R}(\cdot)$. Specifically, along all dimensions of $\mathcal{X}$, the planner equates the marginal trade-offs in instantaneous utility and infection rates to the planner's marginal utility benefit of R, $\mathcal{V}^{*\prime}(R)$.

The equilibrium implementation rule $X^{eq}(R)$ satisfies the static first-order condition

$$\nabla_1 \mathcal{U}(X^{eq}, X^{eq}) = \mathcal{V}_r^{eq}(R, R)\, \nabla_1 \mathcal{R}(X^{eq}, X^{eq}), \tag{2}$$

where ∇_1 denotes the gradient with respect to the individual action x. The static equilibrium condition shows that the same margin of substitution between different dimensions operates in the private decisions, but the trade-offs that individual agents are facing may be different from the ones faced by the planner. In other words, agents privately maximize static utility subject to an upper bound on private infection risks.

The equilibrium implementation rule is efficient, if and only if $X^{eq}(R) = X^{*}(R)$ and $\mathcal{V}^{*}(R) = \mathcal{V}^{eq}(R, R)$, *i.e.* the solution to the static planner's and equilibrium first-order conditions coincide. Taking the ratio of these two first-order conditions for all $i \in \{1, \ldots, K\}$, we obtain:

$$\left. \frac{1 + \frac{\partial \mathcal{U}(X,X)}{\partial X_i} \Big/ \frac{\partial \mathcal{U}(x,X)}{\partial x_i}}{1 + \frac{\partial \mathcal{R}(x,X)}{\partial X_i} \Big/ \frac{\partial \mathcal{R}(x,X)}{\partial x_i}} \right|_{x=X=X^{*}(R)} = \frac{\mathcal{V}^{*\prime}(R)}{\mathcal{V}_r^{eq}(R, R)}. \tag{3}$$

Let us for now take $\mathcal{V}^{*\prime}(R) / \mathcal{V}_r^{eq}(R, R)$ as given. Equation (3) states that the planner's implementation rule is an equilibrium if and only if the ratio of social to private marginal rates of substitution at the planner's solution is equal to $\mathcal{V}^{*\prime}(R) / \mathcal{V}_r^{eq}(R, R)$ along each dimension. When $\mathcal{V}^{*\prime}(R) = \mathcal{V}_r^{eq}(R, R)$, equation (3) reduces to the usual equalization of private and social marginal rates of substitution between utility and infection risks, *i.e.* utility spill-overs of individual choices $\frac{\partial \mathcal{U}(X,X)}{\partial X_i} / \frac{\partial \mathcal{U}(x,X)}{\partial x_i}$ are exactly equal to infection risk spill-overs $\frac{\partial \mathcal{R}(x,X)}{\partial X_i} / \frac{\partial \mathcal{R}(x,X)}{\partial x_i}$ along each dimension. When $\mathcal{V}^{*\prime}(R) > \mathcal{V}_r^{eq}(R, R)$, utility spill-overs must dominate infection risk spill-overs along each dimension. When instead $\mathcal{V}^{*\prime}(R) < \mathcal{V}_r^{eq}(R, R)$, infection risk spill-overs must dominate utility spill-overs along each dimension.

Departures from equation (3) are akin to a cross-sectional efficiency or mis-allocation wedge, resulting in a gap between the planner's and the equilibrium efficiency frontier between infection risk and instantaneous utility due to a mis-allocation of economic activity across sectors.

We do not need to take a stance on the direction of static spill-overs: actions can have positive or negative economic or infection risk spill-overs. In particular agents in our model may fail to internalize that (*i*) reducing activities as privately optimal precaution against infection risk exposes others to negative economic spill-overs and (*ii*) that their own economic activity exposes others to increased infection risks. Equation (3) shows that the planner doesn't weigh them in terms of their absolute, but their relative strengths.

What does this result tell us about optimal policy design? Consider a planner who can impose restrictions $\widehat{\mathcal{X}} \subset \mathcal{X}$ on the choice sets of agents to bring the equilibrium in line with the planner's solution. Our static efficiency and equilibrium conditions offer a foundation to Budish (2020) who suggests that economic policy during the pandemic should be framed as maximizing instantaneous utilities subject to an upper bound constraint on infection risk.[9] Applying this principle to static policy decisions suggests the following simple policy principles:

(*i*) Policy should restrict *critical* activities that generate strong infection risk externalities but marginal economic spill-overs such as socializing, going out to restaurants, large scale entertainment or inessential long-distance travel.

(*ii*) Policy should also protect or subsidize *essential* activities that have strong positive economic spill-overs, especially if they have weak infection risk externalities, for example banking and payment systems, healthcare, public education. If such activities require customer interactions, then policy should focus on managing or mitigating infection risks, for example by substituting face-to-face with online interactions or by mandating the use of face masks.

In particular policy must very carefully manage activities that are both economically essential and critical from an infection risk point of view (high negative infection rate spill-overs), since these are the sectors that have the strongest impact on whether the economic-infection risk trade-off is resolved efficiently.

These principles are summarized in Figure 2.

Of course many of the policy measures adopted in the last few months intuitively follow these principles, but it is worth underlining that (*i*) the rationale for active policy interventions still depends on the presence of economic or infection risk spill-overs that are not fully internalized by private decisions, and (*ii*) the relative size of economic and infection risk spill-overs, not their absolute magnitudes determines to what extent sectors should be restricted or subsidized.

[9] Budish (2020) goes a step further and argues that the upper bound constraint on infection risk should keep the basic reproduction rate of the virus below 1. In the next sections we show how this specific upper-bound constraint emerges from the properties of a fast-moving pandemic with asymptomatic infections.

Figure 2: Policy Design

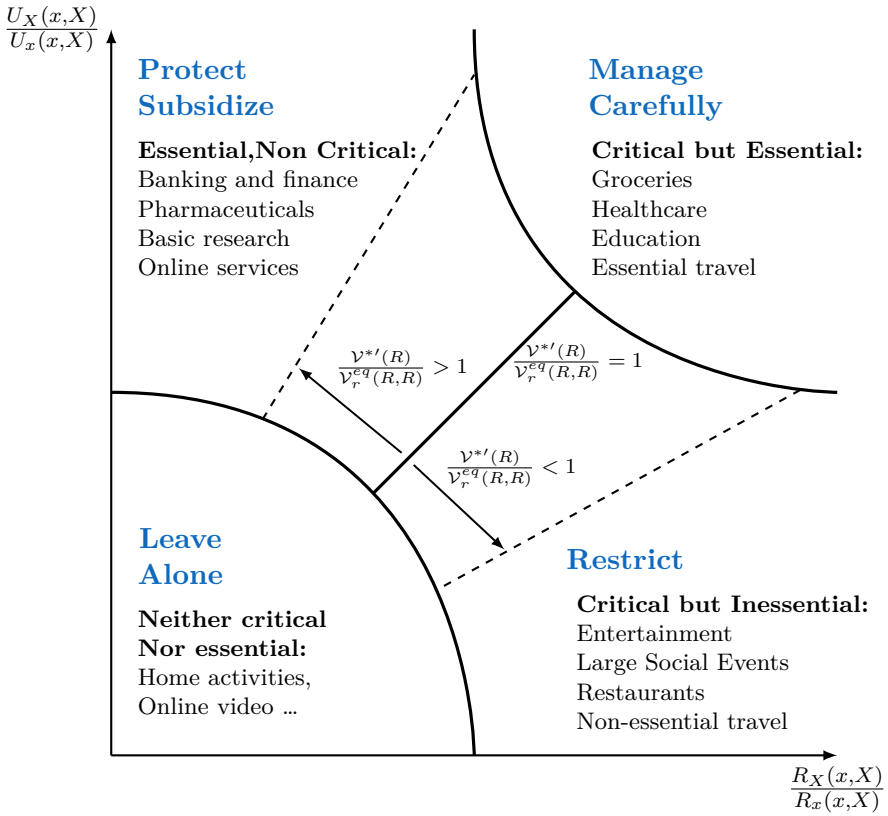

1.4 Dynamic Equilibrium and Efficiency Conditions

Taking first-order conditions,[10] the planner's optimal choice of R is characterized by

$$\mathcal{V}^{*\prime}(R) = \Phi^*(\pi) \equiv \frac{\beta}{1-\beta}\pi(i)\,\pi(s)\left(\frac{\partial v^*(\pi_{+1})}{\partial \pi(s)} - \frac{\partial v^*(\pi_{+1})}{\partial \pi(i)}\right)\Big|_{R=R^*(\pi)}. \tag{4}$$

The equilibrium infection rate $R^{eq}(\cdot)$ is represented by

$$\mathcal{V}_r^{eq}(R,R) = \Phi^{eq}(\pi) \equiv \frac{\beta}{1-\beta}\pi(i)\,\pi(s)\left(\frac{\partial \hat{v}(\pi_{+1},\pi_{+1})}{\partial \pi^k(s)} - \frac{\partial \hat{v}(\pi_{+1},\pi_{+1})}{\partial \pi^k(i)}\right)\Big|_{R=R^{eq}(\pi)}. \tag{5}$$

These dynamic first-order conditions equate the social or private marginal utility benefit of a higher target infection rate R (i.e. the marginal disutility of social distancing) to a social or private marginal cost $\Phi^*(\cdot)$ or $\Phi^{eq}(\cdot)$ that we interpret as the planner's or equilibrium *shadow cost of infection risk*. This shadow cost is equal to the marginal social or private cost of an additional infection $\frac{\partial v^*(\pi_{+1})}{\partial \pi(s)} - \frac{\partial v^*(\pi_{+1})}{\partial \pi(i)}$ (or $\frac{\partial \hat{v}(\pi_{+1},\pi_{+1})}{\partial \pi^k(s)} - \frac{\partial \hat{v}(\pi_{+1},\pi_{+1})}{\partial \pi^k(i)}$), scaled by the inverse of the time preference rate $\beta/(1-\beta)$, and the product of the proportion of infected and susceptible agents. This product measures the rate of interactions between these two groups, which scales the primitive infection risk in our model. $\Phi^*(\cdot)$ and $\Phi^{eq}(\cdot)$ are functions of the current state

[10] We exploit that R affects π_{+1} linearly as a one-for-one increase in $\pi(i)$ and reduction in $\pi(s)$.

π.[11]

In the appendix, we show that $v^*(\pi)$ and $\hat{v}(\pi,\pi)$ can be represented as probability-weighted expectations of life-time utilities:

$$\begin{aligned} v^*(\pi) &= \pi(s)\, v_s^*(\pi) + \pi(i)\, v_i^*(\pi) + (1-\pi(s)-\pi(i))\, v_r^*(\pi) \\ \hat{v}\left(\pi^k,\pi\right) &= \pi^k(s)\,\hat{v}_s\left(\pi^k,\pi\right) + \pi^k(i)\,\hat{v}_i\left(\pi^k,\pi\right) + \left(1-\pi^k(s)-\pi^k(i)\right)\hat{v}_r\left(\pi^k,\pi\right), \end{aligned}$$

where $v_s^*(\pi)$, $v_i^*(\pi)$ and $v_r^*(\pi)$ denote the life-time utility of a susceptible, infected and recovered agent at the planner's solution and $\hat{v}_s\left(\pi^k,\pi\right)$, $\hat{v}_i\left(\pi^k,\pi\right)$, and $\hat{v}_r\left(\pi^k,\pi\right)$ denote the life-time values of being in state s, i, or r, given a private belief π^k and aggregate state π. What's more, it follows from a simple envelope condition[12] that the private marginal cost of infection satisfies

$$\left(\frac{\partial\hat{v}\left(\pi^k,\pi\right)}{\partial\pi^k(s)} - \frac{\partial\hat{v}\left(\pi^k,\pi\right)}{\partial\pi^k(i)}\right)\Bigg|_{\pi^k=\pi} = \hat{v}_s(\pi,\pi) - \hat{v}_i(\pi,\pi) \tag{6}$$

In words, the private marginal cost of infection equals the direct cost of an infection, or of moving an agent from state s to state i.

In contrast, the social marginal cost of infection satisfies

$$-\frac{\partial v^*(\pi)}{\partial\pi(i)} = v_r^*(\pi) - v_i^*(\pi) - \pi(s)\frac{\partial v_s^*(\pi)}{\partial\pi(i)} - \pi(i)\frac{\partial v_i^*(\pi)}{\partial\pi(i)} - (1-\pi(s)-\pi(i))\frac{\partial v_r^*(\pi)}{\partial\pi(i)} \tag{7}$$

$$-\frac{\partial v^*(\pi)}{\partial\pi(s)} = v_r^*(\pi) - v_s^*(\pi) - \pi(s)\frac{\partial v_s^*(\pi)}{\partial\pi(s)} - \pi(i)\frac{\partial v_i^*(\pi)}{\partial\pi(s)} - (1-\pi(s)-\pi(i))\frac{\partial v_r^*(\pi)}{\partial\pi(s)} \tag{8}$$

The expression $-\frac{\partial v^*(\pi)}{\partial\pi(i)}$ measures the social marginal value of recovery, *i.e.* of shifting an agent from state i to state r. This marginal value consists of the direct benefit of recovery $v_r^*(\pi) - v_i^*(\pi) > 0$ that an agent enjoys by recovering from the disease, and the indirect effects a marginal decrease of the infection rate has on susceptible, infected and recovered agents. These terms capture externalities of two types. First, a reduction in the infection rate permits to ease confinement and increase the utility of all agents. Second, the term $-\frac{\partial v_s^*(\pi)}{\partial\pi(i)}$ reflects dynamic infection externalities: reducing the infection rate lowers infection risks for other susceptible agents in the future.

The expression $-\frac{\partial v^*(\pi)}{\partial\pi(s)}$ measures the social marginal value of immunization, *i.e.* of shifting an agent from state s to state r. Again this marginal value consists of a direct benefit of immunization $v_r^*(\pi) - v_s^*(\pi) > 0$, and indirect effects through which lowering the share of susceptibles affects the rest of the population. These expressions reveal the presence of a second externality: higher immunization reduces the need for economic restrictions. More precisely,

[11] In the equilibrium first-order condition, the risk of being infected multiplies the aggregate infection rate $\pi(i)$ with the individual probability of being susceptible $\pi^k(s)$, which equals $\pi(s)$ at a symmetric equilibrium.

[12] Define $\widetilde{v}\left(\pi^k,\pi\right) = \pi^k(s)\,\hat{v}_s(\pi,\pi) + \pi^k(i)\,\hat{v}_i(\pi,\pi) + \left(1-\pi^k(s)-\pi^k(i)\right)\hat{v}_r(\pi,\pi)$. Since $\widetilde{v}\left(\pi^k,\pi\right)$ is linear in π^k and $\hat{v}\left(\pi^k,\pi\right) \geq \widetilde{v}\left(\pi^k,\pi\right)$ for π^k close to π, with equality when $\pi^k=\pi$, Benveniste and Scheinkman's (1979) theorem applies.

the only way to control the epidemic in our baseline model is to reach herd immunity in the long-run. A larger number of immune agents brings the economy closer to herd immunity, and therefore closer to the point where economic restrictions can be lifted, which raises the present values of all agents.

By subtracting the marginal value of immunization from the marginal value of recovery, we obtain $\frac{\partial v^*(\pi)}{\partial \pi(s)} - \frac{\partial v^*(\pi)}{\partial \pi(i)}$, the social marginal cost of an additional infection. This social marginal cost also combines a direct cost of infection $v_s^*(\pi) - v_i^*(\pi)$ with indirect costs coming from the spill-over effects of the additional infection for other agents: increasing infection risks for other susceptibles, but relaxing future economic restrictions.[13]

At equilibrium, the private marginal cost of an additional infection equals the direct cost, but takes as given the future dynamics of the aggregate population state. In contrast to the planner's solution, the equilibrium does not internalize the probability-weighted indirect effects of an additional infection on the continuation values of each type. In other words, agents at equilibrium do not internalize dynamic immunization and infection externalities.

In the appendix we show that the direct cost of infection is strictly positive and uniformly bounded away from 0. On the other hand, we are unable to uniformly bound or sign the externality terms, so that $\frac{\partial v^*(\pi)}{\partial \pi(s)} - \frac{\partial v^*(\pi)}{\partial \pi(i)}$ may indeed be arbitrarily close to 0, if the externalities partly offset the direct costs of infection.

1.5 Efficient Decentralization

Summarizing the two pairs of first-order conditions, Proposition 1 provides necessary and sufficient conditions for efficiency of the Markov perfect equilibrium.

Proposition 1. *The planner's solution $X^*(\pi)$ is a Markov-perfect equilibrium, if and only if (i) $X^*(R) = X^{eq}(R)$ and (ii) $\mathcal{V}_R^{eq}(R,R) = \Phi^*(\pi) - \Phi^{eq}(\pi)$ for all π and $R = R^*(\pi)$.*

Proposition 1 gives two necessary and sufficient conditions for efficiency in the reduced form game: (i) Efficient implementation ($X^{eq}(R) = X^*(R)$ and $\mathcal{V}^{eq}(R,R) = \mathcal{V}^*(R)$), and (ii) "Offsetting static and dynamic spill-overs". The efficient implementation conditions were discussed above as requiring that static economic and infection risk spill-overs are fully internalized by the equilibrium (i.e. no mis-allocation across different sectors or dimensions of economic activities). In the one-dimensional case ($K = 1$), this condition is automatically satisfied.

The condition of offsetting static and dynamic spill-overs requires that static economic and infection risk spill-overs that are summarized by $\mathcal{V}_R^{eq}(R,R)$ exactly offset dynamic immunization and infection externalities that are summarized by $\Phi^*(\pi) - \Phi^{eq}(\pi)$, or $\mathcal{V}_R^{eq}(R^*(\pi), R^*(\pi)) =$

[13] This interpretation of marginal effects is adopted from Garibaldi, Moen, and Pissarides (2020), though they do not distinguish between the direct and indirect effects, and they stop well short of fully characterizing the dynamics of these externalities.

$\Phi^*(\pi) - \Phi^{eq}(\pi)$ for all π.[14] This condition is equivalent to

$$\frac{\mathcal{V}^{*\prime}(R)}{\mathcal{V}_r^{eq}(R,R)} = \frac{\Phi^*(\pi)}{\Phi^{eq}(\pi)}$$

for all π and $R = R^*(\pi)$, which in turn allows us to link the right-hand side of the static efficiency condition (3) to the dynamic first-order conditions (4) and (5).

Proposition 1 shows that the equilibrium choice of R can be efficient only if by coincidence the static economic and infection risk spill-overs at the planner's solution are exactly offset by dynamic infection or immunization externalities. This is an extremely stringent condition that, even if it holds momentarily at a point in time, is bound to be violated almost everywhere in a dynamic context. In particular, the offsetting spill-overs condition requires that any two population states that deliver the same policy $R^*(\pi)$ also generate exactly the same dynamic spill-overs. As we shall see, the natural policy response to the pandemic is hump-shaped, first tightening then relaxing policy restrictions. But since the relative importance of immunization and infection risk externalities evolves as the pandemic progresses, dynamic spill-overs, summarized by $\Phi^*(\pi) - \Phi^{eq}(\pi)$, will necessarily evolve over the course of the pandemic. These dynamics will be discussed at length in the next section.

2 Dynamic Equilibrium and Optimal Policy

We now link the recursive optimality and equilibrium conditions to the dynamics of π that are generated by the SIR model to characterize the equilibrium and optimal policy dynamics over the course of the pandemic, starting from an arbitrary initial infection rate $\pi_0(i) > 0$.

We assume that $\gamma + \delta \in \left(\underline{R}, \overline{R}\right)$. If $\overline{R} > \gamma + \delta$, the basic reproduction rate $\mathcal{R}_0 = \overline{R}/(\gamma + \delta)$ at the pre-pandemic equilibrium exceeds 1, and hence the initial infection, however small, can grow exponentially and take hold within the population if there is no behavioral response to the infection risk. However, since $\gamma + \delta > \underline{R}$ the disease may immediately be contained and eventually be suppressed with a value of R sufficiently close to $\underline{R}$.

The dynamics of $\pi_t(i)$ satisfy

$$\pi_{t+1}(i) = \frac{R_t \pi_t(s) + 1 - \gamma - \delta}{1 - \delta \pi_t(i)} \pi_t(i).$$

For constant R_t, $\pi_t(i)$ exhibits a hump-shaped profile, which is at first increasing, and then decreasing once $R_t \pi_t(s) \leq \gamma + \delta(1 - \pi_t(i))$. For $R(\cdot) \in \{R^*(\cdot), R^{eq}(\cdot)\}$, the effective reproduction rate as a a function of the population state π, is

$$\mathcal{R}_0(\pi) = \frac{R(\pi)\pi(s)}{\gamma + \delta}.$$

[14] This condition is obtained by taking the difference between the planner's and equilibrium first-order conditions $\mathcal{V}^{*\prime}(R) = \Phi^*(\pi)$ and $\mathcal{V}_r^{eq}(R,R) = \Phi^{eq}(\pi)$, evaluated at the planner's solution $R = R^*(\pi)$, and noting that $\mathcal{V}^{*\prime}(R) = \mathcal{V}_R^{eq}(R,R) + \mathcal{V}_r^{eq}(R,R)$.

Since $1-\pi_t(s)-\pi_t(i)$ is increasing and bounded, the population state π_t must converge to a limit π_∞ at which $\pi_\infty(i)=0$, $\pi_\infty(s)\in(0,1)$, and Λ_t converges to a finite limit

$$\Lambda_\infty=\frac{\gamma}{\gamma+\delta\left(1-\pi_\infty(s)\right)}\in\left(\frac{\gamma}{\gamma+\delta},1\right).$$

Let $\{R_t^*,\pi_t^*\}$ and $\{R_t^{eq},\pi_t^{eq}\}$ denote the sequential planner's solution and equilibrium for a given initial distribution π_0. $\{R_t^*,\pi_t^*\}$ and $\{R_t^{eq},\pi_t^{eq}\}$ must satisfy

$$R_t^*=R^*(\pi_t^*)\text{ and }\pi_{t+1}^*=\left(1-\delta\pi_t^*(i)\right)^{-1}\cdot T\left(R^*(\pi_t^*)\,\pi_t^*(i)\right)\cdot\pi_t^*$$
$$R_t^{eq}=R^{eq}(\pi_t^{eq})\text{ and }\pi_{t+1}^{eq}=\left(1-\delta\pi_t^{eq}(i)\right)^{-1}\cdot T\left(R^{eq}(\pi_t^{eq})\,\pi_t^{eq}(i)\right)\cdot\pi_t^{eq}.$$

2.1 Flattening the Curve

Combining the above dynamics with the two first-order conditions (4) and (5) implies that equilibrium and optimal policies generate a hump-shaped response to infection risks that follows the dynamics of $\pi(i)$. This result follows from $\Phi^*(\pi),\Phi^{eq}(\pi)\propto(\beta/(1-\beta))\,\pi(i)$, *i.e.* the shadow value of infection risk, and hence the marginal disutility of equilibrium and optimal social distancing policies, are proportional to the fraction of currently infected agents $\pi(i)$, and the short-run and long-run properties of the population dynamics implied by the S-I-R model. This immediately leads to Proposition 2 which states that equilibrium and optimal policies *flatten the infection curve* in the short run, but do not suppress new infections before the population reaches a state of *herd immunity* in the long run:

Proposition 2. *Starting from any small positive initial fraction $\pi_0(i)>0$ of infected agents in the population, the sequential planner's solution and equilibrium $\{R_t^*,\pi_t^*\}$ and $\{R_t^{eq},\pi_t^{eq}\}$ both satisfy the following properties:*

(i) ***Flatten the Curve (Short Run):*** *Starting from R_0^* and R_0^{eq} arbitrarily close to $\overline{R}$, both policy sequences are initially decreasing to "flatten the curve" and delay infections.*

(ii) ***Herd Immunity (Long-run):*** *In the long run, R_t^* and R_t^{eq} converge to $\overline{R}$, and the economy returns to the pre-pandemic equilibrium in a state of herd immunity:*

$$\pi_\infty^*(s),\pi_\infty^{eq}(s)\leq\frac{\gamma+\delta}{\overline{R}}\text{ and }\Lambda_\infty^*,\Lambda_\infty^{eq}\leq\Lambda\left(\overline{R}\right)\equiv\frac{\gamma\overline{R}}{(\gamma+\delta)\left(\overline{R}-\delta\right)}.$$

Proposition 2 highlights two properties of the equilibrium and optimal policy response.

First, the optimal and the equilibrium policies both *flatten the infection curve* by moving away from the utility maximizing action towards the infection risk minimizing one at the onset of the pandemic. As a result, infections peak later and at a lower level than without a behavioral response. The rationale for flattening the infection curve emerges without reference to the usual medical arguments in favor of such policies: flattening the curve neither serves to gain time until a vaccine or cure is found, nor does it serve to decongest the medical sector.[15] Instead

[15] We will discuss these channels as extensions to our baseline model in Section 3.

this shape of the optimal policy is a result of its economic benefits: Flattening the curve slows the propagation of the infection, which improves the survival rates for each individual agent.

Second, in the long run, the epidemic must subside, and both equilibrium and planner's solutions revert back to the pre-pandemic equilibrium. This however is possible only once a sufficiently large number of agents has been infected and recovered from the disease to establish *herd immunity*. But this also bounds the number of agents that can be saved in the long run, since for each γ agents that recover from the disease, δ will have died.

It is not too surprising that the equilibrium dynamics converge towards herd immunity since private incentives for social distancing disappear as $\pi(i)$ converges back to 0. It may be more surprising to see that the same long-run outcome emerges as part of the planner's solution, since the planner faces a long-run trade off between $\mathcal{V}^*(R)$ and Λ_∞, and she could in principle generate a permanent first-order increase in survival probabilities by implementing a small permanent distortion that has second-order marginal utility costs. However, the planner also factors in the delay between the marginal cost of reducing R today and the marginal benefit of higher future mortality. With discounting, this delay explains why it is optimal to slow the propagation of the pandemic, but not permanently suppress it before reaching herd immunity to raise the agents' long-run survival probability.

2.2 Fast Propagation: The Hammer and the Dance

We sharpen the characterization of equilibrium and optimal policy dynamics by considering the limit when $\beta \to 1$. Interpreting $\beta/(1-\beta)$ as the speed of propagation, this limit focuses on a case where the propagation of the disease is arbitrarily fast. Since $\Phi^*(\pi), \Phi^{eq}(\pi) \propto (\beta/(1-\beta))\,\pi(i)$, the shadow price of infection risk grows arbitrarily large when the infection rate $\pi(i)$ and the marginal cost of infection are bounded away from 0.

Starting from any $\pi_0(i) > 0$, the equilibrium dynamics then go through two phases that we will refer to as the *Hammer* and the *Dance*:[16] a first phase ("the Hammer") of extreme social distancing with R_t^{eq} close to $\underline{R}$ that serves to bring the rate of infections $\pi(i)$ close to 0, and a second phase ("the Dance") in which $\pi(i)$ stays close to 0, and the pandemic gradually progresses towards herd immunity. In this second phase, the rate of infections $\pi(i)$ can neither grow nor decay exponentially, i.e. $\mathcal{R}_0(\pi_t)$ is close to 1. As a consequence, the basic reproduction rate $\mathcal{R}_0$ converges to 1: the equilibrium policy contains but does not suppress the pandemic.

The social planner's solution displays the same hammer and dance dynamics, provided that the social marginal cost of infection is uniformly bounded away from 0. These dynamics are summarized in Proposition 3.

[16] These labels refer to Pueyo (2020) who proposes these phases as a possible strategy for deconfinement.

Proposition 3 (The Hammer and the Dance)**.** *For any $\eta > 0$, there exists $\xi > 0$, such that with $\beta > 1 - \xi$, equilibrium $\{R_t^{eq}, \pi_t^{eq}\}$ has the following structure:*

(i) $R_t^{eq} < \underline{R} + \eta$ whenever $\pi_t(i) > \eta$.

(ii) Starting from $\pi_0(i) > \eta$, equilibrium policy dynamics consist of two phases:

***1. The Hammer:** An initial phase of massive confinement in which R_t^{eq} are kept below $\underline{R} + \eta$ until $\pi_t(i) < \eta$.*

***2. The Dance:** A subsequent phase of gradual deconfinement, in which $\pi_t(i)$ remains stabilized within $(0, \eta)$, while $\pi_t(s)$ gradually declines at rate less than $R_t^{eq}\eta$, and R_t^{eq} stays close to $(\gamma + \delta)/\pi_t(s)$. The Dance ends when $\pi_t(s)$ reaches the herd immunity threshold $(\gamma + \delta)/\overline{R}$, $\pi_t(i)$ converges to 0, and R_t^{eq} converges to $\overline{R}$.*

Optimal policy $\{R_t^*, \pi_t^*\}$ follows a similar path, provided that the social marginal cost of an infection, $\frac{\partial v^*(\pi_{+1})}{\partial \pi(s)} - \frac{\partial v^*(\pi_{+1})}{\partial \pi(i)}$, is uniformly bounded away from 0.

Figure 3 illustrates the equilibrium and optimal policy dynamics with fast propagation described by Proposition 3. We report the equilibrium (blue) and Pareto (green) frontiers and the (red) line determining the infection risk at which the basic reproduction rate equals 1. At equilibrium, the shadow price of infection risks $\Phi^{eq}(\pi)$ becomes arbitrarily large, whenever $\pi(i)$ is sufficiently far from 0. When faced with such a situation, agents enact a massive voluntary confinement ("*The Hammer*"), with $R^{eq}(\pi_t)$ arbitrarily close to the extreme confinement equilibrium $\underline{R}$ until the proportion of infected agents is controlled within a narrow band $\pi(i) \in (0, \eta)$.[17] This corresponds to the shift depicted by arrow H^{eq} in Figure 3. In the second phase, the equilibrium is delicately balanced to keep $\pi(i)$ within this narrow band ("*The Dance*"), letting the epidemic slowly progress until eventually herd immunity is reached and it is allowed to fizzle out. This corresponds to the displacement D^{eq} along the equilibrium frontier in Figure 3. During the Dance phase, the equilibrium policy $R^{eq}(\pi_t)$ cannot stray far from $(\gamma + \delta)/\pi_t(s)$ —the level indicated by the red line at which the basic reproduction rate $\mathcal{R}_0(\pi_t)$ of the infection equals 1. The policy then gradually reverts back to $\overline{R}$ as the red line moves to the right and $\pi_t(s)$ converges towards herd immunity.[18]

The optimal policy follows a similar pattern as the equilibrium, but with one important qualification: whereas the private marginal cost of infection is uniformly bounded away from 0, the social marginal cost $\frac{\partial v^*(\pi_{+1})}{\partial \pi(s)} - \frac{\partial v^*(\pi_{+1})}{\partial \pi(i)}$ is not. As we will discuss below, immunization externalities are important at the onset of the pandemic. However, these immunization externalities diminish as the pandemic progresses, so eventually the social marginal cost of infections becomes large enough to trigger the same Hammer and Dance sequencing as the equilibrium.

The path of policy during the Dance phase is dictated by the speed at which the epidemic

[17] This Hammer phase is not necessary if the pandemic starts from an initial share of infections below η.

[18] The Dance phase is not required if the initial population state has such a high initial infection rate $\pi_0(i)$ that the first phase is already sufficient to establish herd immunity.

Figure 3: Optimal Deconfinement (The Dance, Proposition 3)

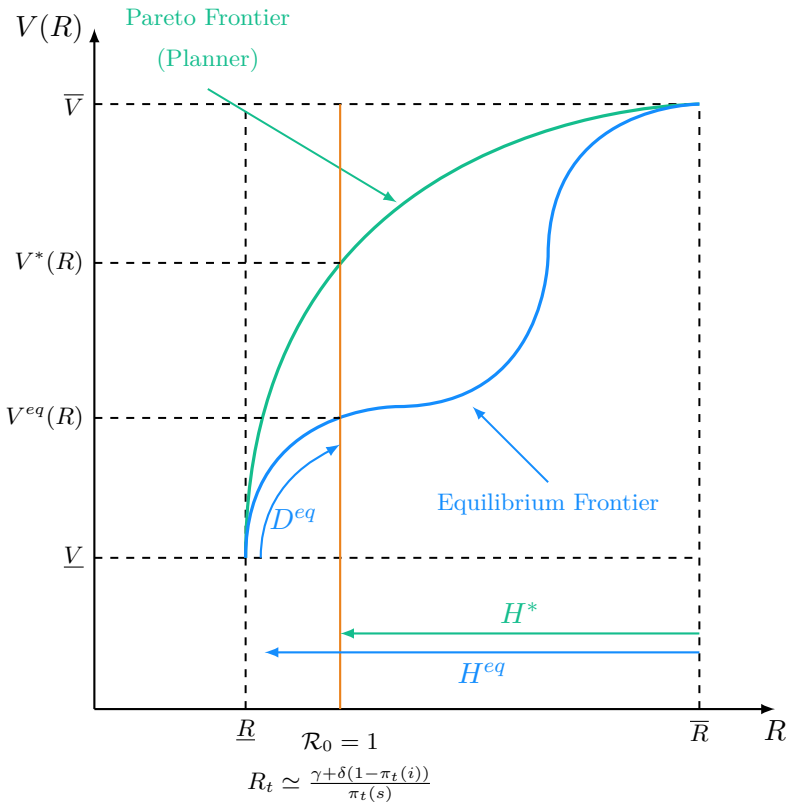

progresses towards herd immunity: any new burst of infections would immediately result in a renewed strengthening of social distancing measures to bring infections back down, while any decay in infections would result in incentives to relax social distancing measures - and thus bring the rate of new infections back up. These results illustrate how the shadow costs of infection risk stabilize the optimal and equilibrium path of policy. Since the equilibrium and optimal policy contain but do not suppress the infection, $\mathcal{R}_0(\pi_t)$ is kept close to 1. Proposition 3 thus offers a formal justification for the intuitive recommendation that optimal policy should focus on maximizing economic welfare $\mathcal{V}^*(R_t)$ while keeping the pandemic under control, treating $\mathcal{R}_0(\pi_t) \leq 1$ or $R_t \leq (\gamma+\delta)/\pi_t(s)$ as a de facto constraint on policy decisions (Budish, 2020).

2.3 The instantaneous propagation limit

So far, our theoretical results did not identify any major differences between equilibrium and optimal policy dynamics. We now provide a closed form characterization to the equilibrium and the planner's solution for the continuous time limit of the discrete time model with arbitrarily fast propagation. This allows us to discuss important differences between the two.

To distinguish between time discounting and the speed of propagation, let $\tau \equiv \Delta t$ denote *calendar time*, and let $\beta = e^{-\rho\Delta}$, for a fixed time discount rate ρ. We index all equilibrium variables by Δ, consider their limit in calendar time as $\Delta \to 0$, holding constant the infection,

recovery and death probabilities $R_t\pi_t(i)$, γ and δ per time interval Δ, and write their continuous-time limits as a function of calendar time τ.[19] In this limit, the infection has the potential to propagate instantaneously in calendar time.

Planner's Solution. As noted above, the planner could, in principle, opt for permanent restriction policies that bound $R^*(\cdot)$ permanently away from $\overline{R}$ to lower the long-run mortality rate. Define R^* as the long-run optimal policy that maximizes $\Lambda(R) \cdot \mathcal{V}^*(R)$, where $\Lambda(R)$ denotes the long-run survival rate associated with a constant policy R:

$$R^* = \arg\max_{R\in[\underline{R},\overline{R}]} \frac{\gamma R}{(\gamma+\delta)(R-\delta)}\mathcal{V}^*(R) \iff \frac{\mathcal{V}^{*\prime}(R^*)R^*}{\mathcal{V}^*(R^*)} = \frac{\delta}{R^*-\delta}$$

Proposition 4 shows that optimization of this long-run trade-off emerges as the solution to the planner's problem in the instantaneous propagation limit.

Proposition 4. ***Convergence to Long-run Optimum:*** *The planner's optimal policy converges to the long run optimum:* $\lim_{\Delta\to 0} R^*(\tau) = R^*$, $\lim_{\Delta\to 0}\pi(i,\tau) = 0$ *and* $\lim_{\Delta\to 0}\pi(s,\tau) = (\gamma+\delta)/R^*$ *for all* $\tau > 0$.

Proposition 4 shows that the long-run trade-off between mortality risk and economic distortions re-emerges in the instantaneous propagation limit. At instant $\tau = 0$, the social planner lets the pandemic progress and applies an instantaneous "Hammer" to immediately bring the pandemic to a level of infection and recovery associated with the long-run optimum (Shift H^* in Figure 3). This phase ends with $\pi(i,\tau)$ arbitrarily close to 0 and $\pi(s,\tau)$ arbitrarily close to $(\gamma+\delta)/R^*$ (they reach 0 and $(\gamma+\delta)/R^*$ at the limit when $\Delta \to 0$). For $\tau > 0$, the optimal policy then keeps the pandemic in a "never-ending Dance", in which policy is permanently kept at R^* to fully suppress further infections.

Why does the planner let the pandemic progress to the level associated with R^*, but no further? The planner controls the speed at which infections progress during the Dance phase. Since any new infections lead to quasi-instantaneous death or recovery, $\pi(s,\tau)$ and $\Lambda(\tau)$ immediately converge to the long-run values consistent with a given policy $R^*(\tau)$. But then, at each point in time, the planner faces the same quasi-static trade-off between economic distortion and survival probability, which has a static optimum at R^*. Hence, in the limit as $\Delta \to 0$, it must be optimal for the planner to stall the Dance phase *immediately* and *permanently* at the long-run optimal policy R^*. Nevertheless, the Hammer remains important at $\tau = 0$: after an initial propagation, the planner applies a quick but powerful hammer to bring $R^*(\tau)$ to its long-run level R^* *from below*: If instead the policy was set to $R^*(\tau) = R^*$ from the start, the epidemic will overshoot the long-run optimum, so that $\pi(s,\tau) < (\gamma+\delta)/R^*$ for $\tau > 0$.

[19] In this section, we will denote $\pi(i,\tau)$ (resp. $\pi(s,\tau)$) the share of infected (resp. susceptible) agents at time τ. Likewise, $R(\tau)$ will denote the infection risk at time τ.

The fact that the planner keeps policy permanently at the long-run optimum R^* illustrates another interesting aspect of optimal policy: the faster the propagation, the slower the optimal recovery. With discounting or finite speed of propagation, the trade-off between economic distortions and mortality risk is no longer instantaneous during the Dance phase, so the longer it takes today's infections to pass through to higher future mortality, the more the planner is willing to let the pandemic progress. Hence the actual speed of progression is inversely related to its potential speed of progression at the planner's solution: in response to fast propagation of the epidemic, the planner slows or stalls its long-run resolution into the very distant future.

Equilibrium Dynamics. Our next proposition shows that the equilibrium dynamics are radically different from the planner's solution. Let $\widetilde{R}(\tau)$ be defined by the solution to

$$e^{-\rho\tau} = \frac{\int_{\widetilde{R}(\tau)}^{\overline{R}} \frac{1}{R^2 \mathcal{V}_r^{eq}(R,R)} e^{\int_R^{\overline{R}} \frac{1}{R'-\delta} \frac{\delta \mathcal{V}^{eq}(R',R')}{\mathcal{V}_r^{eq}(R',R')R'^2} dR'} dR}{\int_{\widetilde{R}(0)}^{\overline{R}} \frac{1}{R^2 \mathcal{V}_r^{eq}(R,R)} e^{\int_R^{\overline{R}} \frac{1}{R'-\delta} \frac{\delta \mathcal{V}^{eq}(R',R')}{\mathcal{V}_r^{eq}(R',R')R'^2} dR'} dR},$$

for a given value of $\widetilde{R}(0) \geq \underline{R}$. $\widetilde{R}(\tau)$ is continuous and increasing with $\lim_{\tau\to 0} \widetilde{R}(\tau) = \widetilde{R}(0)$ and $\lim_{\tau\to\infty} \widetilde{R}(\tau) = \overline{R}$.

Proposition 5. ***Equilibrium Convergence:*** *The Markov-Perfect equilibrium policy converges to* $\widetilde{R}(\tau)$*:* $\lim_{\Delta\to 0} R^{eq}(\tau) = \widetilde{R}(\tau)$*,* $\lim_{\Delta\to 0} \pi(i,\tau) = 0$ *and* $\lim_{\Delta\to 0} \pi(s,\tau) = (\gamma+\delta)/\widetilde{R}(\tau)$ *for all* $\tau > 0$*, where* $\widetilde{R}(0) \geq \gamma + \delta$ *is uniquely determined from the initial infection rate* $\pi(i,0) > 0$*.*

Proposition 5 characterizes the dynamic equilibrium at the instantaneous propagation limit. For $\tau > 0$, this characterization is derived from a pair of first-order ODEs, one that stems from the agents' Hamilton-Jacobi-Bellman equation and optimality condition for $R^{eq}(\tau)$, and one that summarizes the dynamics of $\pi(s,\tau)$ in continuous time. Combining these two ODEs yields a second-order ODE whose solution is given $\widetilde{R}(\tau)$. Importantly, aside from an initial jump coming from the mass of initial infections, the equilibrium does not display any discontinuities, but rather a continuous, gradual progression towards herd immunity: At any point in time, agents trade off the marginal instantaneous utility gains and increases in infection and mortality risks. Any discrete jump in infection rates and mortality would generate a very strong hold-out motive that offsets such a possibility. At the same time, the pandemic cannot stall before reaching herd immunity because as soon as infection risks subside, agents will respond through a higher choice of R. There is thus a unique equilibrium path which consists of an immediate "Hammer" to bring the initial infection rate $\pi(i,0)$ under control, followed by gradual convergence towards a state of herd immunity.

The comparison with the planner's solution at the instantaneous propagation limit thus reveals two major sources of inefficiency in the equilibrium response to the pandemic: First, the planner's solution is actually far more permissive than the equilibrium policy at the onset,

allowing for an immediate large jump in infections to reach the long-run optimum. Second the planner's solution is more restrictive than the equilibrium in the long-run, which implies that long-run mortality is higher at the equilibrium than at the planner's solution. These results suggest that dynamic infection externalities dominate during the recovery, while immunization externalities dominate at the onset of the pandemic.

Externalities. Recall that for any population state π, the ratio of planner´s vs. equilibrium shadow costs Φ^*/Φ^{eq} provides a measure of dynamic externalities. Proposition 6 formalizes the preceding observations about immunization vs. infection externalities at the instantaneous propagation limit.

Proposition 6. ***Externalities at the instantaneous propagation limit:***

(i) Along the planner's solution, $\lim_{\Delta\to 0}\Phi^*(\tau)>0$ *but* $\lim_{\Delta\to 0}\Phi^{eq}(\tau)=0$ *for all* $\tau>0$.

(ii) Along the Markov-perfect equilibrium, $\lim_{\Delta\to 0}\Phi^{eq}(\tau)>0$, *but* $\lim_{\Delta\to 0}\Phi^*(\tau)=0$ *for all* $\tau\in(0,\tau')$, *such that* $\widetilde{R}(\tau')=R^*$.

Part (i) of Proposition 6 shows that at the long-run optimum, the social marginal cost of infection far exceeds the private cost. To prove this result, note that $\Phi^*(\tau)=\mathcal{V}^{*\prime}(R^*)>0$, but $\Phi^{eq}(\tau)\propto\pi(s)\pi(i)\beta/(1-\beta)$ and since the pandemic is stalled at the long-run optimum, it must be the case that $\lim_{\Delta\to 0}\pi(i)\beta/(1-\beta)=0$. Part (ii) of Proposition 6 shows that along the Markov-perfect equilibrium, the private costs of additional infections exceed the social costs early on. To see this, note that $\Phi^{eq}(\tau)=\mathcal{V}_r^{eq}(R^{eq}(\tau),R^{eq}(\tau))>0$, while $\Phi^*(\tau)\propto\frac{\partial v^*(\pi_{+1})}{\partial\pi(s)}-\frac{\partial v^*(\pi_{+1})}{\partial\pi(i)}$. But since the planner's solution immediately converges to the long-run optimum, as long as the latter remains attainable, i.e. whenever $\pi(s)>(\gamma+\delta)/R^*$ or $R^{eq}(\tau)<R^*$, $v^*(\pi)$ doesn't vary with π and $\frac{\partial v^*(\pi_{+1})}{\partial\pi(s)}-\frac{\partial v^*(\pi_{+1})}{\partial\pi(i)}\to 0$.

These results are best understood by comparing the aggregate consequences of each additional infection in the initial and final stages of the pandemic. During the final stages, the residual risk of infection is low. At the same time the equilibrium policy keeps $\mathcal{R}_0$ close to 1, hence any new infection is also at the start of an arbitrarily long future infection chain. The social marginal cost of an additional infection thus corresponds to the cost of infecting the entire subsequent infection chain, since preventing the first infection would have avoided all the subsequent ones along the chain with a very high probability. By not internalizing the costs of these follow-up infections, the equilibrium is less restrictive than the planner's solution during the recovery phase, resulting in a higher level of mortality than optimal.

More specifically, if the basic reproduction rate of the virus $\mathcal{R}_0$ equals 1, then the long-run optimum is unstable: a fraction $\pi(s)\varepsilon$ of additional infections generates $\pi(s)\varepsilon(1-\varepsilon)$ follow-up infections, which in turn trigger $\pi(s)\varepsilon(1-\varepsilon)^2$ infections of their own, and so on.[20] Any

[20]The factor $1-\varepsilon$ appears because the initial perturbation reduces the fraction of infected agents by a factor $1-\varepsilon$, and hence the $\mathcal{R}_0$ from 1 to $1-\varepsilon$.

infinitesimal perturbation $\pi(s)\varepsilon$ of additional infections therefore generates a large discrete mass proportional to $\pi(s)$ of new infections and deaths, which moreover occur instantaneously at the instantaneous propagation limit. Therefore, while the private cost of infection is bounded and equal to $\frac{\delta}{\gamma+\delta}\mathcal{V}^*(R^*)$, the product of agents' life-time utility $\mathcal{V}^*(R^*)$ with the probability that infection results in death, the social cost of an infection late in the pandemic tend to be infinitely larger than the private costs.

In contrast, the social costs of an additional infection are much lower early on in the pandemic. On the one hand, the planner internalizes that a large fraction of the population must eventually suffer an infection, and therefore avoiding infections early on at best postpones them, i.e. the expected utility gains from preventing early infections remain temporary. On the other hand, the planner internalizes that allowing infections early on immunizes the survivors and potentially cuts other infection chains later on. Not surprisingly, no agent is privately eager to participate in this effort to build herd immunity by exposing themselves to an infection. Agents thus have a much stronger private motive to hold out early on during the pandemic until the worst is over: hence immunization externalities are much stronger, and infection externalities weaker at the onset of the pandemic. This explains why in our benchmark model the planner favors a rapid progression towards the long-run optimum, while agents at equilibrium have a strong motive to hold out. But if everyone engages in such behavior, social distancing becomes self-defeating and the pandemic is slowed to a point where the recovery takes much longer, amplifying the economic costs of the pandemic crisis.

Importantly, the strength of immunization externalities relies on the assumption that (i) the initial peak of infections has no additional adverse effects, for example through medical congestion, and (ii) building herd immunity is the only viable long-term exit strategy from the pandemic. As we will discuss in the numerical section below, our analytical solution of the instantaneous propagation limit offers a benchmark that we can use to factor in these additional elements. Our characterization results still remain useful for these extensions, since all of them can be interpreted as natural modifications of the comparison between private and social costs of infections in first-order conditions (4) and (5).

2.4 Numerical Illustration

Figures 4 and 5 illustrate the economic benefits of flattening the curve and the associated economic decisions alongside with the reproduction rate $\mathcal{R}_0$.[21]

During the first 15 weeks of the epidemic, the optimal and equilibrium dynamics consist

[21] These dynamics were obtained assuming specific functional forms and parameters detailed in Appendix B. We focus on a special case in which the static equilibrium implementation is efficient and there are no static externalities ($\mathcal{V}_R^{eq}(R,R)=0$). Figure 15 in Appendix C reports the same dynamics with a higher mortality rate where the results are even more striking.

Figure 4: Flattening the Curve: Population Dynamics

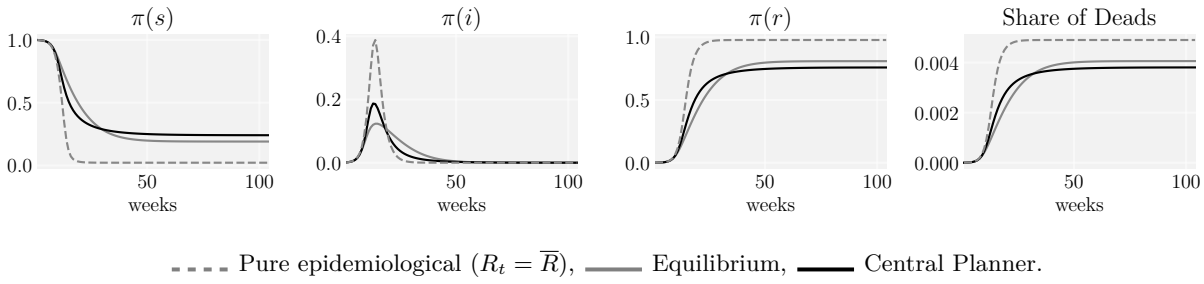

of flattening the curve of infections. Relative to the pure epidemiological model that ignores endogenous economic interactions, the equilibrium and optimal policies substantially dampen the rate of infection early on and bring the initial peak of infections to a lower level —the equilibrium more so than the optimal policy. Social distancing lowers the basic reproduction rate $\mathcal{R}_0$ relative to the pure S-I-R model during this period, resulting in a peak infection rate $\pi(i)$ of infected agents that is less than half as high as in the pure epidemiological model. This reduces the long-run rate of mortality to near the minimum level necessary to establish herd immunity.

Figure 5: Flattening the Curve: Policies

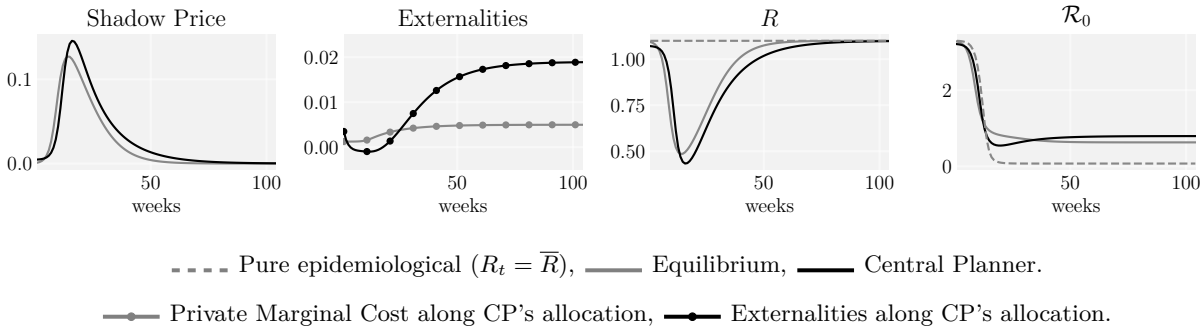

The first panel of Figure 5 shows that social and private incentives are roughly aligned during the initial curve-flattening phase: the optimal and equilibrium shadow costs of infection rise in tandem. However, the planner's solution is less restrictive early in the course of the pandemic than the equilibrium. Indeed, the planner's shadow cost is slightly lower, implying that the infection rate peaks at a higher level and the $\mathcal{R}_0$ drops slightly more slowly, than in equilibrium.

The second phase of the epidemic corresponds to the controlled deconfinement period, during which the shadow cost of infections slowly decreases. During this phase, the shadow cost is significantly higher for the central planner, implying that the socially optimal $\mathcal{R}_0$ continues its sharp drop below 1 before reaching its long-run value from below. Even though the planner's solution has already seen a higher fraction of infections, and thus fewer agents remaining susceptible, the higher shadow cost (first panel of figure 5) implies that social distancing measures

are more stringent at the planner's solution and kept in place for longer (third panel). The equilibrium recovery thus starts from a higher level of susceptible agents, but by relaxing social distancing measures faster than optimal, it results in a more gradual decline of infections during the recovery, and a higher level of long-run mortality.

This persistent difference between the competitive equilibrium and central planner allocation can be best understood in light of the offsetting spillovers condition stated in Proposition 1. With a "hump-shaped" policy that is the natural response to the pandemic at both the equilibrium and the planner's solution (see Figure 5), this offsetting spill-overs condition can hold only if dynamic spill-overs at a given value $R = R^*(\pi)$ are the same at the onset of the pandemic and during the recovery phase. This cannot happen with the evolution of dynamic spill-overs reported in the second Panel of Figure 5, because the relative importance of immunization spill-overs decreases and the relative importance of infection spill-overs increases as the pandemic progresses. The direct private cost of a new infection (gray marked curve) is always positive and increasing. In the long-run, this direct cost is dwarfed by the additional cost created by the dynamic infection externalities (black marked curve): individuals do not internalize the effect of being infected today on future infection risks for other agents. As a result, the total marginal cost of an additional infection is about five times as large as the direct cost faced by the agents. In the short-run, however, the dynamic externalities are reversed: they are quickly decreasing and become negative. As a result, the social marginal cost of a new infection is temporarily smaller than the private cost – it is, in fact, close to zero around week 10. This illustrates the immunization externality discussed in Proposition 6: Agents do not internalize that being infected today brings the economy closer to long-run herd immunity. The optimal path lets the epidemic spread faster than in equilibrium in the very early stages, which reduces economic costs during the recovery without raising long-run mortality. Once the pandemic has immunized a sufficient number of agents, the optimal policy shifts towards controlling further infections to keep long-run mortality under control, while the economy fully recovers.

At equilibrium instead, agents respond to the onset of the pandemic with strong voluntary confinement to "wait out the storm". But this results in a hold-out externality that has the nature of a zero sum game: If everyone waits out the storm, then the pandemic progresses very slowly, infections take longer to materialize, and agents stay locked up for longer than necessary. In addition, once herd immunity builds up and agents gradually exit their confinement, they do not internalize infection externalities. Therefore the long-run mortality at equilibrium eventually exceeds mortality at the planner's solution, even though it was way lower early on.

Figure 6 shows that both utility (relative to $\overline{V}$) and consumption[22] both suffer a sharp drop during the curve-flattening phase due to the economic lockdown, followed by a slow recovery during the deconfinement. In the long-run, the economy converges back to the optimum with

[22]Given the utility function (see Appendix B), a measure of consumption is simply given by $(\mathcal{U} - \underline{V})/(\overline{V} - \underline{V})$.

utility $\bar{V}$. Mirroring the discussion in the previous paragraph, the fall in consumption optimally chosen by the central planner is initially more gradual, but it reaches a lower depth and the deconfinement starts later, than in equilibrium. Our simulation suggests a ca. 20% drop in consumption at equilibrium. The planner's solution reaches its trough slightly later with a 24% consumption drop.

Figure 6: Utility and Consumption

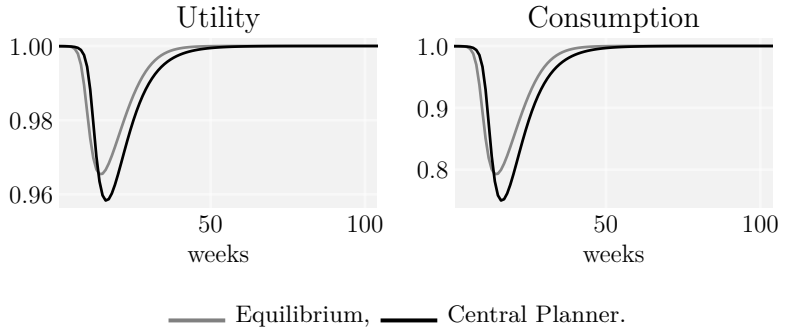

One corollary of these results is that policy delays have only minimal welfare costs and no long-term consequences in the beginning, provided that the planner can compensate for the delay through a stronger "catch-up" intervention without over-shooting long-run herd immunity.[23]

Figure 7: Instantaneous Propagation Limit (Proposition 3)

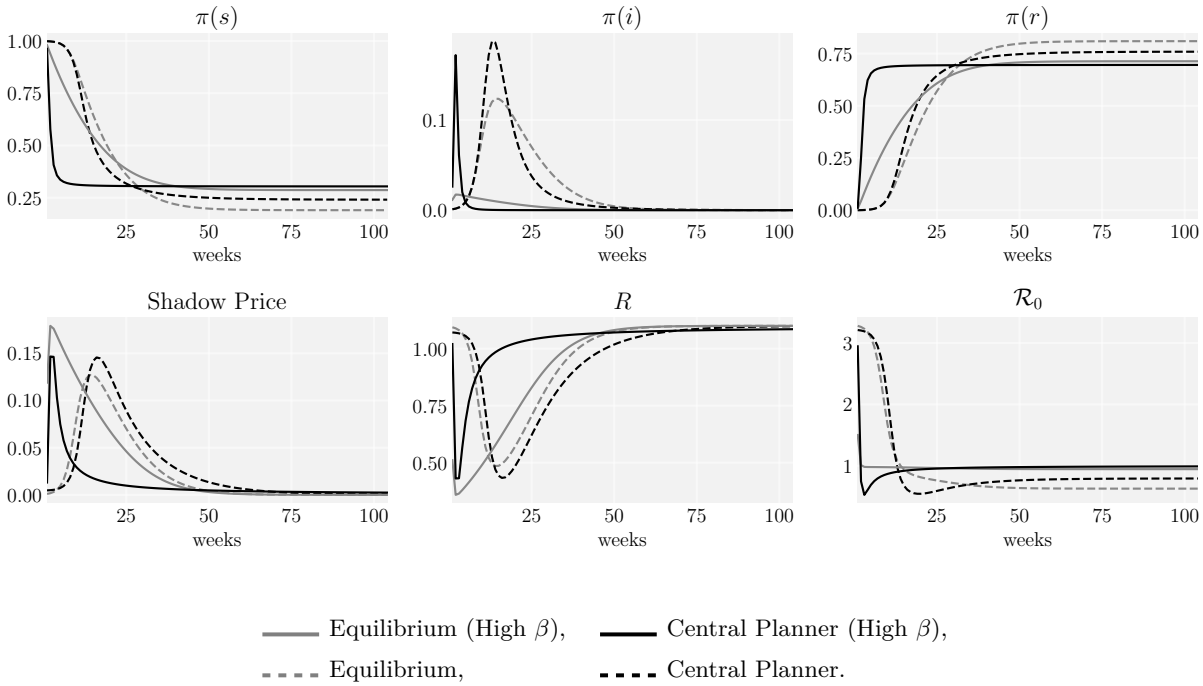

Figure 7 illustrates the instantaneous propagation limit (propositions 3, 4 and 5). We

[23]We present the corresponding simulation in the online appendix. The numerical results are consistent with the observation in Prop. 6 that the social cost of infection converges to 0 at the instantaneous propagation limit, as long as the planner's long-run optimum is attainable.

compare the benchmark equilibrium and optimal paths reported in Figure 4 and 5 to the case where the pandemic propagates seven times faster: each infection and recovery takes place over a few days, rather than a few weeks, but the results are reported on the same calendar time scale as the benchmark.[24] The pandemic peaks much faster at the planner's solution and is quickly stabilized at a level very close to the long-run optimum: the basic reproduction rate $\mathcal{R}_0$ temporarily drops well below 1 to then recover quickly to 1, letting the pandemic continue its very gradual progression towards herd immunity. The equilibrium instead slows down new infections from the beginning, at a rate that just offsets the increase in propagation speed. We clearly see the far slower speed of convergence, the associated higher economic costs (measured by the shadow cost of infection risks which equals the marginal disutility of social distancing), and the higher long-run mortality.[25]

3 Extensions

We explore six extensions of our baseline model: static externalities, medical sector congestion, transitory immunization, potential development of a vaccine, the use of face masks, and testing and tracing.[26] The first four extensions keep that static equilibrium and Pareto frontiers unchanged but alter the balance of static and dynamic externalities in the planner's and agents' dynamic first-order condition. The last two extensions introduce policy tools to manage the static efficiency frontiers but keep the dynamic first-order conditions unchanged.

3.1 Static and Dynamic Spill-overs

Static Spill-overs: We now explore how the presence of static (infection and economic) externalities affects the dynamics. These static externalities introduce a wedge between the private and social marginal disutility of social distancing, i.e. $\mathcal{V}_R^{eq}(R,R) \neq 0$. In order to organize the discussion, let us consider the benchmark utility function we used in the simulations of Section 2.4

$$\left(\frac{\mathcal{V}^{eq}(r,R)-\underline{V}}{\bar{V}-\underline{V}}\right)^2+\alpha\left(\frac{\bar{R}-r}{\bar{R}-\underline{R}}\right)^2+(1-\alpha)\left(\frac{\bar{R}-R}{\bar{R}-\underline{R}}\right)^2=1$$

where α is a parameter that controls for the strength and type of static externalities at play in the confinement game. The case $\alpha=1$ corresponds to the situation where static externalities offset

[24] A higher mortality rate magnifies these results. See Figure 16 in Appendix C.

[25] The role of the long-run optimum as a point of convergence for the planner's solution is less easy to see from these figures but can be summarized with a few numbers. In our numerical example, herd immunity is reached when $\pi(s)=0.303$. The long-run optimum is attained when $\pi(s)=0.314$. At the benchmark calibration, $\pi(s)=0.339$ after 40 weeks and $\pi(s)=0.305$ after 100 weeks at the planner's solution. With faster propagation instead, $\pi(s)=0.316$ after 20 weeks, and progresses only to $\pi(s)=0.313$ after 100 weeks. $\pi(s)$ thus stays close to the long run optimum for a very long time.

[26] The interested reader is left to refer to the online appendix for analytical details pertaining to extensions.

each other and $\mathcal{V}_R^{eq}(R,R) = 0$, private and social marginal costs of social distancing coincide. When $\alpha > 1$, agents do not internalize that their actions contemporaneously increase infection risks for others, and private marginal costs of social distancing exceed social marginal costs ($\mathcal{V}_R^{eq}(R,R) < 0$). When $\alpha < 1$, agents do not internalize the adverse static utility consequences of their actions for other agents, so social marginal costs exceed private marginal costs of social distancing ($\mathcal{V}_R^{eq}(R,R) > 0$).

Figure 8 reports the equilibrium dynamics obtained when economic externality dominates ($\alpha = 0.5$) and when infection risks externality does ($\alpha = 1.5$), and compares them to the planner's solution which remains unchanged, and to the equilibrium without static externalities ($\alpha = 1$).

Figure 8: Varying α

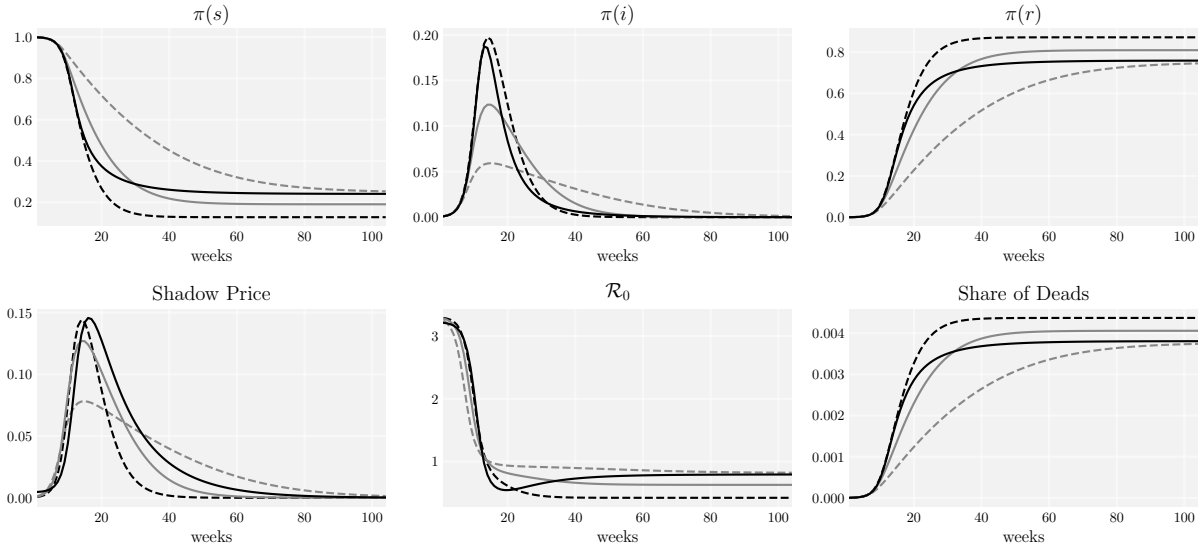

When the infection risks externality dominates ($\alpha = 1.5$), static and dynamic spill-overs are offsetting at the beginning, but reinforcing during the recovery: the fact that agents do not fully internalize static infection risks partly offsets the dynamic immunization externality, resulting in an equilibrium policy that tracks the first-best more closely until reaching the infection peak. However, past the peak, the static infection risk externality reinforces the dynamic infection externality, resulting in even faster deconfinement and a higher level of long-run mortality.

The economic externality ($\alpha = 0.5$) instead reinforces the immunization externality at the onset, resulting in a lower peak infection rate and a slower path of deconfinement. Offsetting spill-overs instead appear in the long run, where static economic externalities offset the dynamic infection externality, reducing the speed of recovery and long-run mortality.

Medical Sector Congestion: We now introduce a congestion externality in the medical sector by letting the death rate be increasing in the fraction of infected agents.[27] This congestion therefore plays through the S-I-R dynamics and affects the dynamic evolution of shadow costs. It therefore acts as a dynamic externality.

The results[28] are presented in Figure 9.[29] Compared to the baseline model, the shadow

Figure 9: Congestion Effects

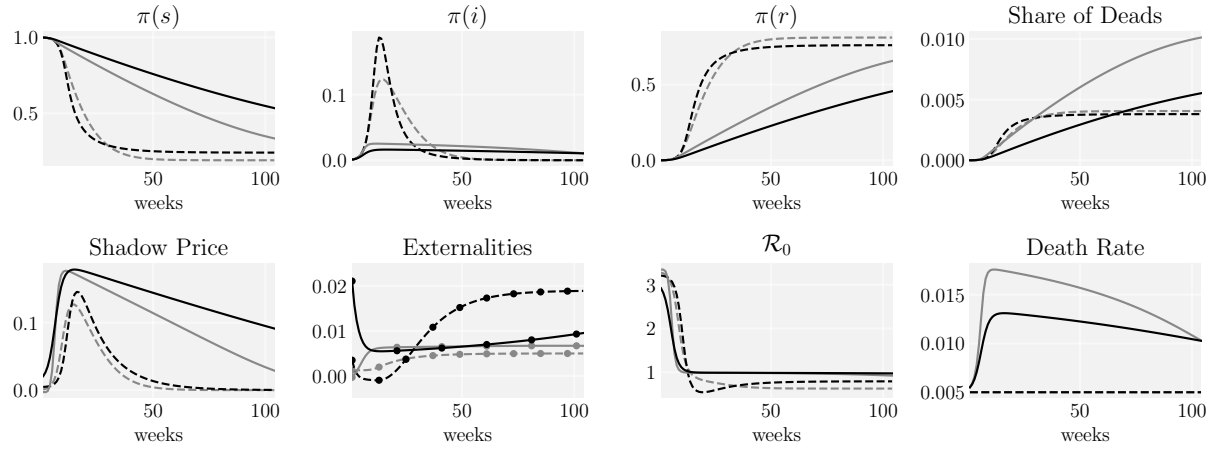

Central Planner: Benchmark, Congestion; **Equilibrium:** Benchmark, Congestion. **Private Marginal Cost:** Benchmark, Congestion, **Externalities:** Benchmark, Congestion.

price of infections rises to a much higher level, and the infection rate peaks at a much smaller level. This is the case both in the equilibrium and for the central planner, but for two different reasons. In the central planner's optimum, the infection externality is now far more costly, as an additional infection today raises the future death rate more than linearly due to the congestion externality. The second lower panel of Figure 9 shows clearly that the dynamic externality is initially much larger than in our baseline model, indicating that the infection externality strongly dominates the herd-immunity externality even in the early stages of the epidemic. As a result, the total marginal cost of an additional infection remains large and positive throughout the duration of the epidemic. Because of this very high infection externality, the social planner implements a much stronger initial "hammer" phase, with a much lower $\mathcal{R}_0$ at time 0, than in the equilibrium and in the baseline model. Since this effect is not internalized by private agents, the rate of infections and, in turn, the death rate grow much faster in equilibrium than along

[27] In this and all subsequent cases, we go back to a setting where static externalities offset each other ($\alpha = 1$).

[28] For our simulations we assume that the conditional death rate $\delta(\cdot)$ takes the form

$$\delta(\pi(i)) = \underline{\delta} + \exp(\varphi\pi(i)) - 1,$$

where $\underline{\delta}$ corresponds to the conditional death rate that prevails in the model without congestion. In the spirit of Piguillem and Shi (2020), we set parameter φ so that when the economy reaches an infection rate of 1% the unconditional death rate in the economy doubles. This leads to the value $\varphi = 0.1682$.

[29] Figure 17 in the Appendix show the same graphs over a longer horizon.

the optimal allocation. It is for this reason (rather than because of the infection externality) that the private shadow price of infections shoots up and catches up with that of the central planner. But the subsequent reduction in economic activity occurs too late: the share of deaths rises much faster in equilibrium and reaches a much higher long-run level than in the central planner's solution (see Figure 17 in the Appendix).

Adding medical sector congestion highlights the importance of the lower-bound condition on the social marginal costs of infection risks: Proposition 3 shows that bounding this social marginal costs away from 0 is a necessary condition for immediate, strong policy interventions to be optimal. This principle extends to the model with medical congestion, but here we see that medical congestion amplifies the social costs of infection due to its impact on mortality.[30] The congestion externality thus works against the immunization externality in the short run, and it strengthens the infection externality in the long run, resulting in slower long-run convergence.

Transitory Immunization: Our baseline model assumes that, once recovered, agents have acquired permanent immunity. This assumption is however controversial and largely debated in the epidemiological literature (see World Health Organization (2020), and Kissler et al. (2020) among others). We therefore relax this assumption and consider the case where a recovered agent can, with probability ν, return to the pool of susceptible agents and be infected again in the future. This therefore leaves the possibility that the disease reoccurs with a non zero probability in the future. These dynamics are depicted in Figure 10, where immunity is expected to last for 2 years.[31]

Figure 10: Transitory Immunity

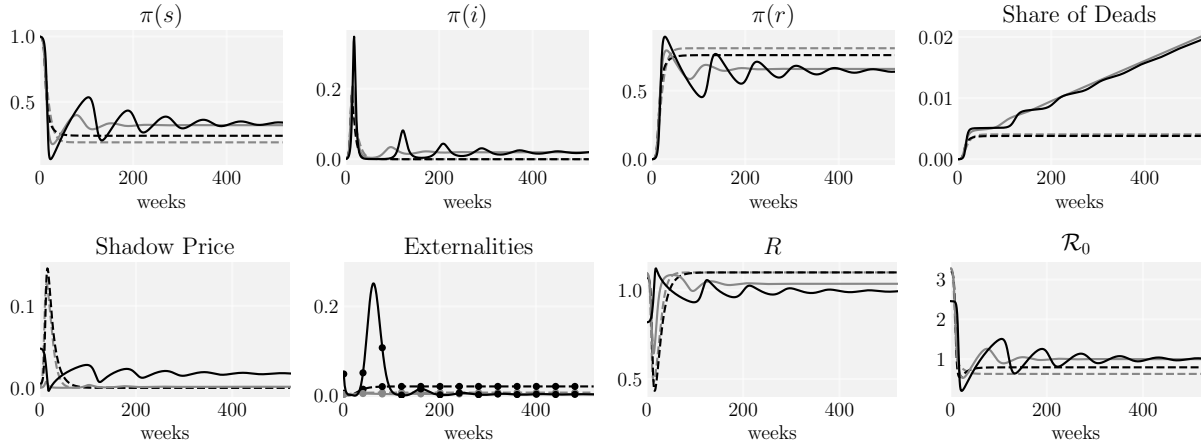

Central Planner: Permanent, Transitory; **Equilibrium:** Permanent, Transitory. **Private Marginal Cost:** Permanent, Transitory, **Externalities:** Permanent, Transitory.

[30] Medical congestion keeps the planner's and agent's first-order conditions unchanged, but alters the characterization of dynamic spill-overs.

[31] This amounts to a probability ν of 1%.

The dynamics displays recurring waves of infections, which eventually converge to a steady-state where infections are endemic. This experiment exemplifies the key difference between the competitive equilibrium and the central planner's allocation: In an equilibrium, agents take future aggregate dynamics as given and hence favor confinement when infection risks are rising and deconfinement when infection risks are falling. The central planner internalizes these aggregate dynamics, as reflected in the evolution of the shadow cost of infections: While the equilibrium shadow price only focuses on the current state of infections, tracking $\pi_t(i)$ very closely, the planner's shadow price is "forward-looking" and tracks $\pi_t(s)$ more closely as immunization gradually erodes between successive infection waves. Because the central planner understands that infections are endemic, she assigns a higher cost to infection risks than the individual agent along the dynamics, resulting in stronger permanent social distancing measures as the equilibrium. Consequently, the death toll, although it keeps increasing over time, is lower in the central planner's allocation than in the competitive equilibrium.

Hoping for Vaccines or Cures: We now assume that, in each period, there is a positive probability that a vaccine be discovered and, consequently, all the susceptible individuals be moved immediately to the recovered state. We set this probability such that a vaccine is expected to arrive after one year.[32] This vaccine comes too late for agents who are already infected. Like the medical congestion model, vaccines primarily alter the computation of shadow costs and dynamic externalities, while the baseline first-order conditions only changes marginally, discounting $\Phi^*(\pi)$ and $\Phi^{eq}(\pi)$ by a factor $1-\xi$, where ξ denotes the weekly probability of discovering a vaccine. The results are presented in Figure 11.

Figure 11: Possibility of a Vaccine

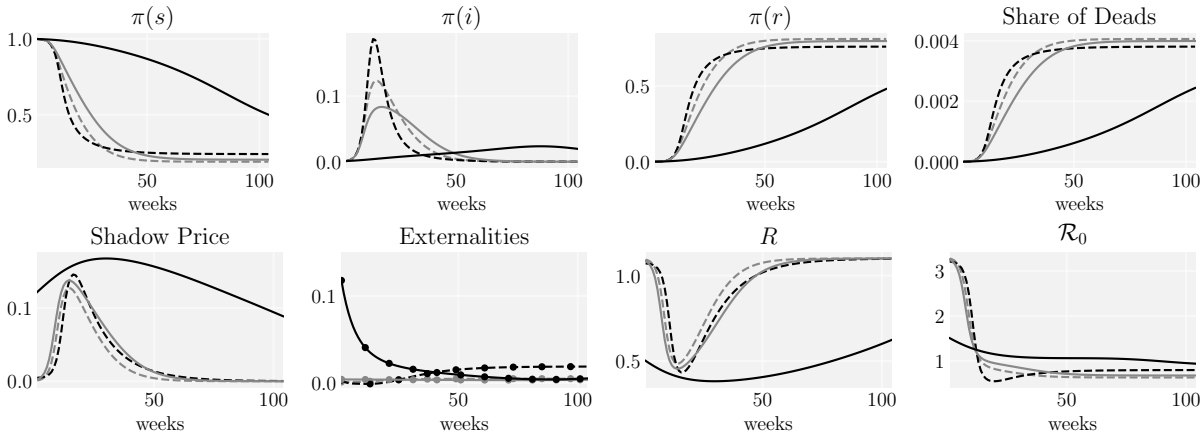

Central Planner: Benchmark, Vaccine; **Equilibrium:** Benchmark, Vaccine. **Private Marginal Cost:** Benchmark, Vaccine, **Externalities:** Benchmark, Vaccine.

[32] This leads to set ξ to $1/52 = 0.0192$. Note that a cure would have a similar effect in the model by instantaneously lowering the death rate for infected agents, so this parameter can be viewed as the overall arrival rate for a "game-changing" long-term exit strategy.

The dynamics of equilibrium resemble those of the baseline calibration. Individuals understand that a vaccine may be found, so they are even more willing to hold out and reduce economic activity in the short run, but the private benefits of a vaccine are too remote to significantly change their behavior.

The planner however follows a markedly different path relative to the baseline dynamics of Figures 4 and 5. As in the case of medical sector congestion, the shadow price of infection risks is initially much higher than in equilibrium, leading to a very strong "hammer" phase at time 0 that immediately brings $\mathcal{R}_0$ down and saving as many lives as possible until the discovery of a vaccine. Doing so allows the planner to delay the peak of infections, as shown by the dynamics of $\pi(i)$, in the hope that a vaccine is discovered before herd immunity is reached. Correspondingly, the externalities are bounded away from zero (see second lower panel of Figure 11): the prospect of a vaccine mutes the immunization externality and gives the planner a better long-term perspective than aiming for herd immunity in the early stages of the pandemic.

Because the equilibrium doesn't internalize this value of delaying the pandemic to develop a vaccine, it reaches a much higher level of long-run mortality than the planner's solution.

As a robustness check, we report the effects of a vaccine that is available at a 2-year horizon in the appendix (see Figure 18). The effects turn out to be much smaller than with a 1-year horizon: the planner is willing to delay the peak infection rate merely by a few weeks, before "giving up hope" for an early vaccine and letting the pandemic run its course towards herd immunity. This highlights that the horizon at which cures or vaccines are expected to be available is really important to determine whether the optimal policy response in the short run should focus on acquiring herd immunity or holding out for a better exit strategy.

3.2 Managing the Pareto Frontier

Face masks: Here we add the use of face masks to the set of static decision variables. Wearing a face mask confers no direct utility or disutility but reduces an agent's infection risk by a factor $f(m, M) \in (0, 1]$, where m denotes the agent's own use of masks, and M denotes aggregate mask usage. We assume that $f(0, 0) = 1$, $f_m(m, M) + f_M(m, M) \leq f_m(m, M) \leq 0$, with individual and aggregate decreasing returns. Masks are competitively supplied at price $P(M) = C'(M)$, where $C(M)$ displays decreasing returns to scale, and $\lim_{M \to 0} C'(M) = 0$.[33]

Analysis of the static equilibrium and optimality conditions for mask usage already generates several interesting insights:[34]

1. Mask usage yields an important static substitution effect: By flattening the utility-infection risk trade-off, face masks allow agents and the planner to relax social distancing

[33] In the simulations, we will use $f(m, M) = e^{-3m-3M}$ and $C(M) = M^2/2$.

[34] A detailed discussion of the planner's problem and equilibrium characterization with face masks is presented in the online appendix.

measures. This substitution effect becomes stronger the smaller is the shadow cost of infection risks.

2. Mask usage may be inefficiently low and require a proportional subsidy or regulatory measures to offset the mask use externality f_M.

3. We can represent the equilibrium shadow cost of infection risks $\Phi^{eq}(\cdot)$ as a function of the price $P(M)$ and usage M of face masks in equilibrium, and therefore obtain a market-based indicator for concurrent shadow costs of infection risks.

Figure 12 illustrates how the introduction of face masks changes the planner's and equilibrium frontiers. The use of face masks expands and flattens the set of attainable payoff in any given period. The planner's Pareto frontier is thus strictly larger, and allows agents to reach higher welfare for a given level of infection risks. However, because of externalities, the new equilibrium frontier is not guaranteed to be strictly higher than the old one.

Figure 12: Introduction of Face Masks

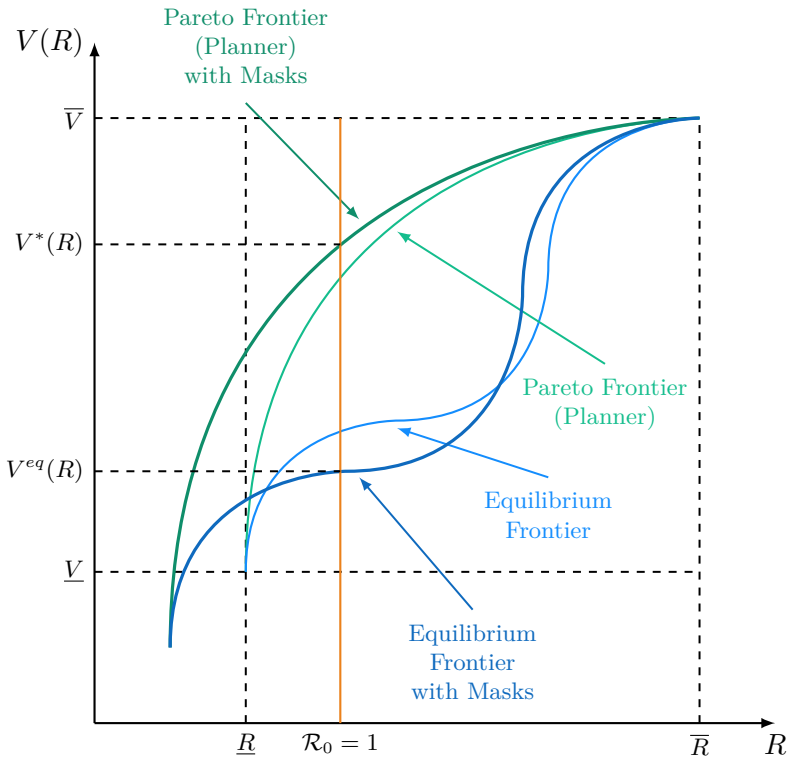

This figure already provides an intuitive understanding of how face masks change the dynamics of equilibrium and optimal policy. First, face masks give the planner and agents at equilibrium an option to push infection risks even below $\underline{R}$, thus resulting in yet faster control of the epidemic. Second, during the "Dance" phase, the use of masks serves to relax the Pareto frontier and achieve a higher level of utility for a given level of infection risks. Recall that the dynamics of infection risks during the deconfinement path are governed by a basic reproduction coefficient close to 1 (the red line in the Figure). Reducing infection risks through the use of

face masks during this phase allows the planner to increase R one for one. Therefore, face masks are a *short-run complement* to relaxing economic restrictions, since for a given state of epidemic progression, they allow for a higher level of economic activity. On the other hand, face masks do not improve on the long-run convergence towards a full recovery with herd immunity, since incentives for mask usage disappears once the economy approaches a complete recovery to $\overline{R}$.[35]

Figure 13: Masks

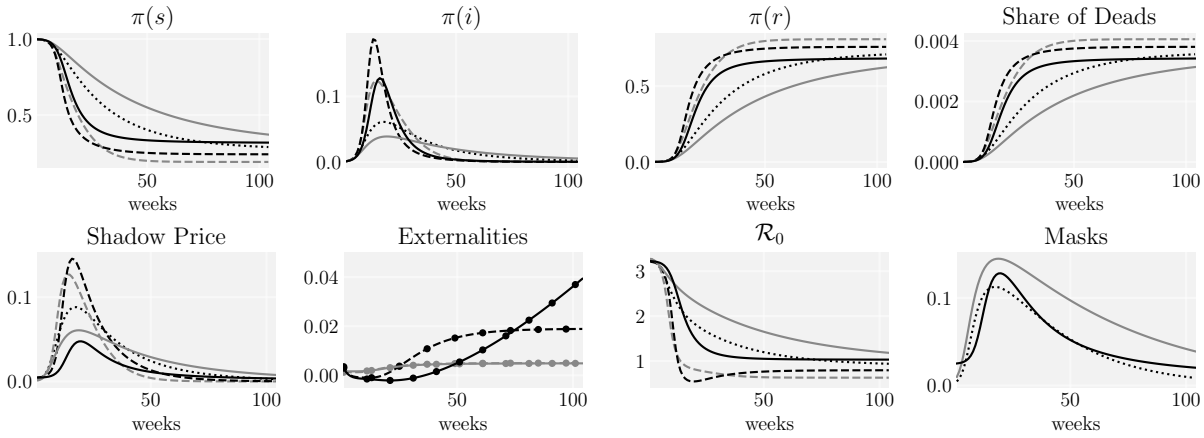

Figure 13 compares our benchmark to the model with face masks, with or without a subsidy correcting for the static externality associated with the use of masks. As shown in the shadow cost panel (lower left panel of Figure 13), usage of face masks significantly lowers the economic costs of the pandemic. In the planner's problem, face masks reduce the consumption trough from about 24% to 1%. The use of masks, rather than economic restrictions and costly social distancing measures, are thus used by the planner to contain the pandemic, both at its peak and during the recovery. The substitution effects towards higher economic activity are especially important during deconfinement, *i.e.* for a given bound on infection risks, face masks allow a deconfinement at a higher level of economic activity than at the benchmark. The planner's solution also features a lower peak infection rate and a lower level of long run mortality with face masks: long-term benefits of mask usage therefore balance economic and mortality costs.

In equilibrium, face masks also lower the peak infection rates as much or more than at the planner's solution, but the hold-out externality implies that it takes even longer to reach herd immunity. Therefore, while the peak economic costs are lower, the pandemic is longer-lasting, and has higher economic costs during the recovery. Interestingly, these delays are exacerbated by the subsidy for face mask usage: correcting the static externality may therefore worsen the dynamic spill-overs associated with face masks. What's more, the long-run infection

[35] Similar arguments also apply to the equilibrium, except that here the face masks may locally depress economic activity further if the new equilibrium frontier lies below the original one due to the importance of spill-overs.

externalities are amplified, resulting in a larger mortality gap between the planner's solution and the equilibrium (See second lower panel of Figure 13).

Testing and Contact-tracing: We also consider an extension of our model to testing and contact-tracing. By testing and quarantining anyone with a positive test result, one can reduce the number of undetected infections to $\pi(i)\,(1-\Pr(test|i))$, where $\Pr(test|i)$ denotes the fraction of infected agents that have had a positive test result and are thus in quarantine, which we interpret as a temporary exit from the game. This has a similar effect of improving the static efficiency frontier as face masks, as illustrated on Figure 14.

Figure 14: Frontiers with Testing

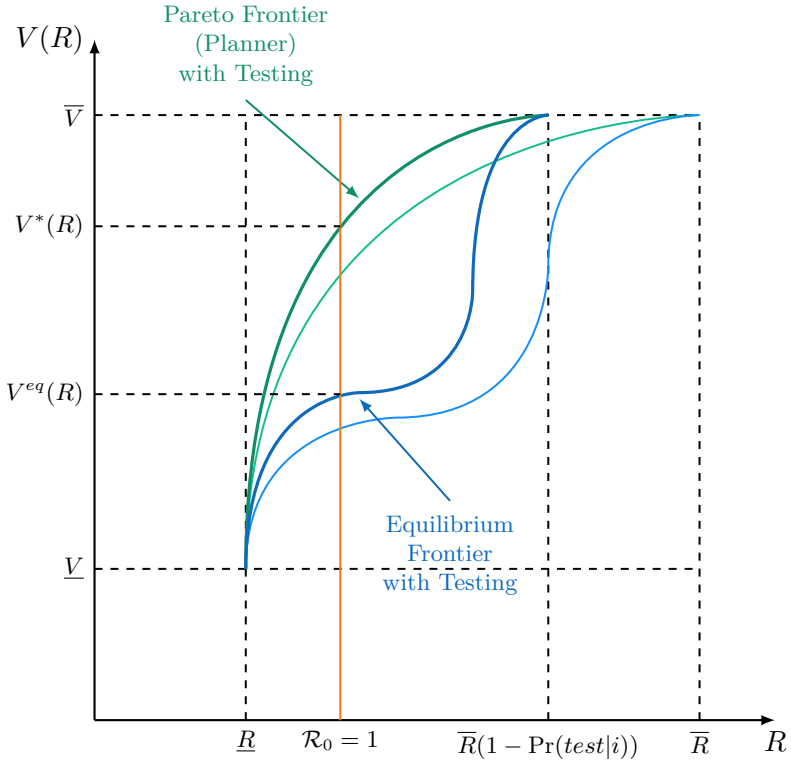

Like face masks, testing shifts the static and dynamic trade-offs between infection risks and instantaneous utilities and allows for strong substitution towards higher economic activity by offering better control over follow-up infections.[36] There is one major difference, however:

[36] Adding testing and quarantining into the model comes with two challenges. First, we need an additional state variable to keep track of the fraction of agents in quarantine. Second, testing alters agents' beliefs about their own health status, if they are informed of a negative test result. Hence, we need to keep track of heterogeneity across agents according to their test history. Berger, Herkenhoff, and Mongey (2020) show how to include additional state variables in an SIR model to capture the information generated through testing. Piguillem and Shi (2020) integrate such a structure into a simple dynamic planner's problem with capacity constraints in the medical sector, but focus on simple testing and quarantine policies. Eichenbaum, Rebelo, and Trabandt (2020b) extend their baseline model to allow for testing. Like us, these papers emphasize the potential for testing to relax untargeted

Testing lowers the threshold level of recoveries that is required to establish herd immunity, permanently eradicate the virus, and permanently eliminate economic restrictions. This is seen by the shift to the left of the Pareto and equilibrium frontiers for $\overline{V}$.

The benefits from testing are proportional to the fraction of infected agents that are quarantined, $\Pr(test|i)$. If tests are conducted purely randomly and with limited test capacity, then the fraction of positive test results is negligible when the fraction of infections vanishes. In the online appendix, we discuss how a systematic application of contact tracing from agents who tested positively allows to increase $\Pr(test|i)$ and quarantine a significant positive fraction of infections. Hence testing and contact-tracing offers long-term benefits in the form of a better long-term exit strategy.[37]

4 Concluding remarks

This paper provides a parsimonious framework to characterize the response to a pandemic crisis akin to the 2020 COVID-19 epidemic both in a *laissez-faire* equilibrium and at the central planner allocations. As our main take-away, we stress that health and economic objectives are not mutually exclusive, since controlling the pandemic is a necessary condition for restoring economic prosperity. In our model, fast propagation of the pandemic gives rise to the same 'Hammer-and-Dance' policy dynamics that were originally motivated by epidemiological considerations, such as protecting the medical sector or gaining time to develop a vaccine or cure: a relatively short period of a few weeks, during which rather extreme measures such as a strict lockdown of most sectors of activity are put in place to control the spread of the pandemic as quickly as possible (the *Hammer*) is followed by a longer recovery phase during which economic restrictions are gradually relaxed, while maintaining control of the pandemic at all times (the *Dance*). These policy dynamics arise from a simple cost-benefit comparison that equates the marginal economic costs of social distancing or lockdown measures to the marginal health benefits of avoiding infections. Fast speed of propagation amplifies the marginal health benefits of lockdown measures, which implies either that costly lockdown measures are implemented as happens during the Hammer, or that the infection rate is kept close to zero to offset the high speed of propagation, as happens during the Dance.

One important lesson we draw from the theoretical model is that the central planner is

quarantine measures. However, they do not combine testing with contact-tracing, which is key to maximize the containment potential from testing and quarantine policies. Our analysis in the appendix takes advantage of certain properties of the instantaneous propagation limit to side-step the technical challenges, but stops short of fully simulating the model like with the other applications and extensions.

[37] Pollinger (2020) shows that a combination of extended testing, tracing, and quarantines in combination with confinement can offer a fast exit from the pandemic, *i.e.* convergence to permanent containment without herd immunity, for any initial level of susceptible and infected agents.

forward-looking and fully internalizes the impact of current choices on the entire future dynamics. Agents at equilibrium instead act myopically, trading off concurrent utility and infection risks, but taking the aggregate future infection dynamics as given. Hence they do not internalize that risking an infection today may alter the future course of the pandemic, either by exposing others to higher future infection risks, or by contributing to long-run herd immunity. A robust feature of our analysis has been that infection externalities dominate during the recovery, and how well they are regulated is an important determinant for long-run mortality. In comparison, immunization externalities may be important in the short-run, resulting in private hold-out motives that slow down the recovery phase and impose excessively high short-run economic pain, unless they are offset by congestion in the medical sector or the value of delaying infections to find a cure or vaccine.

We made several simplifying assumptions to keep the model parsimonious and obtain stark results. For example, we assumed that the static equilibria of the economic and confinement stage games are efficient (assumptions 1 and 2). These assumptions imply that any rationale for policy interventions must arise from the trade-off between health and economic objectives. Most of the qualitative predictions of our model would not change, but the case for interventionist policies would only be strengthened without these assumptions.

The two characteristic features of the Covid-19 pandemic, fast propagation and asymptomatic transmission, were introduced by focusing on a limiting case of instantaneous propagation and through the assumption that agents do not know their own infection status (only death is observable). The symmetric information assumption allowed us to abstract from additional sources of heterogeneity related to demographic factors, different exposure risks by occupation or economic characteristics. By focusing on a static economic interaction, we are leaving aside dynamic economic adjustment processes and spill-overs e.g. through investment, labor market adjustment, dynamic sectoral adjustment or balance sheet effects. In addition, we abstracted from richer epidemiological models that would have allowed for symptomatic transmissions and managing the flow of patients. Introducing cross-sectional heterogeneity, dynamic economic interactions or a richer epidemiological model would require expanding the state space to account for additional heterogeneity or further dynamic linkages and thus increase the complexity of the model, but the basic logic of the hammer-and-dance result with fast propagation should hold as long as (i) for each agent at each moment, the equilibrium (optimal) decision equates a marginal disutility of social distancing to a private (social) shadow cost of infection risks, and (ii) this shadow cost of infection risks scales with the speed of propagation and the current infection rate.

By evaluating policies from *behind the veil of ignorance*, we abstracted from conflicts of interest between different types of agents. For example, if agents knew their health status, those who have recovered would like to return to the pre-pandemic equilibrium as fast as possible,

those who are infected also have nothing to gain from further confinement, but those who are still susceptible to infection would benefit from confinement policies. Likewise, introducing heterogeneity in the utility vs. mortality trade-off would also result in a conflict of interest over optimal policy, or call for "smart" policies that condition on individual characteristics.

We adopted a utilitarian welfare criterion that aligns private and social preferences over symmetric utility and mortality outcomes, but one could equally well adopt other welfare criteria. For instance, a Rawlsian criterion that maximizes the welfare of the worst-off agents might place a higher weight on survival probabilities. Strong myopia or hyperbolic discounting by policymakers or economic agents would have the opposite effects: it would place too strong a concern on immediate economic prosperity, especially when immediate infection risks are small. In general, a policymaker who does not weigh much mortality risk will relax social distancing measures, while one that over-weighs concerns about mortality risks will tighten them. One way to interpret such changes to the planner's objective function is that they introduce an additional externality between the planner's and equilibrium shadow costs of infection risks.

Finally, we have simplified the analysis by assuming common knowledge of the model parameters and perfect foresight about aggregate dynamics. Relaxing these assumptions would result in the additional need to infer the current state of the economy along with the optimal policy design. Since costs of straying from the optimal infection rate are asymmetric during deconfinement, we conjecture that optimal policy should stay on the side of caution, especially when fast propagation quickly results in a resurgence of new infections. This can be achieved by tracking the shadow cost of infection risks that summarizes the pressure to maintain social distancing measures in place as the pandemic progresses.

We conclude with a few remarks about optimal policy design. A direct implication of our main hammer-and-dance result is that equilibrium and optimal policy prioritize controlling new infections over safeguarding economic prosperity. Optimal policy can therefore be thought of as a standard welfare maximization problem subject to a constraint on the pandemic transmission rate, along the lines of Budish (2020). This constraint is inversely related to the current infection rate, resulting in very tight restrictions during the hammer phase while new infections are high, and then relaxing these restrictions during the recovery phase where the effective reproduction rate is kept below unity. This can be achieved by tracking the shadow cost of infection risks that encapsulates the trade-off between economic activity and infection risks and future mortality. This perspective on policy design suggests the following general policy guidelines:

1. *The Hammer:* early, decisive action is warranted if it helps save lives in the long run. If it merely delays infections in the short run, it just lengthens the recovery and inflicts higher economic costs on the population.

2. The social value of herd immunity depends on whether there exist better long-term alter-

natives (a cure or a vaccine, as envisaged in the preceding extensions), and at what time horizon.

3. *The Dance:* optimal deconfinement eases economic restrictions as much as possible while keeping the spread of the pandemic under control by keeping $\mathcal{R}_0$ below 1. In this context, it is important to manage static externalities so as to not make the pandemic more painful than necessary - equate private and social marginal costs of lockdowns across all sectors. Moreover, instruments that help control the spread of infections, like facemasks or testing-and-tracing policies, allow to relax the static trade-offs between economic activity and infection risks and thereby serve to ease economic restrictions.

4. Beware of infection externalities during deconfinement: *i)* do not count on private agents to fully understand the risks they pose to others; *ii)* do not pay for faster economic recovery with a higher mortality rate.

In many respects, these policy prescriptions are consistent with the ones suggested by epidemiologists. This comes as no surprise, as the dynamics of optimal policy are dictated by the dynamics of the epidemic, and an economic recovery is possible only if the pandemic is kept under control. Our analysis shows that these policy guidelines follow from simple economic principles equating the marginal costs and benefits of social distancing measures. Our results further provide an economic justification for using indicators like current infection rates or the measured $\mathcal{R}_0$ to guide regional or local policy makers, or mandating face masks or extensive testing-and-tracing policies as a complement to easing economic restrictions.[38] Health and economic objectives are thus not mutually exclusive.[39]

References

Alvarez, Fernando E, David Argente, and Francesco Lippi. (2020). "A Simple Planning Problem for COVID-19 Lockdown." Working Paper 26981, National Bureau of Economic Research. URL http://www.nber.org/papers/w26981.

[38] During the recovery, many countries, including France and Germany, have started to report weekly measures of $\mathcal{R}_0$, as well as weekly infection rates $\pi(i)$ at the regional or sub-regional level, and base policy adjustments on whether $\mathcal{R}_0$ is kept above or below 1 and whether $\pi(i)$ is below a threshold of 50 per 100000 inhabitants (0.05%) to re-instate tighter social distancing measures. This approach is very much consistent with our results.

[39] We first articulated many of these results and policy prescriptions in late April and early May of 2020, for example in Assenza et al. (2020a) or Assenza et al. (2020b), when most countries were still close to the peak of the pandemic and just beginning to enter the Dance phase. As of July 2020, we are struck by how much our results anticipated subsequent developments of the pandemic and policy responses, in particular the importance of managing dynamic infection externalities and avoiding a resurgence of cases during the recovery. We leave it to future work to evaluate through the lens of our model how well different countries managed to cope with the pandemic and its economic consequences.

Assenza, Tiziana, Fabrice Collard, Martial Dupaigne, Patrick Fève, Christian Hellwig, Sumudu Kankanamge, and Nicolas Werquin. (2020)a. "The Hammer and the Dance: Equilibrium and Optimal Policy during a Pandemic Crisis." Discussion Paper 14731, Centre for Economic Policy Research. URL https://cepr.org/active/publications/discussion_papers/dp.php?dpno=14731.

———. (2020)b. "The Hammer and the Dance: Health and Economic Objectives are not Mutually Exclusive after all!" VOX-eu column 14731, Centre for Economic Policy Research. URL https://voxeu.org/article/health-and-economic-objectives-are-not-mutually-exclusive.

Atkeson, Andrew. (2020). "What Will Be the Economic Impact of COVID-19 in the US? Rough Estimates of Disease Scenarios." Working Paper 26867, National Bureau of Economic Research. URL http://www.nber.org/papers/w26867.

Benveniste, Lawrence M. and José A. Scheinkman. (1979). "On the Differentiability of the Value Function in Dynamic Models of Economics." *Econometrica* 47 (3):727–732.

Berger, David W, Kyle F Herkenhoff, and Simon Mongey. (2020). "An SEIR Infectious Disease Model with Testing and Conditional Quarantine." Working Paper 26901, National Bureau of Economic Research. URL http://www.nber.org/papers/w26901.

Bethune, Zachary A and Anton Korinek. (2020). "Covid-19 Infection Externalities: Trading Off Lives vs. Livelihoods." Working Paper 27009, National Bureau of Economic Research. URL http://www.nber.org/papers/w27009.

Budish, Eric B. (2020). "$R < 1$ as an Economic Constraint: Can We "Expand the Frontier" in the Fight Against Covid-19?" *University of Chicago, Becker Friedman Institute for Economics Working Paper* (2020-31).

Eichenbaum, Martin S, Sergio Rebelo, and Mathias Trabandt. (2020)a. "The Macroeconomics of Epidemics." Working Paper 26882, National Bureau of Economic Research. URL http://www.nber.org/papers/w26882.

———. (2020)b. "The Macroeconomics of Testing and Quarantining." Discussion Paper DP14688, CEPR. URL http:https://portal.cepr.org/discussion-paper/16312.

Farboodi, Maryam, Gregor Jarosch, and Robert Shimer. (2020). "Internal and external effects of social distancing in a pandemics." *Covid Economics* 9 (24 April):22–58. URL https://cepr.org/sites/default/files/news/CovidEconomics9.pdf.

Ferguson, Neil M., Daniel Laydon, and Gemma Nedjati-Gilani et al. (2020). "Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand." Report. URL doi:https://doi.org/10.25561/77482.

Garibaldi, Pietro, Espen R. Moen, and Christopher A Pissarides. (2020). "Modelling contacts and transitions in the SIR epidemics model." *Covid Economics* 5 (16 April):1–20. URL https://cepr.org/sites/default/files/news/CovidEconomics5.pdf.

Gonzalez-Eiras, Martín and Dirk Niepelt. (2020). "On the optimal 'lockdown' during an epidemic." *Covid Economics* 7 (20 April):68–87. URL https://cepr.org/sites/default/files/news/CovidEconomics7.pdf.

Hall, Robert E, Charles I Jones, and Peter J Klenow. (2020). "Trading Off Consumption and COVID-19 Deaths." Working Paper 27340, National Bureau of Economic Research. URL http://www.nber.org/papers/w27340.

Jones, Callum J, Thomas Philippon, and Venky Venkateswaran. (2020)a. "A note on efficient mitigation policies." *Covid Economics* 4 (14 April):25–46. URL https://cepr.org/sites/default/files/news/CovidEconomics4.pdf.

———. (2020)b. "Optimal Mitigation Policies in a Pandemic: Social Distancing and Working from Home." Working Paper 26984, National Bureau of Economic Research. URL http://www.nber.org/papers/w26984.

Kermack, William O. and Anderson G. McKendrick. (1927). "A Contribution to the Mathematical Theory of Epidemics." *Proceedings of the Royal Society A* 115 (772):700–721.

Kissler, Stephen M., Christine Tedijanto, Edward Goldstein, Yonatan H. Grad, and Marc Lipsitch. (2020). "Projecting the transmission dynamics of SARS-CoV-2 through the post-pandemic period." *Science* 368 (6493):860–868. URL https://science.sciencemag.org/content/368/6493/860.

Krueger, Dirk, Harald Uhlig, and Taojun Xie. (2020). "Macroeconomic Dynamics and Reallocation in an Epidemic." Working Paper 27047, National Bureau of Economic Research. URL http://www.nber.org/papers/w27047.

Lavezzo, Enrico, Elisa Franchin, Constanze Ciavarella, and Gina Cuomo-Dannenburg et al. (2020). "Suppression of a SARS-CoV-2 outbreak in the Italian municipality of Vo'." *Nature* URL https://doi.org/10.1038/s41586-020-2488-1.

Li, Ruiyun, Sen Pei, Bin Chen, Yimeng Song, Tao Zhang, Wan Yang, and Jeffrey Shaman. (2020). "Substantial undocumented infection facilitates the rapid dissemination of novel coro-

navirus (SARS-CoV-2)." *Science* 368 (6490):489–493. URL https://science.sciencemag.org/content/368/6490/489.

Moghadasa, Seyed M., Meagan C. Fitzpatrick, Pratha Sahb, Abhishek Pandeyb, Affan Shoukatb, Burton H. Singerd, and Alison P. Galvani. (2020). "The implications of silent transmission for the control of COVID-19 outbreaks." *Proceedings of the National Academy of Science* URL https://doi.org/10.1073/pnas.2008373117.

Oran, Daniel P. and Eric J. Topol. (2020). "Prevalence of Asymptomatic SARS-CoV-2 Infection." *Annals of Internal Medicine* URL https://doi.org/10.7326/M20-3012.

Piguillem, Facundo and Liyan Shi. (2020). "The Optimal COVID-19 Quarantine and Testing Policies." Tech. Rep. 20/04, EIEF Working Papers.

Pollinger, Stefan. (2020). "Optimal case detection and social distancing policies to suppress Covid-19." *Covid Economics* 23 (28 May):152–187. URL https://cepr.org/sites/default/files/news/CovidEconomics23.pdf.

Pueyo, Tomas. (2020). "Coronavirus: The Hammer and the Dance." blog post, Medium. URL https://medium.com/@tomaspueyo/coronavirus-the-hammer-and-the-dance-be9337092b56.

Silverman, Justin D., Nathaniel Hupert, and Alex D. Washburne. (2020). "Using Influenza Surveillance Networks to Estimate State-Specific Prevalence of SARS-CoV-2 in the United States." *Science Translational Medicine* URL https://stm.sciencemag.org/content/early/2020/06/22/scitranslmed.abc1126.

Toxvaerd, Flavio. (2020). "Equilibrium social distancingn." *Covid Economics* 15 (7 May):110–133. URL https://cepr.org/sites/default/files/news/CovidEconomics15.pdf.

World Health Organization. (2020). "Immunity Passports in the Context of COVID-19." Scientific Brief 24 April 2020, World Health Organization.

A Appendix: Derivations and Proofs

Proof of proposition 1: Necessity of $X^{eq}(R) = X^*(R)$ and $\mathcal{V}^{eq}(R,R) = \mathcal{V}^*(R)$ is obvious. We thus show that given $X^{eq}(R) = X^*(R)$ and $\mathcal{V}^{eq}(R,R) = \mathcal{V}^*(R)$, condition (ii) is necessary and sufficient for efficiency. Condition (ii) is necessary since $\mathcal{V}_R^{eq}(R,R) + \mathcal{V}_r^{eq}(R,R) \neq \mathcal{V}^{*\prime}(R)$, together with $\mathcal{V}^{eq}(R) = \mathcal{V}^*(R)$ implies that $\mathcal{V}^{eq}(R \pm \varepsilon, R) > \mathcal{V}^*(R)$ for some small perturbation ε. But this contradicts that R was an equilibrium of the reduced form game. Conditions (i) and (ii) together are also sufficient. From condition (i), if $\mathcal{V}^{eq}(R,R) = \mathcal{V}^*(R)$ and $\mathcal{V}^{eq}(R',R') \leq \mathcal{V}^*(R')$ for R' in a neighborhood of R, then $\mathcal{V}^{*\prime}(R) = \mathcal{V}_R^{eq}(R,R) + \mathcal{V}_r^{eq}(R,R)$. But then, it follows that $\Phi^{eq}(\pi) = \Phi^*(\pi) - \mathcal{V}_R^{eq}(R,R) = \mathcal{V}^{*\prime}(R) - \mathcal{V}_R^{eq}(R,R) = \mathcal{V}_r^{eq}(R,R)$, and hence that R is implemented in a Nash equilibrium of the hybrid game.

Value functions: Here we complete the characterization of continuation values and shadow prices. Fix a sequence $\{R_t, \mathcal{U}_t\}_{t=0}^{\infty}$ of infection rate choices and instantaneous utilities, and let V_t denote the expected life-time utility of agents and π_t the population state in period t. V_t and π_t must satify

$$V_t = (1-\beta)\mathcal{U}_t + \beta(1 - \delta\pi_t(i))V_{t+1}$$
$$\pi_{t+1} = (1-\delta\pi_t(i))^{-1} \cdot T(R_t\pi_t(i)) \cdot \pi_t$$

The sequence of population states $\{\pi_t\}_{t=0}^{\infty}$ is uniquely determined from the initial state π_0 and the sequence of target infection rates $\{R_t\}_{t=0}^{\infty}$. Let V_t^s, V_t^i, V_t^r denote the life-time utility of an agent who at date t is in state s, i, or r. These life-time utilities satisfy

$$V_t^r = (1-\beta)\mathcal{U}_t + \beta V_{t+1}^r = (1-\beta)\sum_{s=0}^{\infty}\beta^s \mathcal{U}_{t+s} \in \left[\underline{V}, \overline{V}\right]$$

$$\begin{aligned} V_t^i &= (1-\beta)\mathcal{U}_t + \beta(1-\gamma-\delta)V_{t+1}^i + \beta\gamma V_{t+1}^r \\ &= (1-\beta)\sum_{s=0}^{\infty}\beta^s\left(\frac{\gamma}{\gamma+\delta} + \frac{\delta}{\gamma+\delta}(1-\gamma-\delta)^s\right)\mathcal{U}_{t+s} \in \left[\frac{\gamma}{\gamma+\delta}V_t^r, V_t^r\right] \end{aligned}$$

$$\text{with } V_t^i - \frac{\gamma}{\gamma+\delta}V_t^r \in \left[\frac{\delta}{\gamma+\delta}\frac{1-\beta}{1-\beta+\beta(\gamma+\delta)}\underline{V}, \frac{\delta}{\gamma+\delta}\frac{1-\beta}{1-\beta+\beta(\gamma+\delta)}\overline{V}\right]$$

and $\lim_{\beta\to 1} V_t^i = \frac{\gamma}{\gamma+\delta}V_t^r$, and

$$\begin{aligned} V_t^s &= (1-\beta)\mathcal{U}_t + (1 - R_t\pi_t(i))\beta V_{t+1}^s + \beta R_t\pi_t(i)V_{t+1}^i \\ &= V_t^i + (\gamma+\delta)\sum_{s=0}^{\infty}\beta^{s+1}\left(V_{t+s+1}^i - \frac{\gamma}{\gamma+\delta}V_{t+s+1}^r\right)\prod_{k=0}^{s-1}(1 - R_{t+k}\pi_{t+k}(i)) \in \left[V_t^i, V_t^r\right] \end{aligned}$$

The sequence of expected life-time utilities satisfies

$$V_t = \pi_t(s)V_t^s + \pi_t(i)V_t^i + (1 - \pi_t(s) - \pi_t(i))V_t^r,$$

where V_t^s, V_t^i, V_t^r denote the life-time utility of an agent who at date t is in state s, i, or r.[40] Applying this decomposition yields the representation

$$v^*(\pi) = \pi(s)v_s^*(\pi) + \pi(i)v_i^*(\pi) + (1 - \pi(s) - \pi(i))v_r^*(\pi)$$

of the planner's value function, and

$$\hat{v}\left(\pi^k, \pi\right) = \pi^k(s)\hat{v}_s\left(\pi^k, \pi\right) + \pi^k(i)\hat{v}_i\left(\pi^k, \pi\right) + \left(1 - \pi^k(s) - \pi^k(i)\right)\hat{v}_r\left(\pi^k, \pi\right)$$

for the equilibrium value function, along with the subsequent characterizations of externalities. In addition, we check that

$$\frac{\delta}{1-\beta+\beta(\gamma+\delta)}\overline{V} \geq V_t^s - V_t^i \geq \frac{\delta}{1-\beta+\beta(\gamma+\delta)}\underline{V}(1-\beta)\sum_{s=0}^{\infty}\beta^s\prod_{k=0}^{s-1}(1 - R_{t+k}\pi_{t+k}(i))$$

[40] It is easy to prove this step by substituting this guess into the above recursion for V_t.

where $\prod_{k=0}^{s-1}(1-R_{t+k}\pi_{t+k}(i))$ is the probability of remaining without infection from t until $t+s$. Since these probabilities are uniformly bounded away from 0 under the SIR dynamics, $V_t^s - V_t^i$ is then uniformly bounded away from 0.

Finally, for given $\pi_0(i)$, we can construct an upper bound for $\frac{\partial v^*(\pi_{+1})}{\partial \pi(s)} - \frac{\partial v^*(\pi_{+1})}{\partial \pi(i)}$ by computing the envelope conditions from the planner's problem:

$$\begin{aligned}\frac{\partial v^*(\pi_t)}{\partial \pi(s)} &= (1-R_t\pi_t(i))\beta\frac{\partial v^*(\pi_{t+1})}{\partial \pi(s)} + R_t\pi_t(i)\beta\frac{\partial v^*(\pi_{t+1})}{\partial \pi(i)} \\ &= \frac{\Lambda_{t+1}\pi_{t+1}(s)}{\Lambda_t\pi_t(s)}\beta\frac{\partial v^*(\pi_{t+1})}{\partial \pi(s)} + \left(1-\frac{\Lambda_{t+1}\pi_{t+1}(s)}{\Lambda_t\pi_t(s)}\right)\beta\frac{\partial v^*(\pi_{t+1})}{\partial \pi(i)}\end{aligned}$$

$$\begin{aligned}\frac{\partial v^*(\pi_t)}{\partial \pi(i)} &= -\beta\delta v^*(\pi_{t+1}) + \frac{(\delta-R_t)\pi_t(s)}{1-\delta\pi_t(i)}\beta\frac{\partial v^*(\pi_{t+1})}{\partial \pi(s)} + \frac{1-\gamma-\delta+R_t\pi_t(s)}{1-\delta\pi_t(i)}\beta\frac{\partial v^*(\pi_{t+1})}{\partial \pi(i)} \\ &= -\beta\delta v^*(\pi_{t+1}) + \left(\frac{\pi_{t+1}(s)-\pi_t(s)}{\pi_t(i)}\right)\beta\frac{\partial v^*(\pi_{t+1})}{\partial \pi(s)} + \frac{\pi_{t+1}(i)}{\pi_t(i)}\beta\frac{\partial v^*(\pi_{t+1})}{\partial \pi(i)} \\ &= -\beta\delta v^*(\pi_{t+1}) + \frac{\pi_{t+1}(i)}{\pi_t(i)}\beta\left(\frac{\partial v^*(\pi_{t+1})}{\partial \pi(i)} - \frac{\partial v^*(\pi_{t+1})}{\partial \pi(s)}\right) + \beta\frac{\partial v^*(\pi_{t+1})}{\partial \pi(s)}\frac{1-\gamma-\delta+\delta\pi_t(s)}{1-\delta\pi_t(i)}\end{aligned}$$

and therefore

$$\begin{aligned}\frac{\partial v^*(\pi_t)}{\partial \pi(s)} - \frac{\partial v^*(\pi_t)}{\partial \pi(i)} &= \beta\delta v^*(\pi_{t+1}) + \beta\frac{\partial v^*(\pi_{t+1})}{\partial \pi(i)}\frac{\gamma}{\Lambda_{t+1}} \\ &\quad + \beta\left(\frac{\partial v^*(\pi_{t+1})}{\partial \pi(s)} - \frac{\partial v^*(\pi_{t+1})}{\partial \pi(i)}\right)\left(\frac{\pi_{t+1}(i)}{\pi_t(i)} + \frac{\Lambda_{t+1}\pi_{t+1}(s)}{\Lambda_t\pi_t(s)} - 1 + \frac{\gamma}{\Lambda_{t+1}}\right) \\ &\leq \sum_{s=0}^{\infty}\beta^{s+1}\frac{\pi_{t+s}(i)}{\pi_t(i)}\left(\delta v^*(\pi_{t+s+1}) + \beta\frac{\gamma}{\Lambda_{t+s+1}}\frac{\partial v^*(\pi_{t+s+1})}{\partial \pi(i)}\right) \leq \sum_{s=0}^{\infty}\frac{\pi_{t+s}(i)}{\pi_t(i)}\delta\overline{V}\end{aligned}$$

given that $\frac{\partial v^*(\pi_{t+s+1})}{\partial \pi(i)} < 0$, $\beta < 1$ and $v^*(\pi_{t+s+1}) < \overline{V}$. Now, $\sum_{s=0}^{\infty}\pi_{t+s}(i)\delta \leq \sum_{s=0}^{\infty}\pi_{t+s}(i)\delta\Lambda_{t+s}/\Lambda_\infty = (\Lambda_t - \Lambda_\infty)/\Lambda_\infty$. For given $\pi_0(i) > 0$, it follows that $\frac{\partial v^*(\pi_t)}{\partial \pi(s)} - \frac{\partial v^*(\pi_t)}{\partial \pi(i)}$ has a uniform (in β) upper bound. On the other hand, $\frac{\partial v^*(\pi_t)}{\partial \pi(s)} - \frac{\partial v^*(\pi_t)}{\partial \pi(i)}$ can not readily be bounded from below. ■ Q.E.D.

Proposition 2: Immediate, given the first-order conditions in the text. ■ Q.E.D.

Proposition 3: Fix K and K' such that $K' \geq \hat{v}_s^{eq}(\pi) - \hat{v}_i^{eq}(\pi) \geq \frac{\overline{R}}{\gamma+\delta}K$. This implies

$$\frac{1-\beta}{\beta}\Phi^{eq}(\pi) = \pi(s)\pi(i)\left(\hat{v}_s^{eq}(\pi) - \hat{v}_i^{eq}(\pi)\right) \geq \frac{\overline{R}}{\gamma+\delta}K\pi(s)\pi(i)$$

Therefore, the equilibrium policy satisfies $\frac{1-\beta}{\beta}\mathcal{V}_r^{eq}(R,R) \geq \frac{\overline{R}}{\gamma+\delta}\pi(s)K\pi(i)$, or $\frac{1-\beta}{\beta}\mathcal{V}_r^{eq}(R,R) \geq K\pi(i)$, whenever $\pi(s) > (\gamma+\delta)/\overline{R}$. Now fix $\kappa < \frac{1-\delta}{1-\gamma-\delta+\overline{R}} < 1$ and $\eta > 0$. There exists $\xi > 0$, such that whenever $\max\left\{\beta, \underline{V}/\overline{V}\right\} > 1-\xi$ and $R > \underline{R}+\eta$, $\frac{1-\beta}{\beta}\mathcal{V}_r^{eq}(R,R) < K\kappa\eta$, and therefore $\mathcal{V}_r^{eq}(R,R) < \Phi^{eq}(\pi)$, for $\pi(i) \geq \kappa\eta$. Therefore, starting with $\pi(i) \geq \kappa\eta$, policy remains within $[\underline{R}, \underline{R}+\eta]$, until $\pi_t(i) \leq \kappa\eta$.

Suppose next that $\pi_t(i) \leq \kappa\eta$. It then follows that

$$\pi_{t+1}(i) = \pi_t(i)\frac{1-\gamma-\delta+R_t\pi_t(s)}{1-\delta\pi_t(i)} \leq \eta\kappa\frac{1-\gamma-\delta+\overline{R}}{1-\delta\eta} < \eta$$

Therefore, $\pi_t(i) \leq \kappa\eta$ implies $\pi_{t+1}(i) < \eta$, and whenever $\pi_{t+1}(i) \in (\kappa\eta, \eta)$, it must be the case that $R_{t+1} \leq \underline{R}+\eta$, and hence $\pi_{t+2}(i) \leq \pi_{t+1}(i) \leq \eta$. Therefore, once $\pi_t(i) \leq \kappa\eta$, we must have $\pi_{t+s}(i) \leq \eta$ for all subsequent periods at the planner's solution.

In addition, there exists $\kappa' > 0$, such that

$$\frac{1-\beta}{\beta}\mathcal{V}_r^{eq}(R,R) > K'\kappa'\eta \text{ for } R < \frac{\gamma+\delta}{\pi(s)}$$

Therefore $\mathcal{V}_r^{eq}(R,R) > \Phi^{eq}(\pi)$ for $R \leq \frac{\gamma+\delta}{\pi(s)}$ and $\pi(i) < \kappa'\eta$, which implies that it must be optimal to set $R_t^* > \frac{\gamma+\delta}{\pi(s)}$ and $R_t^{eq} > \frac{\gamma+\delta}{\pi(s)}$. But then $\pi_{t+1}(i) > \pi_t(i)$. But then, it follows that at the equilibrium, $\pi_t(i)$ cannot permanently escape from the set $(\kappa'\eta, \kappa\eta)$.

Exactly the same steps also apply to the planner's solution, provided that we can find constant K such that $\frac{\partial v^*(\pi)}{\partial \pi(s)} - \frac{\partial v^*(\pi)}{\partial \pi(i)} \geq \frac{\overline{R}}{\gamma+\delta}K$. ■ Q.E.D.

Proposition 4: Written in calendar time, the planner's problem is

$$\max_{\{R(n\Delta)\}} \left(1-e^{-\rho\Delta}\right) \sum_{n=0}^{\infty} e^{-\rho\Delta n} \frac{\gamma}{\gamma+\delta\left(1-\pi(i,\Delta n)-\pi(s,\Delta n)\right)} \mathcal{V}^*\left(R(n\Delta)\right),$$

subject to the law of motion for π. Proposition 3 implies that for any $\eta > 0$, there exists $\overline{\Delta} > 0$ and finite N, such that $\pi(i, n\Delta) < \eta$ and $\left\| R(n\Delta) - \frac{\gamma+\delta}{\pi(s,n\Delta)} \right\| < \eta$ for $\Delta \leq \overline{\Delta}$ and $n > N$. It then follows that

$$\begin{aligned}
&\lim_{\Delta\to 0} \max_{\{R_{n\Delta}\}} \frac{1-e^{-\rho\Delta}}{\Delta} \sum_{n=0}^{\infty} \Delta e^{-\rho\Delta n} \frac{\gamma}{\gamma+\delta\left(1-\pi(i,\Delta n)-\pi(s,\Delta n)\right)} \mathcal{V}^*\left(R_{n\Delta}\right) \\
&= \max_{R(\tau)} \rho \int_0^{\infty} e^{-\rho\tau} \frac{\gamma}{\gamma+\delta\left(1-\pi(s,\tau)\right)} \mathcal{V}^*\left(R(\tau)\right) d\tau, \text{ where } \pi(s,\tau) = \frac{\gamma+\delta}{R(\tau)} \text{ for } \tau > 0.
\end{aligned}$$

The expression $\frac{\gamma}{\gamma+\delta(1-\pi(s,\tau))}\mathcal{V}^*(R(\tau))$ with $\pi(s,\tau) = \frac{\gamma+\delta}{R(\tau)}$ reaches a maximum when $R(\tau) = R^*$. Therefore the planner's objective function is bounded from above by the long-run optimal policy $R(\tau) = R^*$ for any $\tau > 0$.

We complete the proof of part (i) by constructing a policy path $\{R_{n\Delta}\}$ that displays fast (geometric) convergence to R^*, and therefore enables the planner therefore to attain the long-run optimum at the continuous time limit. Consider the path $R(n\Delta) = \gamma/\pi(s,\Delta n)$. Since

$$\Lambda(\Delta(n+1))\,\pi(i,\Delta(n+1)) = \Lambda(\Delta n)\,\pi(i,\Delta n)(1-\delta) = \Lambda(0)\,\pi(i,0)(1-\delta)^{n+1},$$

the total measure of surviving agents converges to $1-\pi(i,0)$ as $n\to\infty$, the total measure of agents who have recovered converges to $(\gamma/\delta)\cdot\pi(i,0)$, and the proportion of susceptible agents converges to

$$\lim_{n\to\infty} \pi(s,\Delta n) = \frac{1-\frac{\gamma+\delta}{\delta}\pi(i,0)}{1-\pi(i,0)}$$

Set $\overline{\pi}(i)$ such that

$$\frac{1-\frac{\gamma+\delta}{\delta}\overline{\pi}(i)}{1-\overline{\pi}(i)} = \frac{\gamma+\delta}{R^*} \Longleftrightarrow \overline{\pi}(i) = \left(\frac{R^*}{\gamma+\delta}-1\right)\frac{\delta}{R^*-\delta}.$$

If $\pi(i,0) \leq \overline{\pi}(i)$, there exists a path that converges geometrically to a limit with $\lim_{n\to\infty}\pi(s,\Delta n) = \frac{\gamma+\delta}{R^*}$. If instead $\pi(i,0) > \overline{\pi}(i)$, then the planner reaches $\pi(i,n\Delta) \leq \overline{\pi}(i)$ by setting $R(n\Delta) = \underline{R}$ for a finite number of periods. Therefore the planner's objective function and optimal policy converge to the long-run optimum at the instantaneous propagation limit. ■ Q.E.D.

Proposition 5: We characterize the continuous time limit of the equilibrium conditions.

Value functions and First-order conditions: Notice first that for any $\tau > 0$,

$$\begin{aligned}
V^r(\tau) &= \lim_{\Delta\to 0} \rho\Delta \sum_{n=0}^{\infty} e^{-\rho\Delta n} \mathcal{V}^{eq}\left(R(\tau+n\Delta), R(\tau+n\Delta)\right) = \rho\int_{\tau}^{\infty} e^{-\rho\left(\tau'-\tau\right)} \mathcal{V}^{eq}\left(R\left(\tau'\right), R\left(\tau'\right)\right) d\tau' \\
V^i(\tau) &= \lim_{\Delta\to 0} \rho\Delta \sum_{n=0}^{\infty} e^{-\rho\Delta n} \left(\frac{\gamma}{\gamma+\delta} + \frac{\delta}{\gamma+\delta}(1-\gamma-\delta)^n\right) \mathcal{V}^{eq}\left(R(\tau+n\Delta), R(\tau+n\Delta)\right) = \frac{\gamma}{\gamma+\delta} V^r(\tau)
\end{aligned}$$

For $V^s(\tau)$, we have

$$V^s(\tau) = \left(1-e^{-\rho\Delta}\right)\mathcal{V}^{eq}\left(R(\tau), R(\tau)\right) + e^{-\rho\Delta} V^i(\tau+\Delta) + e^{-\rho\Delta} R(\tau)\,\xi^{\Delta}(\tau)\,\Delta\left(V^s(\tau+\Delta) - V^i(\tau+\Delta)\right)$$

where $\xi^{\Delta}(\tau) = \pi(i,\tau)/\Delta$. Taking the limit as $\Delta \to 0$, this yields the Hamilton-Jacobi-Bellman equation

$$\rho V^s(\tau) = \rho \mathcal{V}^{eq}(R(\tau), R(\tau)) + V^{s\prime}(\tau) - R(\tau)\xi(\tau)\left(V^s(\tau) - V^i(\tau)\right)$$

where $\xi(\tau) = \lim_{\Delta\to 0}\xi^{\Delta}(\tau)$, and $V^{s\prime}(\cdot)$ denotes the time derivative of $V^s(\cdot)$.

The first-order condition for $R(\tau)$ is

$$\mathcal{V}_r^{eq}(R(\tau), R(\tau)) = \Phi^{eq}(\tau) = \frac{e^{-\rho\Delta}}{1-e^{-\rho\Delta}}\xi^{\Delta}(\tau)\Delta\pi(s,\tau)\left(V^s(\tau) - V^i(\tau)\right)$$

Taking limits as $\Delta \to 0$, we obtain $\lim_{\Delta\to 0}\pi(s,\tau) = \frac{\gamma+\delta}{R(\tau)}$,

$$\lim_{\Delta\to 0}\Phi^{eq}(\tau) = \frac{1}{\rho}\xi(\tau)\frac{\gamma+\delta}{R(\tau)}\left(V^s(\tau) - V^i(\tau)\right)$$

and therefore the FOC at the limit satisfies

$$\xi(\tau)\left(V^s(\tau) - V^i(\tau)\right) = \rho\mathcal{V}_r^{eq}(R(\tau), R(\tau))R(\tau)\frac{1}{\gamma+\delta}.$$

Substituting the FOC into the H-J-B equation for $V^s(\cdot)$, we obtain

$$\rho V^s(\tau) = V^{s\prime}(\tau) + \rho\mathcal{V}^{eq}(R(\tau)) - \rho\frac{R(\tau)^2}{\gamma+\delta}\mathcal{V}_r^{eq}(R(\tau), R(\tau))$$

which yields the solution

$$V^s(\tau) = \rho\int_{\tau}^{\infty} e^{-\rho(\tau'-\tau)}\left(\mathcal{V}^{eq}\left(R\left(\tau'\right), R\left(\tau'\right)\right) - \frac{R(\tau')^2}{\gamma+\delta}\mathcal{V}_r^{eq}\left(R\left(\tau'\right), R\left(\tau'\right)\right)\right)d\tau'$$

and

$$V^s(\tau) - V^i(\tau) = \frac{\rho}{\gamma+\delta}\int_{\tau}^{\infty} e^{-\rho(\tau'-\tau)}\left(\delta\mathcal{V}^{eq}\left(R\left(\tau'\right), R\left(\tau'\right)\right) - R\left(\tau'\right)^2\mathcal{V}_r^{eq}\left(R\left(\tau'\right), R\left(\tau'\right)\right)\right)d\tau'$$

Derivation of ODE: Substituting $V^s(\tau) - V^i(\tau)$ into the FOC we obtain

$$\mathcal{V}_r^{eq}(R(\tau), R(\tau))R(\tau) = \xi(\tau)\int_{\tau}^{\infty} e^{-\rho(\tau'-\tau)}\left(\delta\mathcal{V}^{eq}\left(R\left(\tau'\right), R\left(\tau'\right)\right) - R\left(\tau'\right)^2\mathcal{V}_r^{eq}\left(R\left(\tau'\right), R\left(\tau'\right)\right)\right)d\tau'.$$

Taking time derivatives on both sides yields

$$\begin{aligned}&\frac{R'(\tau)}{R(\tau)}\left(1 + \frac{\mathcal{V}_{rr}^{eq}(R(\tau), R(\tau)) + \mathcal{V}_{rR}^{eq}(R(\tau), R(\tau))}{\mathcal{U}_r^{eq}(R(\tau), R(\tau))}R(\tau)\right)\\ &= \frac{\xi'(\tau)}{\xi(\tau)} + \rho - \frac{\delta\mathcal{V}^{eq}(R(\tau), R(\tau)) - R(\tau)^2\mathcal{V}_r^{eq}(R(\tau), R(\tau))}{\mathcal{V}_r^{eq}(R(\tau), R(\tau))R(\tau)}\xi(\tau)\end{aligned}$$

In addition, from the dynamics of $\pi(s,\tau)$,

$$\frac{\pi(s,\tau+\Delta) - \pi(s,\tau)}{\pi(s,\tau)} = \frac{(\delta - R(\tau))\Delta\xi^{\Delta}(\tau)}{1-\delta\Delta\xi^{\Delta}(\tau)}$$

and therefore, taking the limit as $\Delta \to 0$,

$$\frac{\pi'(s,\tau)}{\pi(s,\tau)} = (\delta - R(\tau))\xi(\tau).$$

Combining this with $\frac{\pi'(s,\tau)}{\pi(s,\tau)} = -\frac{R'(\tau)}{R(\tau)}$ gives

$$\xi(\tau) = \frac{R'(\tau)}{R(\tau)(R(\tau)-\delta)}.$$

Taking time derivatives we obtain

$$\frac{\xi'(\tau)}{\xi(\tau)} = \frac{R''(\tau)}{R'(\tau)} - \frac{R'(\tau)}{R(\tau)} - \frac{R'(\tau)}{R(\tau)-\delta}$$

Substituting $\xi(\tau)$ and $\frac{\xi'(\tau)}{\xi(\tau)}$ into the FOC yields the following second-order ODE for $R(\cdot)$:

$$\begin{aligned}\frac{R''(\tau)}{R'(\tau)} &= \frac{R'(\tau)}{R(\tau)}\left(1+\frac{\mathcal{V}_{rr}^{eq}(R(\tau),R(\tau))+\mathcal{V}_{rR}^{eq}(R(\tau),R(\tau))}{\mathcal{V}_{r}^{eq}(R(\tau),R(\tau))}R(\tau)\right)\\ &+\frac{R'(\tau)}{R(\tau)}+\frac{R'(\tau)}{R(\tau)-\delta}-\rho+\frac{\delta\mathcal{V}^{eq}(R(\tau),R(\tau))-R(\tau)^2\mathcal{V}_{r}^{eq}(R(\tau),R(\tau))}{\mathcal{V}_{r}^{eq}(R(\tau),R(\tau))R(\tau)^2}\frac{R'(\tau)}{R(\tau)-\delta}\\ &=\frac{R'(\tau)}{R(\tau)}\left(2+\frac{\mathcal{V}_{rr}^{eq}(R(\tau),R(\tau))+\mathcal{V}_{rR}^{eq}(R(\tau),R(\tau))}{\mathcal{V}_{r}^{eq}(R(\tau),R(\tau))}R(\tau)\right)+\frac{R'(\tau)}{R(\tau)-\delta}\frac{\delta\mathcal{V}^{eq}(R(\tau),R(\tau))}{\mathcal{U}_{r}^{eq}(R(\tau),R(\tau))R(\tau)^2}-\rho.\end{aligned}$$

Boundary conditions: $\pi(s,\tau)$ is given at $\tau=\tau_0>0$. To compute $\lim_{\tau\to 0}\pi(s,\tau)$, consider optimality conditions for time $n\Delta$, with finite n, as $\Delta\to 0$. The FOC is

$$\mathcal{V}_{r}^{eq}(R(n\Delta),R(n\Delta))=\Phi^{eq}(n\Delta)=\frac{e^{-\rho\Delta}}{1-e^{-\rho\Delta}}\xi^{\Delta}(n\Delta)\Delta\pi(s,n\Delta)\left(V^{s}(n\Delta)-V^{i}(n\Delta)\right)$$

Since $\xi^{\Delta}(n\Delta)\Delta=\pi(i,n\Delta)>(1-\gamma-\delta)\pi(i,0)$ for any $\pi(i,0)>0$, the RHS is of order $o\left(\Delta^{-1}\right)$, and therefore $\lim_{\Delta\to 0}R(n\Delta)=\underline{R}$ for any finite n. From this we obtain an initial jump in $\pi(i,0)$ to $\pi(i,0_+)=0$ and $\pi(s,0)$ from $1-\pi(i,0)$ to $\pi(s,0_+)$, where $\pi(s,0_+)$ and $\pi(i,0_+)$ are defined as the limit of the sequence

$$\pi_{n+1}(s)=\frac{1-\underline{R}\pi_n(i)}{1-\delta\pi_n(i)}\pi_n(s)\text{ and }\pi_{n+1}(i)=\frac{1-\gamma-\delta+\underline{R}\pi_n(s)}{1-\delta\pi_n(i)}\pi_n(i)$$

as $n\to\infty$. This defines the initial condition for $\lim_{\tau\to 0}\pi(s,\tau)$. In addition, it must be the case that $\lim_{\tau\to\infty}R(\tau)=\overline{R}$.

Solving the ODE: Rewrite the ODE as

$$\frac{R''(\tau)}{(R'(\tau))^2}+\frac{\rho}{R'(\tau)}=\frac{2}{R(\tau)}+\frac{\mathcal{V}_{rr}^{eq}(R(\tau),R(\tau))+\mathcal{V}_{rR}^{eq}(R(\tau),R(\tau))}{\mathcal{V}_{r}^{eq}(R(\tau),R(\tau))}+\frac{1}{R(\tau)-\delta}\frac{\delta\mathcal{V}^{eq}(R(\tau),R(\tau))}{\mathcal{V}_{r}^{eq}(R(\tau),R(\tau))R(\tau)^2}$$

Integrating the RHS w.r.t. R, we obtain

$$\begin{aligned}&\int_{R(0)}^{R(\tau)}\left(\frac{2}{R}+\frac{\mathcal{V}_{rr}^{eq}(R,R)+\mathcal{V}_{rR}^{eq}(R,R)}{\mathcal{V}_{r}^{eq}(R,R)}+\frac{1}{R-\delta}\frac{\delta\mathcal{V}^{eq}(R,R)}{\mathcal{V}_{r}^{eq}(R,R)R^2}\right)dR\\ &=2\left(\log R(\tau)-\log R(0)\right)+\log\left(\mathcal{V}_{r}^{eq}(R(\tau),R(\tau))\right)-\log\left(\mathcal{V}_{r}^{eq}(R(0),R(0))\right)+\int_{R(0)}^{R(\tau)}\frac{1}{R-\delta}\frac{\delta\mathcal{V}^{eq}(R,R)}{\mathcal{V}_{r}^{eq}(R,R)R^2}dR\end{aligned}$$

Integrating the LHS w.r.t. R, we obtain

$$\int_{R(0)}^{R(\tau)}\left(\frac{R''(\tau)}{(R'(\tau))^2}+\frac{\rho}{R'(\tau)}\right)dR=\int_0^{\tau}\left(\frac{R''(\tau')}{R'(\tau')}+\rho\right)d\tau'=\log R'(\tau)-\log R'(0)+\rho\tau$$

Equating the two and exponentiating, we obtain

$$\frac{R'(0)}{R'(\tau)}e^{-\rho\tau}=\frac{R(0)^2\mathcal{V}_{r}^{eq}(R(0),R(0))}{R(\tau)^2\mathcal{V}_{r}^{eq}(R(\tau),R(\tau))}e^{-\int_{R(0)}^{R(\tau)}\frac{1}{R-\delta}\frac{\delta\mathcal{V}^{eq}(R,R)}{\mathcal{V}_{r}^{eq}(R,R)R^2}dR}$$

Once again integrating the LHS w.r.t. R, we obtain

$$\int_{R(0)}^{R(\tau)}\frac{R'(0)}{R'(\tau)}e^{-\rho\tau}dR=R'(0)\int_0^{\tau}e^{-\rho\tau'}d\tau'=\frac{R'(0)}{\rho}\left(1-e^{-\rho\tau}\right)$$

where $R'(0)$ must be determined from the boundary conditions. Integrating the RHS w.r.t. R, we obtain

$$\int_{R(0)}^{R(\tau)}\frac{R(0)^2\mathcal{V}_{r}^{eq}(R(0),R(0))}{R^2\mathcal{V}_{r}^{eq}(R,R)}e^{-\int_{R(0)}^{R}\frac{1}{R'-\delta}\frac{\delta\mathcal{V}^{eq}\left(R',R'\right)}{\mathcal{V}_{r}^{eq}(R',R')R'^2}dR'}dR$$

and combining the two and solving for $R'(0)$, we have

$$\frac{1-e^{-\rho\tau}}{1-e^{-\rho\tau\left(\overline{R}\right)}}=\frac{\int_{R(0)}^{R(\tau)}\frac{1}{R^2\mathcal{V}_{r}^{eq}(R,R)}e^{-\int_{R(0)}^{R}\frac{1}{R'-\delta}\frac{\delta\mathcal{V}^{eq}\left(R',R'\right)}{\mathcal{V}_{r}^{eq}(R',R')R'^2}dR'}dR}{\int_{R(0)}^{\overline{R}}\frac{1}{R^2\mathcal{V}_{r}^{eq}(R,R)}e^{-\int_{R(0)}^{R}\frac{1}{R'-\delta}\frac{\delta\mathcal{V}^{eq}(R',R')}{\mathcal{V}_{r}^{eq}(R',R')R'^2}dR'}dR}$$

where $\tau\left(\overline{R}\right)$ is the time at which R reaches $\overline{R}$.

Verifying the boundary: The last step is to show that $\tau\left(\overline{R}\right)=\infty$. Suppose to the contrary that $\tau\left(\overline{R}\right)<\infty$. Using $\mathcal{V}_r^{eq}\left(R,R\right)\approx\left(R-\overline{R}\right)\left(\mathcal{V}_{rr}^{eq}\left(\overline{R},\overline{R}\right)+\mathcal{V}_{rR}^{eq}\left(\overline{R},\overline{R}\right)\right)$ for R close to $\overline{R}$, we rewrite the ODE as

$$\frac{R''\left(\tau\right)}{\left(R'\left(\tau\right)\right)^2}+\frac{\rho}{R'\left(\tau\right)}=\frac{2}{R\left(\tau\right)}+\frac{1}{\overline{R}-R\left(\tau\right)}\left(K-1\right)$$

for τ close to $\tau\left(\overline{R}\right)$, where

$$K=-\frac{1}{\overline{R}-\delta}\frac{\delta\overline{V}}{\left(\mathcal{V}_{rr}^{eq}\left(\overline{R},\overline{R}\right)+\mathcal{V}_{rR}^{eq}\left(\overline{R},\overline{R}\right)\right)R^2}>0.$$

We guess and verify that $R'\left(\tau\right)=\Gamma\left(\overline{R}-R\left(\tau\right)\right)$. The guess implies that $R''\left(\tau\right)=-\Gamma R'\left(\tau\right)$, and therefore

$$K-1+2\frac{\overline{R}-R\left(\tau\right)}{R\left(\tau\right)}-\left(\frac{R''\left(\tau\right)}{\left(R'\left(\tau\right)\right)^2}+\frac{\rho}{R'\left(\tau\right)}\right)\left(\overline{R}-R\left(\tau\right)\right)=K-\frac{\rho}{\Gamma}+2\frac{\overline{R}-R\left(\tau\right)}{R\left(\tau\right)}$$

which verifies the guess with $\Gamma=\rho/K$ for $R\left(\tau\right)$ sufficiently close to $\overline{R}$. But then it follows that

$$\rho\left(\tau\left(\overline{R}\right)-\tau\right)=\rho\int_{R(\tau)}^{\overline{R}}\frac{1}{R'\left(\tau\right)}dR=K\int_{R(\tau)}^{\overline{R}}\frac{1}{\left(\overline{R}-R\right)}dR=\infty.$$

Therefore the solution to the ODE is given by

$$e^{-\rho\tau}=\frac{\int_{R(\tau)}^{\overline{R}}\frac{1}{R^2\mathcal{V}_r^{eq}(R,R)}e^{-\int_{R(0)}^{R}\frac{1}{R'-\delta}\frac{\delta\mathcal{V}^{eq}\left(R',R'\right)}{\mathcal{V}_r^{eq}(R',R')R'^2}dR'}dR}{\int_{R(0)}^{\overline{R}}\frac{1}{R^2\mathcal{V}_r^{eq}(R,R)}e^{-\int_{R(0)}^{R}\frac{1}{R'-\delta}\frac{\delta\mathcal{V}^{eq}(R',R')}{\mathcal{V}_r^{eq}(R',R')R'^2}dR'}dR}=\frac{\int_{R(\tau)}^{\overline{R}}\frac{1}{R^2\mathcal{V}_r^{eq}(R,R)}e^{\int_{R}^{\overline{R}}\frac{1}{R'-\delta}\frac{\delta\mathcal{V}^{eq}\left(R',R'\right)}{\mathcal{V}_r^{eq}(R',R')R'^2}dR'}dR}{\int_{R(0)}^{\overline{R}}\frac{1}{R^2\mathcal{V}_r^{eq}(R,R)}e^{\int_{R}^{\overline{R}}\frac{1}{R'-\delta}\frac{\delta\mathcal{V}^{eq}(R',R')}{\mathcal{V}_r^{eq}(R',R')R'^2}dR'}dR}.$$

■ Q.E.D.

Proposition 6: Part (i): Note that $\lim_{\Delta\to0}\Phi^*=\mathcal{V}^{*\prime}\left(R^*\right)>0$ for all $\tau>0$. We thus show that $\lim_{\Delta\to0}\Phi^{eq}=0$. At the instantaneous propagation limit with $R\left(\tau\right)=R^*$ for $\tau>0$, $V^s\left(\tau\right)-V^i\left(\tau\right)=\frac{\delta}{\gamma+\delta}\mathcal{V}^*\left(R^*\right)$, since $V^s\left(\tau\right)=V^r\left(\tau\right)=\mathcal{V}^*\left(R^*\right)$ (since there are no new infections), and $V^i\left(\tau\right)=\frac{\gamma}{\gamma+\delta}V^r\left(\tau\right)$, when $\beta\to1$ (recovery and mortality are resolved instantaneously).

Therefore $\lim_{\Delta\to0}\Phi^{eq}=0$ holds if and only $\lim_{\Delta\to0}\pi\left(i,\tau\right)/\left(r\Delta\right)=0$ for all $\tau>0$. To see this must be the case, notice that

$$\Lambda\left(\tau+\Delta n\right)\pi\left(s,\tau+\Delta n\right)-\Lambda\left(\tau\right)\pi\left(s,\tau\right)=\sum_{k=0}^{n-1}\Delta R\left(\tau+\Delta k\right)\Lambda\left(\tau+\Delta k\right)\xi\left(\tau+\Delta k\right),$$

where $\xi\left(\tau+\Delta k\right)=\pi\left(i,\tau+\Delta k\right)/\Delta$. Taking limits as $\Delta\to0$ on both sides with $n=\tau'/\Delta$, and noting that for $\tau>0$, $\Lambda\left(\tau+\Delta k\right)\to\Lambda\left(\tau\right)>0$, $\pi\left(s,\tau+\Delta n\right)\to\pi\left(s,\tau\right)>0$, and $R\left(\tau+\Delta k\right)\to R^*>0$ at the long-run optimum, we obtain

$$0=\lim_{\Delta\to0}\sum_{k=0}^{\tau'/\Delta}\Delta\xi\left(i,\tau+\Delta k\right)=\lim_{\Delta\to0}\int_0^{\tau'-\tau}\xi\left(\tau+s\right)ds$$

but this holds for all $\tau'>\tau>0$, if and only if $\lim_{\Delta\to0}\xi\left(\tau\right)=\lim_{\Delta\to0}\pi\left(i,\tau\right)/\left(r\Delta\right)=0$ almost everywhere.

Part (ii): Along the Markov-perfect equilibrium, $\lim_{\Delta\to0}R\left(\tau\right)=\widetilde{R}\left(\tau\right)$, and therefore $\lim_{\Delta\to0}\Phi^{eq}=\lim_{\Delta\to0}\mathcal{V}_r^{eq}\left(R\left(\tau\right),R\left(\tau\right)\right)=\mathcal{V}_r^{eq}\left(\widetilde{R}\left(\tau\right),\widetilde{R}\left(\tau\right)\right)>0$. We thus show that $\lim_{\Delta\to0}\Phi^*\left(\tau\right)=0$ for any τ, such that $R\left(\tau\right)<R^*$. Fix such a $\tau>0$. By the same arguments as in proposition 4, the planner's solution implements the long-run optimum R^* for any $\tau'>\tau$, provided that this long-run optimum is feasible. This is the case whenever, the fraction of susceptible agents as of date τ satisfies $\lim_{\Delta\to0}\pi\left(s,\tau\right)<\left(\gamma+\delta\right)/R^*$.

But then the planner's value function satisfies $\lim_{\Delta\to0}v^*\left(\pi\right)=\mathcal{V}^*\left(R^*\right)\Lambda\left(R^*\right)/\Lambda\left(\pi\right)$, where $\Lambda\left(R^*\right)/\Lambda\left(\pi\right)$ represents the conditional survival probability. But since $\Lambda\left(\pi\right)^{-1}=1+\left(\frac{\delta}{\gamma}\right)\left(1-\pi\left(i\right)-\pi\left(s\right)\right)$ for all π, it follows that

$$\lim_{\Delta\to0}\left(\frac{\partial v^*\left(\pi\right)}{\partial\pi\left(s\right)}-\frac{\partial v^*\left(\pi\right)}{\partial\pi\left(i\right)}\right)=0\text{ and hence }\lim_{\Delta\to0}\Phi^*\left(\pi\right)=0$$

at any π for which the long-run optimum remains attainable. ■ Q.E.D.

B Parameterization

In this section, we describe our benchmark parameterization of the model. Although, we do not calibrate the model, we want to think of this parameterization as a fair rendition of the pandemic. We assume a utility function for agents $\mathcal{V}^{eq}(r,R)$ that takes the form

$$\left(\frac{\mathcal{V}^{eq}(r,R)-\underline{V}}{\bar{V}-\underline{V}}\right)^2+\alpha\left(\frac{\bar{R}-r}{\bar{R}-\underline{R}}\right)^2+(1-\alpha)\left(\frac{\bar{R}-R}{\bar{R}-\underline{R}}\right)^2=1$$

for $r,R\in\left[\underline{R},\bar{R}\right]$ and $\mathcal{V}^{eq}(r,R)\in\left[\underline{V},\bar{V}\right]$. This elliptic functional form ensures that $\mathcal{V}^{eq}(r,R)$ satisfies the Inada conditions $\lim_{R\to\underline{R}}\mathcal{V}_r^{eq}(r,\underline{R})=\infty$ and $\mathcal{V}_r^{eq}\left(\bar{R},\bar{R}\right)=0$. The central planner's indirect utility function $\mathcal{V}^*(R)$ takes the same form with $\alpha=1$. At the symmetric equilibrium, we have $\mathcal{V}^{eq}(R,R)=\mathcal{V}^*(R)$, *i.e.* the equilibrium coincides with the planner's efficiency frontier (efficient implementation, $X^{eq}(\cdot)=X^*(\cdot)$), and $\mathcal{V}_r^{eq}(R,R)=\alpha\mathcal{V}^{*\prime}(R)$. If $\alpha=1$, the individual utility coincides with that of the central planner. Thus, the two static spill-overs – the economic externality and the infection risk externality – exactly offset each other. If instead $\alpha<1$, we have $\mathcal{V}_r^{eq}(R,R)<\mathcal{V}^{*\prime}(R)$ so the economic externality dominates the infection externality. Conversely, if $\alpha>1$, the infection externality dominates.

Our benchmark parameterization takes a period to be a week ($\beta=0.999$), assumes a mortality rate $\frac{\delta}{\delta+\gamma}=0.5\%$ and a time to resolution $\frac{1}{\delta+\gamma}$ of 3 weeks. These imply $\delta=0.0016$ and $\gamma=0.3317$. The initial infection rate is $\pi_0(i)=0.01\%$. The two bounds $\underline{R},\bar{R}$ defining the interval of possible values for the infection risk R are set such that (i) the basic reproduction coefficient $\mathcal{R}_0$ at the outbreak of the pandemic is 3.3; and (ii) the ratio $\bar{R}/\underline{R}$ is equal to 12, a value compatible with the Chinese experience reported in Pueyo (2020). We normalize $\bar{V}=1$ and set $\underline{V}$ to 5/6. Following Hall, Jones, and Klenow (2020), this value equates $\underline{V}/\bar{V}$ to the value of a year of life equal to $250K\$$, and the maximum instantaneous utility surplus $1-\underline{V}/\bar{V}$ to annual per-capita consumption of $50K\$$. Finally, in our baseline scenario we assume that $\alpha=1$, and we explore the sensitivity of our results to the value of α below in this Section 3.1.

C Additional Figures

Figure 15: Flattening the Curve: Higher Mortality Rate

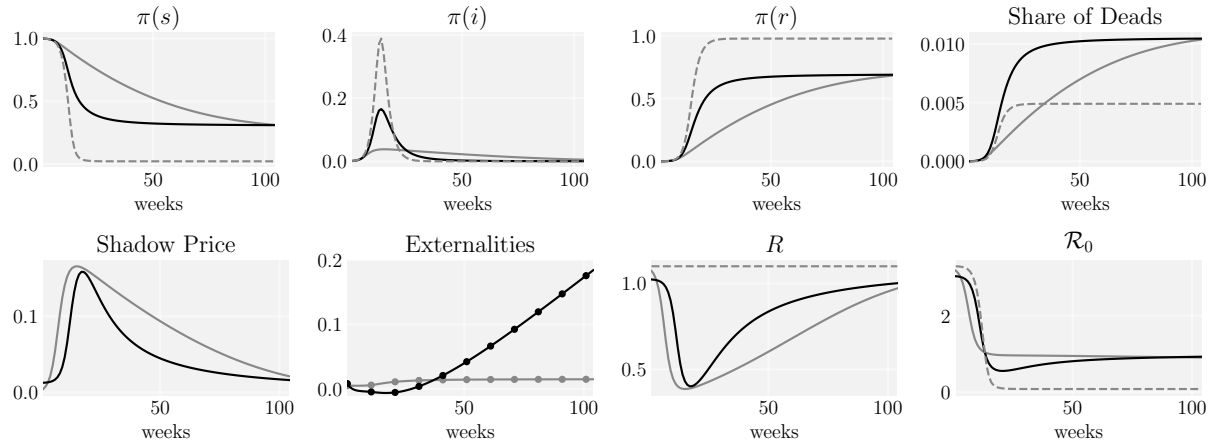

The parameters are the same as the ones chosen for our benchmark parameterization in Section B, except that we have raised the baseline mortality rate $\delta/(\gamma+\delta)$ from 0.5% to 1.5% to better illustrate our main results.

Figure 16: Instantaneous Propagation Limit: Higher Mortality Rate

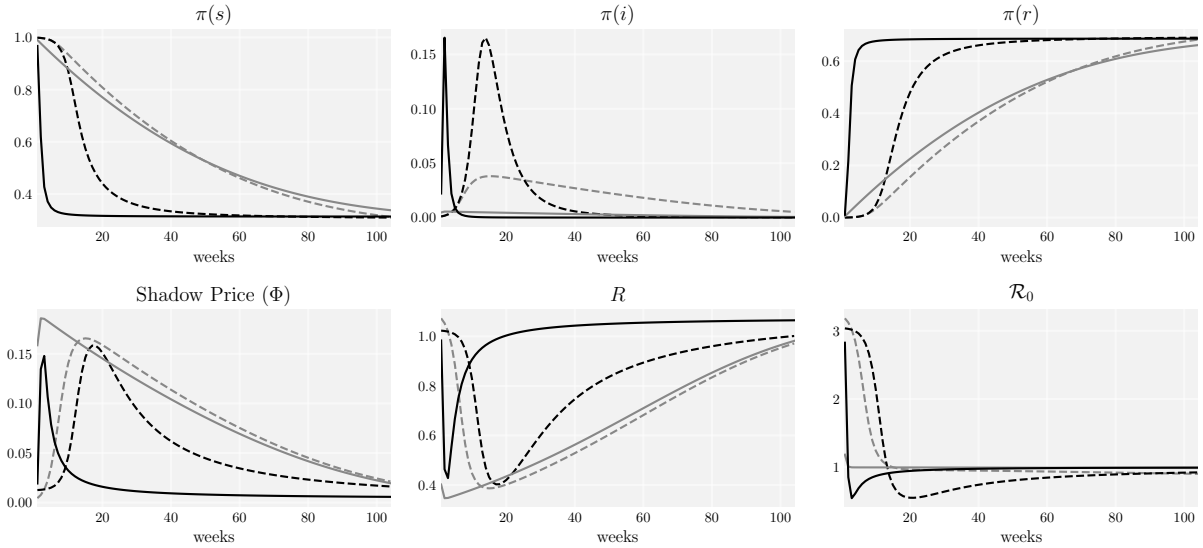

The parameters are the same as the ones chosen for our benchmark parameterization in Section B, except that we have raised the baseline mortality rate $\delta/(\gamma+\delta)$ from 0.5% to 1.5% to better illustrate our main results.

Figure 17: Congestion Effects (Long Horizon)

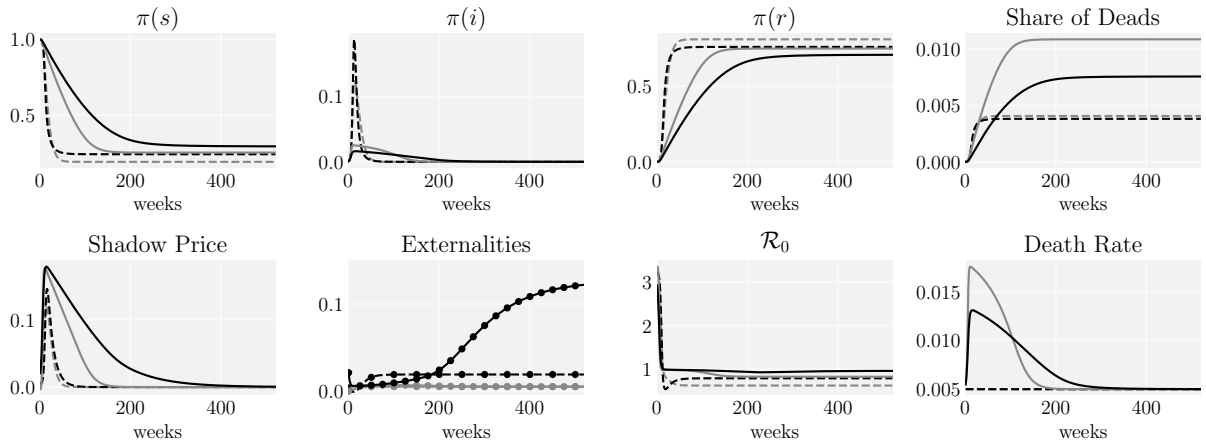

Central Planner: - - - - Benchmark ——— Congestion; **Equilibrium:** - - - Benchmark, ——— Congestion. **Private Marginal Cost:** - - - Benchmark, ——— Congestion, **Externalities:** - - - Benchmark, ——— Congestion.

Figure 18: Possibility of a Vaccine (2 year waiting time)

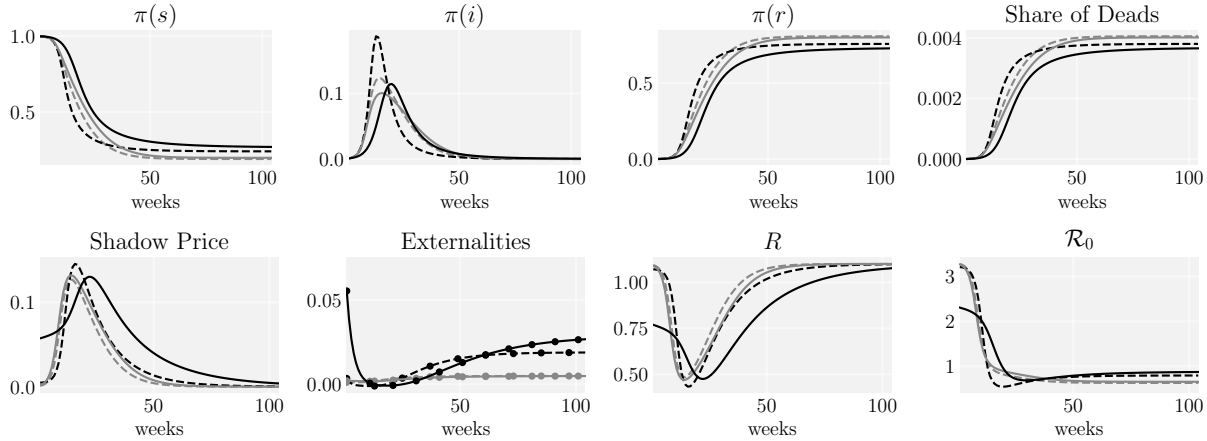

Central Planner: - - - - Benchmark ——— Vaccine; **Equilibrium:** - - - Benchmark, ——— Vaccine. **Private Marginal Cost:** - - - Benchmark, ——— Vaccine, **Externalities:** - - - Benchmark, ——— Vaccine.

Online Appendix
(Not for Publication)

D Varying the value of Life

In this section, we report the change in the dynamics when we vary the utility ratio $\underline{V}/\overline{V}$ from 5/6. In Figure 19 the ratio is raised to 0.98, meaning that both the agents and the central planner assign a higher value to life.[41] Figure 20 reports a similar experiment in which the ratio is lower to 2/3. This experiment makes clear that while the Private Marginal Cost along the Central Planner's allocation is essentially left unaffected by variations in the value of life, this is clearly not the case for externalities. When the value of life is high ($\underline{V}/\overline{V} = 0.98$), dynamic externalities are reinforced implying that the central planner applies a stronger initial hammer to contain faster the pandemic and opts more gradualism when relaxing confinement. As expected, this brings the death toll to a lower level in the long run. The opposite happens when agents and the central planner assign a lower value to life.

Figure 19: Lower Utility Gap ($\underline{V}/\overline{V} = 0.98$)

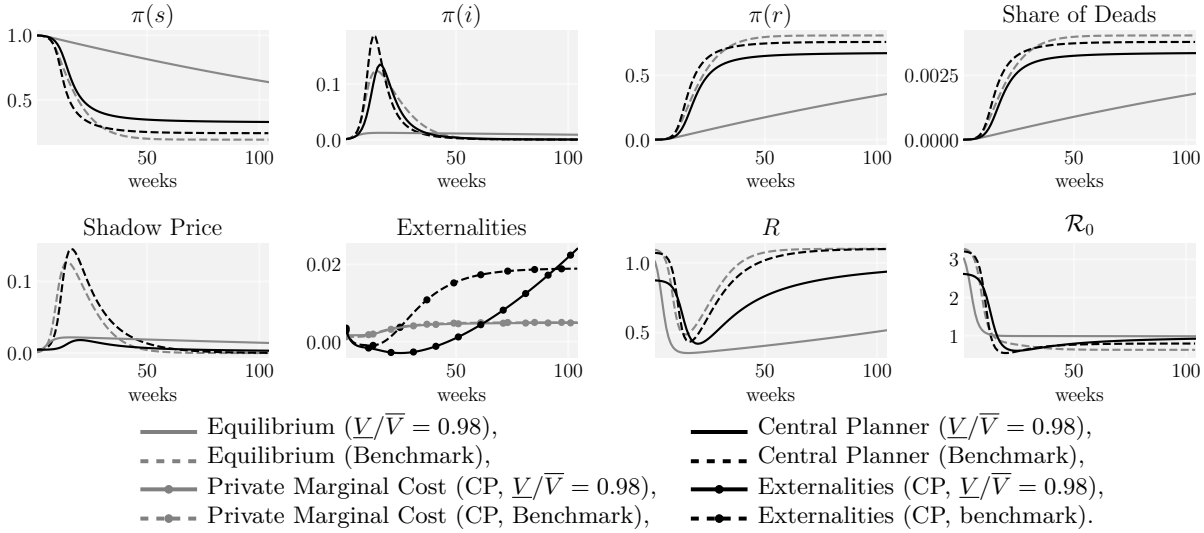

[41] The appendix of the paper also report what happens when the mortality rate is raised 0.5% to 1.5%

Figure 20: Bigger Utility Gap ($\underline{V}/\overline{V} = 2/3$)

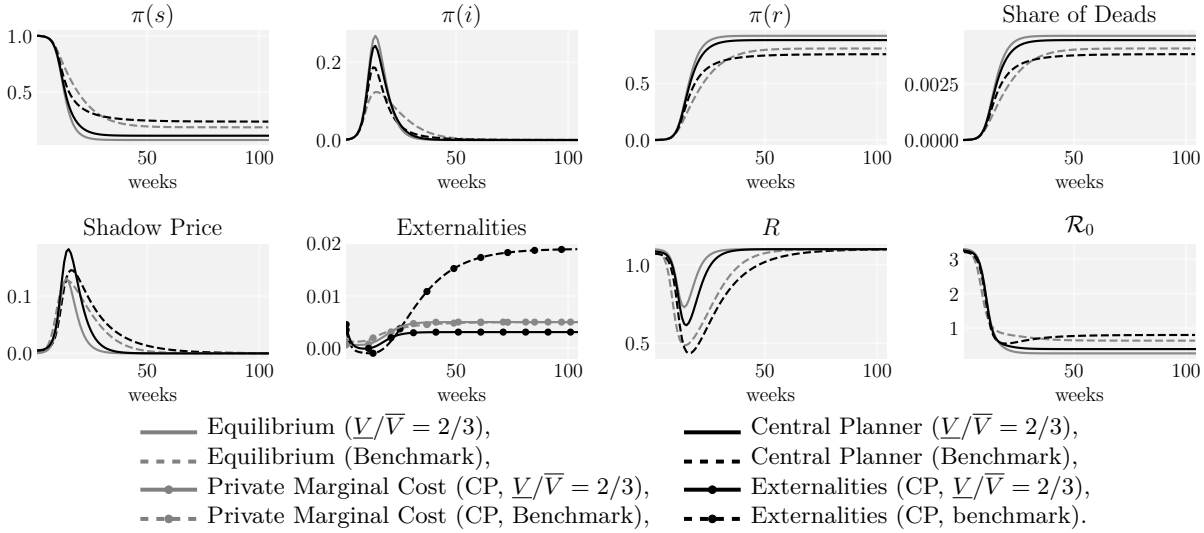

E Extensions

This section reports details pertaining to various extensions we consider in the main paper. In particular, we describe the dynamic system that was solved to obtain both the laissez-faire equilibrium and the central planner allocation.

E.1 Medical Congestion

We first introduce medical congestion in the model by letting the death probability, δ, of an infected agent be an increasing and convex function, $\delta(\pi(i))$, of the aggregate share of infected agents, $\pi(i)$.

The Central Planner's problem then writes

$$
\begin{aligned}
v(\pi_t) = \max_{R_t \in [\underline{R}, \overline{R}]} & (1-\beta)\mathcal{V}^*(R_t) + \beta(1-\delta(\pi_t(i))\pi_t(i))v(\pi_{t+1}) \\
s.t. \quad & \pi_{t+1}(s) = \pi_t(s)\left(\frac{1 - R_t\pi_t(i)}{1-\delta(\pi_t(i))\pi_t(i)}\right) \\
& \pi_{t+1}(i) = \frac{R_t\pi_t(s)\pi_t(i) + (1-\delta-\gamma)\pi_t(i)}{1-\delta(\pi_t(i))\pi_t(i)}
\end{aligned}
$$

The set of first order conditions is given by

$$
\begin{aligned}
\mathcal{V}^{*\prime}(R_t) &= \frac{\mu_t(s) - \mu_t(i)}{1-\beta}\frac{\pi_t(s)\pi_t(i)}{1-\delta_t\pi_t(i)} \\
\mu_t(s) &= \beta(1-\delta_t\pi_t(i))\left[\mu_{t+1}(s)\frac{(1-R_{t+1}\pi_{t+1}(i))}{1-\delta_{t+1}\pi_{t+1}(i)} + \mu_{t+1}(i)\frac{R_{t+1}\pi_{t+1}(i)}{1-\delta_{t+1}\pi_{t+1}(i)}\right] \\
\mu_t(i) &= \beta(1-\delta_t\pi_t(i))\Bigg[\mu_{t+1}(s)\pi_{t+1}(s)\frac{(\delta_{t+1} - R_{t+1} + \delta_{t+1}\zeta_{t+1}(1-R_{t+1}\pi_{t+1}(i)))}{(1-\delta_{t+1}\pi_{t+1}(i))^2} \\
&\qquad + \mu_{t+1}(i)\left(\frac{R_{t+1}\pi_{t+1}(s) + 1 - \gamma - \delta_{t+1}}{(1-\delta_{t+1}\pi_{t+1}(i))^2}\right) \\
&\qquad + \mu_{t+1}(i)\left(\frac{\delta_{t+1}\zeta_{t+1}(R_{t+1}\pi_{t+1}(i)\pi_{t+1}(s) + (1-\gamma)\pi_{t+1}(i) - 1)}{(1-\delta_{t+1}\pi_{t+1}(i))^2}\right) \\
&\qquad - \beta\delta_{t+1}(1+\zeta_{t+1})v(\pi_{t+2})\Bigg]
\end{aligned}
$$

where $\delta_t \equiv \delta(\pi_t(i))$ and $\zeta_t = \delta'(\pi_t(i))\pi_t(i)/\delta(\pi_t(i))$. The quantities $\mu(s)$ and $\mu(i)$ denote respectively the discounted marginal continuation value of $\pi(s)$ and $\pi(i)$:

$$\mu_t(s) = \beta(1-\delta(\pi_t(i))\pi_t(i))v_{\pi(s)}(\pi_{t+1})$$
$$\mu_t(i) = \beta(1-\delta(\pi_t(i))\pi_t(i))v_{\pi(i)}(\pi_{t+1})$$

Equilibrium: The equilibrium writes

$$\begin{aligned}
v(\pi_t^k, \pi_t) = \max_{r_t \in [\underline{R}, \overline{R}]} & (1-\beta)\mathcal{V}^{eq}(r_t, R_t) + \beta(1-\delta(\pi_t(i))\pi_t^k(i))v(\pi_{t+1}^k, \pi_{t+1}) \\
{}_t s.t. \quad & \pi_{t+1}^k(s) = \pi_t^k(s)\left(\frac{1-r_t\pi_t(i)}{1-\delta(\pi_t(i))\pi_t^k(i)}\right) \\
& \pi_{t+1}^k(i) = \frac{r_t\pi_t^k(s)\pi_t(i) + (1-\delta(\pi_t(i))-\gamma)\pi_t^k(i)}{1-\delta(\pi_t(i))\pi_t^k(i)} \\
& \pi_{t+1}(s) = \pi_t(s)\left(\frac{1-R_t\pi_t(i)}{1-\delta(\pi_t(i))\pi_t(i)}\right) \\
& \pi_{t+1}(i) = \frac{R_t\pi_t(s)\pi_t(i) + (1-\delta(\pi_t(i))-\gamma)\pi_t(i)}{1-\delta(\pi_t(i))\pi_t(i)}
\end{aligned}$$

Given that the effect of the share of infected on the death rate is external, the symmetric equilibrium ($r_t = R_t$, $\pi_t^k(i) = \pi_t(i)$, $\pi_t^k(s) = \pi_t(s)$) just writes as

$$\begin{aligned}
(1-\beta)\mathcal{V}_r^{eq}(r_t, R_t) &= (\mu_t(s) - \mu_t(i))\frac{\pi_t(s)\pi_t(i)}{1-\delta_t\pi_t(i)} \\
\mu_t(s) &= \beta(1-\delta_t\pi_t(i))\left[\mu_{t+1}(s)\frac{(1-R_{t+1}\pi_{t+1}(i))}{1-\delta_{t+1}\pi_{t+1}(i)} + \mu_{t+1}(i)\frac{R_{t+1}\pi_{t+1}(i)}{1-\delta_{t+1}\pi_{t+1}(i)}\right] \\
\mu_t(i) &= \beta(1-\delta_t\pi_t(i))\Bigg[\mu_{t+1}(s)\pi_{t+1}(s)\frac{\delta_{t+1}(1-R_{t+1}\pi_{t+1}(i)))}{(1-\delta_{t+1}\pi_{t+1}(i))^2} \\
&\qquad + \mu_{t+1}(i)\frac{1-\delta_{t+1}-\gamma+\delta_{t+1}R_{t+1}\pi_{t+1}(s)\pi_{t+1}(i)}{(1-\delta_{t+1}\pi_{t+1}(i))^2} \\
&\qquad - \beta\delta_{t+1}v(\pi_{t+2}, \pi_{t+2})\Bigg]
\end{aligned}$$

where $\delta_t \equiv \delta(\pi_t(i))$.

For our simulations, the conditional death rate $\delta(\cdot)$ takes the form

$$\delta(\pi(i)) = \underline{\delta} + \exp(\varphi\pi(i)) - 1$$

$\underline{\delta}$ corresponds to the conditional death rate that prevails in the model without congestion. In the spirit of Piguillem and Shi (2020), φ was computed such that when the economy reaches an infection rate of 1% the unconditional death rate in the economy doubles. This leads to a value for φ of 0.1682.

E.2 Transitory Immunity

Our baseline model assumes that once recovered, agents have acquired permanent immunity to the virus, we now relax this assumption and consider that, once recovered, an agent may be shifted back to the pool of susceptible

with probability ν. The central planner problem then writes

$$\begin{aligned}
\nu v(\pi_t) = \max_{R_t \in [\underline{R}, \overline{R}]} & (1-\beta)\mathcal{V}^*(R_t) + \beta(1-\delta\pi_t(i))v(\pi_{t+1}) \\
s.t. \quad \pi_{t+1}(s) &= \frac{\pi_t(s)(1-R_t\pi_t(i)) + \nu\pi_t(r)}{1-\delta\pi_t(i)} \\
\pi_{t+1}(i) &= \frac{R_t\pi_t(s)\pi_t(i) + (1-\delta-\gamma)\pi_t(i)}{1-\delta\pi_t(i)} \\
\pi_{t+1}(r) &= \frac{(1-\nu)\pi_t(r) + \gamma\pi_t(i)}{1-\delta\pi_t(i)}
\end{aligned}$$

The set of first order conditions is given by

$$\begin{aligned}
\mathcal{V}^{*\prime}(R_t) &= \frac{\mu_t(s)-\mu_t(i)}{1-\beta}\frac{\pi_t(s)\pi_t(i)}{1-\delta\pi_t(i)} \\
\mu_t(s) &= \beta(1-\delta\pi_t(i))\left[\mu_{t+1}(s)\frac{(1-R_{t+1}\pi_{t+1}(i))}{1-\delta\pi_{t+1}(i)} + \mu_{t+1}(i)\frac{R_{t+1}\pi_{t+1}(i)}{1-\delta\pi_{t+1}(i)}\right] \\
\mu_t(i) &= \beta(1-\delta\pi_t(i))\Bigg[\mu_{t+1}(s)\frac{\pi_{t+1}(s)(\delta - R_{t+1} + \delta\nu\pi_{t+1}(r)))}{(1-\delta\pi_{t+1}(i))^2} \\
&\qquad + \mu_{t+1}(i)\left(\frac{R_{t+1}\pi_{t+1}(s) + 1 - \gamma - \delta}{(1-\delta\pi_{t+1}(i))^2}\right) \\
&\qquad + \mu_{t+1}(r)\left(\frac{\gamma + \delta(1-\nu)\pi_{t+1}(r)}{(1-\delta\pi_{t+1}(i))^2}\right) - \beta\delta v(\pi_{t+2})\Bigg] \\
\mu_t(r) &= \beta(1-\delta\pi_t(i))\left[\frac{\nu\mu_{t+1}(s) + (1-\nu)\mu_{t+1}(r)}{1-\delta\pi_{t+1}(i)}\right]
\end{aligned}$$

where $\mu(s)$ and $\mu(i)$ denote respectively the discounted marginal continuation value of $\pi(s)$ and $\pi(i)$:

$$\begin{aligned}
\mu_t(s) &= \beta(1-\delta\pi_t(i))v_{\pi(s)}(\pi_{t+1}) \\
\mu_t(i) &= \beta(1-\delta\pi_t(i))v_{\pi(i)}(\pi_{t+1})
\end{aligned}$$

Note that, while the set of FOCs is written maintaining 3 state variables, 2 are actually sufficient to describe the state space of the central planner as, by construction, $\pi_t(r) = 1 - \pi_t(s) - \pi_t(i)$. We however maintained the 3 state variable representation to simplify expressions.

The equilibrium writes

$$\begin{aligned}
v(\pi_t^k, \pi_t) = \max_{r_t \in [\underline{R}, \overline{R}]} & (1-\beta)\mathcal{V}^{eq}(r_t, R_t) + \beta(1-\delta)\pi_t^k(i))v(\pi_{t+1}^k, \pi_{t+1}) \\
{}_t s.t. \quad \pi_{t+1}^k(s) &= \frac{\pi_t^k(s)(1-r_t\pi_t(i)) + \nu\pi_t(r)}{1-\delta\pi_t^k(i)} \\
\pi_{t+1}^k(i) &= \frac{r_t\pi_t^k(s)\pi_t(i) + (1-\delta-\gamma)\pi_t^k(i)}{1-\delta\pi_t^k(i)} \\
\pi_{t+1}^k(r) &= \frac{(1-\nu)\pi_t^k(r) + \gamma\pi_t^k(i)}{1-\delta\pi_t^k(i)} \\
\pi_{t+1}(s) &= \frac{\pi_t(s)(1-R_t\pi_t(i)) + \nu\pi_t(r)}{1-\delta\pi_t(i)} \\
\pi_{t+1}(i) &= \frac{R_t\pi_t(s)\pi_t(i) + (1-\delta-\gamma)\pi_t(i)}{1-\delta\pi_t(i)} \\
\pi_{t+1}(r) &= \frac{(1-\nu)\pi_t(r) + \gamma\pi_t(i)}{1-\delta\pi_t(i)}
\end{aligned}$$

Given that the effect of the share of infected on the death rate is external, the symmetric equilibrium ($r_t = R_t$,

$\pi_t^k(i) = \pi_t(i)$, $\pi_t^k(s) = \pi_t(s)$) just writes as

$$
\begin{aligned}
(1-\beta)\mathcal{V}_r^{eq}(r_t, R_t) &= (\mu_t(s) - \mu_t(i))\frac{\pi_t(s)\pi_t(i)}{1-\delta_t\pi_t(i)} \\
\mu_t(s) &= \beta(1-\delta\pi_t(i))\left[\mu_{t+1}(s)\frac{(1-R_{t+1}\pi_{t+1}(i))}{1-\delta\pi_{t+1}(i)} + \mu_{t+1}(i)\frac{R_{t+1}\pi_{t+1}(i)}{1-\delta\pi_{t+1}(i)}\right] \\
\mu_t(i) &= \beta(1-\delta\pi_t(i))\Bigg[\mu_{t+1}(s)\frac{\delta\pi_{t+1}(s)(1-R_{t+1}\pi_{t+1}(i)) + \delta\nu\pi_{t+1}(r)}{(1-\delta\pi_{t+1}(i))^2} \\
&\qquad + \mu_{t+1}(i)\frac{\delta R_{t+1}\pi_{t+1}(s)\pi_{t+1}(i) + 1 - \delta - \gamma}{(1-\delta\pi_{t+1}(i))^2} \\
&\qquad + \mu_{t+1}(r)\frac{\gamma + \delta(1-\nu)\pi_{t+1}(r)}{(1-\delta\pi_{t+1}(i))^2} - \beta\delta v(\pi_{t+2}, \pi_{t+2})\Bigg] \\
\mu_t(r) &= \beta(1-\delta\pi_t(i))\left[\frac{\nu\mu_{t+1}(s) + (1-\nu)\mu_{t+1}(r)}{1-\delta\pi_{t+1}(i)}\right]
\end{aligned}
$$

where $\mu(s)$ and $\mu(i)$ denote respectively the discounted marginal continuation value of $\pi(s)$ and $\pi(i)$:

$$
\begin{aligned}
\mu_t(s) &= \beta(1-\delta\pi_t(i))v_{\pi(s)}(\pi_{t+1}) \\
\mu_t(i) &= \beta(1-\delta\pi_t(i))v_{\pi(i)}(\pi_{t+1}).
\end{aligned}
$$

E.3 Vaccination and Cure

The possibility of a vaccine and/or a cure is introduced by assuming that a vaccine/cure may be discovered in each and every period with a small probability $\xi > 0$. After discovery, all susceptible agents immediately move from state s to state r (recovered/immune). Hence, the state vector is immediately updated to $(0, \pi(i))$. In such a state, there remains no further externalities such that the selected infection rate $R_t = \overline{R}$, which permits the agent to reach maximal utility $\mathcal{U}_t = \overline{V}$. The Planner value at this state vector, $v^*(0, \pi_t(i))$, is then given by

$$
v^*(0, \pi_t(i)) = \pi_t(i)v_t(i) + (1-\pi_t(i))v_t(r)
$$

where $v_t(i)$ and $v_t(r)$ denote, respectively, the value of an infected and a recovered agent. Given that immunity is permanent, the value of the recovered agent, $v_t(r)$ is simply given by the discounted sum of future utility:

$$
v_t(r) = \sum_{\tau=0}^{\infty}\beta^\tau(1-\beta)\overline{V} = \overline{V}
$$

The value of the infected agent is given by

$$
v_t(i) = \underbrace{(1-\beta)\overline{V}}_{(a)} + \underbrace{\beta\gamma\overline{V}}_{(b)} + \beta(1-\gamma-\delta)\left[\frac{\gamma}{\delta+\gamma}\underbrace{\sum_{\tau=0}^{\infty}\beta^\tau(1-\beta)\overline{V}}_{(c)} + \frac{\delta}{\delta+\gamma}\underbrace{\sum_{\tau=0}^{\infty}(\beta(1-\gamma-\delta))^\tau(1-\beta)\overline{V}}_{(d)}\right]
$$

The term (a) corresponds to the instantaneous utility the infected agent gets in the period. The (b) terms accounts for the fact that with probability γ the agent will recover and therefore enjoy value $v_{t+1}(r) = \overline{V}$ as of the future period. With probability $1-\gamma-\delta$ the agent remains infected. In the future, either the agent will recover, with probability $\gamma/(\delta+\gamma)$, and enjoy the value of a recovered agent (c), or the agent will eventually die, and until she dies she will enjoy a discounted flow of utility given by (d). The value of the infected agent simplifies to

$$
v_t(i) = \frac{1-\beta(1-\gamma)}{1-\beta(1-\gamma-\delta)}\overline{V}
$$

Then, the value of the central planner in the vaccinated state is simply given by

$$
v^*(0, \pi_t(i)) = (1-\omega\pi_t(i))\overline{V} \text{ with } \omega \equiv \frac{\delta}{\frac{1-\beta}{\beta}+\gamma+\delta}.
$$

The central planner problem writes

$$
\begin{aligned}
v(\pi_t) = \max_{R_T \in [\underline{R}, \overline{R}]} & (1-\beta)\mathcal{V}^*(R_t) + \beta(1-\delta\pi_t(i))[(1-\xi)v(\pi_{t+1}) + \xi\overline{V}(1-\omega\pi_t(i))] \\
s.t. \quad & \pi_{t+1}(s) = \pi_t(s)\left(\frac{1-R_t\pi_t(i)}{1-\delta\pi_t(i)}\right) \\
& \pi_{t+1}(i) = \frac{R_t\pi_t(s)\pi_t(i) + (1-\delta-\gamma)\pi_t(i)}{1-\delta\pi_t(i)}
\end{aligned}
$$

The set of first order conditions is then given by

$$
\begin{aligned}
\mathcal{V}^{*\prime}(R_t) &= \frac{\mu_t(s)-\mu_t(i)}{1-\beta}\frac{\pi_t(s)\pi_t(i)}{1-\delta\pi_t(i)} \\
\mu_t(s) &= \beta(1-\xi)(1-\delta\pi_t(i))\left[\mu_{t+1}(s)\frac{(1-R_{t+1}\pi_{t+1}(i))}{1-\delta\pi_{t+1}(i)} + \mu_{t+1}(i)\frac{R_{t+1}\pi_{t+1}(i)}{1-\delta\pi_{t+1}(i)}\right] \\
\mu_t(i) &= \beta(1-\xi)(1-\delta\pi_t(i))\Bigg[\mu_{t+1}(s)\pi_{t+1}(s)\frac{(\delta-R_{t+1})}{(1-\delta\pi_{t+1}(i))^2} + \mu_{t+1}(i)\frac{R_{t+1}\pi_{t+1}(s)+1-\gamma-\delta}{(1-\delta\pi_{t+1}(i))^2} \\
&\qquad - \beta(1-\xi)\delta v(\pi_{t+2}) - \beta\delta\xi\overline{V}\left(1+\omega/\delta - 2\omega\pi_{t+1}(i)\right)\Bigg]
\end{aligned}
$$

where $\mu(s)$ and $\mu(i)$ denote respectively the discounted marginal continuation value of $\pi(s)$ and $\pi(i)$:

$$
\begin{aligned}
\mu_t(s) &= \beta(1-\xi)(1-\delta(\pi_t(i))\pi(i))v_{\pi(s)}(\pi_{t+1}) \\
\mu_t(i) &= \beta(1-\xi)(1-\delta(\pi_t(i))\pi(i))v_{\pi(i)}(\pi_{t+1})
\end{aligned}
$$

The equilibrium writes

$$
\begin{aligned}
v(\pi_t^k, \pi_t) = \max_{r_t \in [\underline{R}, \overline{R}]} & (1-\beta)\mathcal{V}^{eq}(r_t, R_t) + \beta(1-\delta\pi_t^k(i))[(1-\xi)v(\pi_{t+1}^k, \pi_{t+1}) + \xi\overline{V}(1-\omega\pi_t(i))] \\
s.t. \quad & \pi_{t+1}^k(s) = \pi_t^k(s)\left(\frac{1-r_t\pi_t(i)}{1-\delta\pi_t^k(i)}\right) \\
& \pi_{t+1}^k(i) = \frac{r_t\pi_t^k(s)\pi_t(i) + (1-\delta-\gamma)\pi_t^k(i)}{1-\delta\pi_t^k(i)} \\
& \pi_{t+1}(s) = \pi_t(s)\left(\frac{1-R_t\pi_t(i)}{1-\delta\pi_t(i)}\right) \\
& \pi_{t+1}(i) = \frac{R_t\pi_t(s)\pi_t(i) + (1-\delta-\gamma)\pi_t(i)}{1-\delta\pi_t(i)}
\end{aligned}
$$

The set of first order conditions, at a symmetric equilibrium, is given by

$$
\begin{aligned}
\mathcal{V}_r^{eq}(r_t, R_t) &= \frac{\mu_t(s)-\mu_t(i)}{1-\beta}\frac{\pi_t(s)\pi_t(i)}{1-\delta\pi_t(i)} \\
\mu_t(s) &= \beta(1-\xi)(1-\delta\pi_t(i))\left[\mu_{t+1}(s)\frac{(1-R_{t+1}\pi_{t+1}(i))}{1-\delta\pi_{t+1}(i)} + \mu_{t+1}(i)\frac{R_{t+1}\pi_{t+1}(i)}{1-\delta\pi_{t+1}(i)}\right] \\
\mu_t(i) &= \beta(1-\xi)(1-\delta\pi_t(i))\Bigg[\mu_{t+1}(s)\pi_{t+1}(s)\frac{\delta(1-R_{t+1}\pi_{t+1}(i)))}{(1-\delta\pi_{t+1}(i))^2} \\
&\qquad + \mu_{t+1}(i)\frac{1-\delta-\gamma+\delta R_{t+1}\pi_{t+1}(s)\pi_{t+1}(i)}{(1-\delta\pi_{t+1}(i))^2} \\
&\qquad - \beta(1-\xi)\delta v(\pi_{t+2}, \pi_{t+2}) - \beta\delta\xi\overline{V}(1+\omega/\delta - 2\omega\pi_{t+1}(i))\Bigg]
\end{aligned}
$$

where $\mu(s)$ and $\mu(i)$ denote respectively the discounted marginal continuation value of $\pi(s)$ and $\pi(i)$:

$$
\begin{aligned}
\mu_t(s) &= \beta(1-\xi)(1-\delta(\pi_t(i))\pi(i))v_{\pi(s)}(\pi_{t+1}) \\
\mu_t(i) &= \beta(1-\xi)(1-\delta(\pi_t(i))\pi(i))v_{\pi(i)}(\pi_{t+1})
\end{aligned}
$$

In the main text, we consider the case where, on average, a vaccine ought to be available within a year. In Figure 21 we report a similar exercise when the vaccine can be available within a 2 year period of time. The figure indicates that, as the horizon for a vaccine is delayed, the planner is willing to delay the peak infection rate merely by a few weeks, before "giving up hope" for an early vaccine and letting the pandemic run its course towards herd immunity.

Figure 21: Possibility of a Vaccine (2 Years)

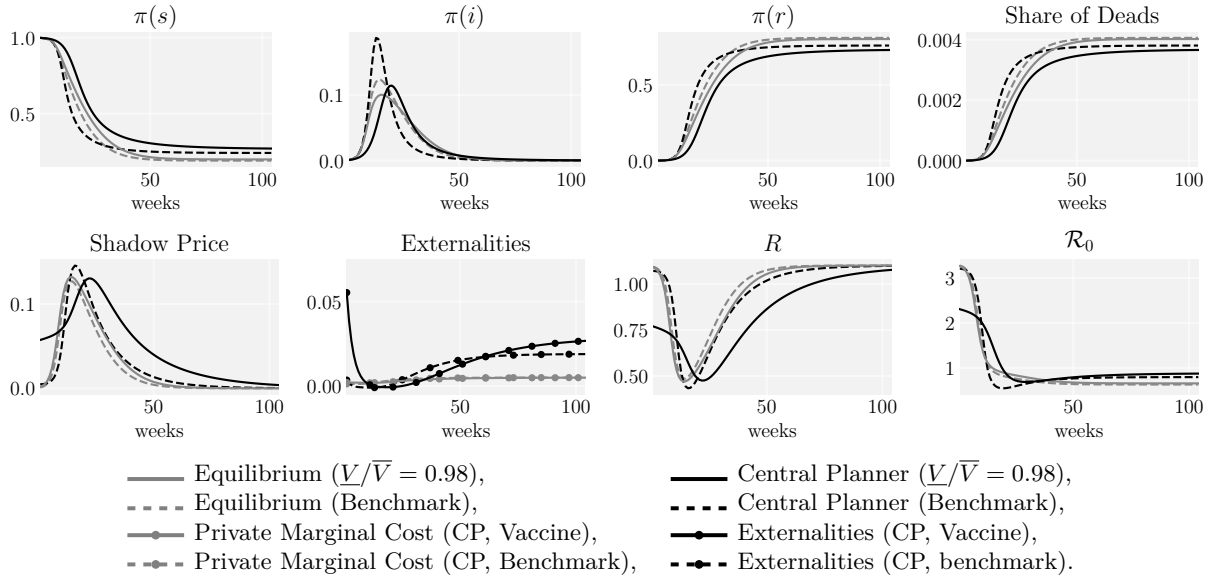

E.4 Face Masks

We now include the use of face masks in the set of static decision variables. Wearing a face mask confers no direct utility or disutility but reduces an agent's infection risk by a factor $f(m, M) \in (\underline{F}, 1]$ with $\underline{F} > 0$, where m denotes the agent's own use of masks, and M denotes aggregate mask usage. Suppose that $f(0,0) = 1$, $f_m(m, M) + f_M(m, M) \leq f_m(m, M) \leq 0$, with Inada condition $\lim_{m \to 0} f_m(m, M) = \infty$, and individual and aggregate decreasing returns:

$$-\frac{(f_m(M, M) + f_M(M, M))\, M}{f(M, M)} \leq -\frac{f_m(m, M)\, m}{f(m, M)} \leq 1$$

Mask production entails a production cost $C(M)$, where $C(M)$ displays decreasing returns to scale. We start by characterizing the consequences of the introduction of masks in the static game, and will then characterize the optimal dynamic behavior.

E.5 Static Game

Planner's solution: With face masks, the planner's within period objective is to maximize

$$\mathcal{U}(X, X) - \phi f(M, M)\, \mathcal{R}(X, X) - C(M).$$

The corresponding FOCs yield

$$\nabla \mathcal{U}(X, X) = \phi f(M, M)\, \nabla \mathcal{R}(X, X)$$
$$C'(M) = -(f_m(M, M) + f_M(M, M))\, \phi \mathcal{R}(X, X)$$

Therefore, the use of masks directly equates the social marginal rate of substitution between instantaneous utility and effective infection risks $f\mathcal{R}(X, X)$ to ϕ. When $f < 1$, this shifts X^* in the direction of maximizing instantaneous utilities. The efficient level of mask usage equates the marginal cost of masks $C'(M)$ to their marginal

benefit $-\left(f_m\left(M,M\right)+f_M\left(M,M\right)\right)\phi\mathcal{R}\left(X,X\right)$.[42] Due to a behavioral response towards utility maximization, the effective infection risk $f\mathcal{R}\left(X,X\right)$ does not decline one-for-one with face masks, but at a rate

$$-\frac{df\mathcal{R}\left(X,X\right)}{df}=\mathcal{R}\left(\phi f\right)\left(1-\mathcal{E}_R\left(\phi f\right)\right)$$

where $\mathcal{E}_R\left(\phi f\right)$ denotes the elasticity of R w.r.t. ϕ, evaluated at ϕf. Thus a fraction $\mathcal{E}_R\left(\phi f\right)\in\left(0,1\right)$ is dissipated by substitution effects. In particular, $\mathcal{E}_R\left(\phi\right)$ is inversely related to ϕ and varies from almost complete dissipation ($\mathcal{E}_R\left(0\right)=1$) when the shadow price of infection risk small, to almost no dissipation at the other extreme where $\phi\to\infty$.

As before, we can trace out a modified Pareto frontier $\widetilde{\mathcal{V}}\left(\widetilde{R}\right)$ between effective infection risk $\widetilde{R}=fR$ and instantaneous utility. Since $\widetilde{\mathcal{V}}'\left(\widetilde{R}\right)=\phi f-\phi C'\left(M\left(f\right)\right)M'\left(f\right)\leq\phi f$, the new Pareto frontier expands the set of attainable payoff and is strictly flatter than the previous one at each $\widetilde{R}$.

Equilibrium: Let P denote the consumer price of a face mask. Then individual agents maximize the following objective function, taking as given the aggregate choices (M,X):

$$\max_{m,x}\mathcal{U}\left(x,X\right)-\phi f\left(m,M\right)\mathcal{R}\left(x,X\right)-Pm$$

which yields the following first-order conditions for a symmetric equilibrium:

$$\begin{aligned}\nabla_1\mathcal{U}\left(X,X\right)&=\phi f\left(M,M\right)\nabla_1\mathcal{R}\left(X,X\right)\\P&=-f_m\left(M,M\right)\phi\mathcal{R}\left(X,X\right)\end{aligned}$$

At equilibrium, agents equate their private marginal rate of substitution to the shadow value of infection risks ϕ multiplied by $f\left(M,M\right)$, and the price of face masks internalizes the private marginal benefit of their use. The new equilibrium corresponds to $X^{eq}\left(\phi f\right)$, the effective infection risk to $f\mathcal{R}\left(X^{eq}\left(\phi f\right)\right)$.

As before, we can compute the equilibrium frontier between instantaneous utility and effective infection risks $\widetilde{R}=fR$. Because of externalities, the new equilibrium frontier is not guaranteed to be strictly higher than the old one. Figure 14 in Section 5.1 in the main text summarizes this discussion and depicts how the introduction of face masks changes the planner's and equilibrium frontiers.

Decentralization: The decentralization of the optimal policy requires the same alignment of private and social marginal rates of substitution to the new shadow value $\phi f\left(M,M\right)$, and in addition it requires a Pigouvian price subsidy for face masks to cover the positive externalities from face mask usage for others:

$$P=\left(1-s\right)C'\left(M\right),\quad\text{where}\quad s=\frac{f_M\left(M,M\right)}{f_m\left(M,M\right)+f_M\left(M,M\right)}.$$

Furthermore, we can represent the shadow price ϕ as

$$\phi=\frac{P}{-f_m\left(M,M\right)\mathcal{R}\left(X,X\right)}$$

i.e. the shadow price of infection risks is a function of the equilibrium price P and the quantity M if face masks. Holding supply of masks constant at $M=\overline{M}$, the shadow price ϕ is proportional to $P/\mathcal{R}\left(\phi f\right)$, which yields an elasticity of ϕ w.r.t. P of $\mathcal{E}_{\phi,P}|_{M=\overline{M}}=1/\left(1-\mathcal{E}_R\left(\phi f\right)\right)>1$. Controlling for supply, the shadow price of infection risk thus fluctuates more than one-for-one with the spot price for face masks. This spot price may thus offer a useful market signal to trace the dynamic evolution of infection risks.

[42] We can translate the use of face masks into a reduced form game over the choice of f, given a cost function $C\left(f\right)$. As in the base-line game, this extends the analysis to other mitigation efforts that have direct utility or monetary costs. Face masks can thus be seen as a broad stand-in for any effort that directly reduces private or aggregate infection risks.

E.6 Dynamics

The impact of face masks on the dynamics of equilibrium and optimal policy is entirely summarized by its effect on shifting the Pareto and equilibrium frontiers. Masks do not fundamentally change the results of propositions 3 and 4 reported in Section 4 in the main text, but modify two points. First, face masks give the planner and agents at equilibrium an option to push infection risks even below $\underline{R}$, thus resulting in yet faster control of the epidemic. Second, during the "Dance" phase, the use of masks serves to relax the Pareto frontier: Since for given $\pi(s)$, effective infection risk fR must stay close to $(\gamma+\delta)/\pi(s)$ during this phase, reducing f through the use of face masks allows the planner to increase R one for one, which relaxes economic restrictions and brings X closer to X^*. Therefore, face masks are a *short-run complement* to relaxing economic restrictions, since for a given state of epidemic progression and infection risk, they allow for a higher level of economic activity.

Similar arguments also apply to the equilibrium, except that here the face masks may locally depress economic activity further if the new equilibrium frontier lies below the original one due to the importance of spill-overs. On the other hand, face masks do not improve on the long-run convergence towards a full recovery with herd immunity, since incentives for mask usage will disappear once the economy approaches a complete recovery to $\overline{R}$.

Face masks also relax the long-run mortality-prosperity tradeoff. The optimal choice of $\widetilde{R} = fR^*$ shifts to

$$\frac{\widetilde{\mathcal{V}}'\left(\widetilde{R}\right)\widetilde{R}}{\widetilde{\mathcal{V}}\left(\widetilde{R}\right)} = \frac{\delta}{\widetilde{R}-\delta}.$$

Since $\widetilde{\mathcal{V}}'(R)/\widetilde{\mathcal{V}}(R) < \mathcal{V}^{*\prime}(R)/\mathcal{V}^*(R)$ for all R, the long run optimum with face masks satisfies $\widetilde{R} < R^*$. Hence at the long-run optimum, the planner transfers some of the static gains from relaxing economic restrictions due to the use of face masks back to lower long-run mortality, i.e. the long-run optimum relaxes economic restrictions less than one-for-one with the reduction in infection risks brought about by face masks.

Therefore, while face masks are strong substitutes for economic restrictions in the short run, the substitutability is weaker at longer horizons, and it may even be reverted in the very long run if by slowing infections, face masks also slow the progression of the epidemic towards herd immunity and a permanent recovery. The epidemic then takes a longer time to progress, and restrictions must thus be kept in place for longer.

To summarize, face masks facilitate the economic recovery as much as they limit new infections. The short-run substitution effects towards higher economic activity are especially important during deconfinement, i.e. for a given bound on infection risks, face masks allow a deconfinement at a higher level of economic activity than at the benchmark. In the long run, the substitutability between use of face masks and economic restrictions is weakened by substitution towards lower long-run mortality, or reversed if slower epidemic progression delays permanent recovery. In addition, they do not improve on the long-run recovery with herd immunity.

If face masks have important positive spill-overs, i.e. they protect others from being infected as much or more than they protect the person wearing a mask, then their provision may need to be subsidized, along with a mandate for their use in public spaces. At the same time, efficient management of face mask use has important side benefits: by lowering the shadow price of infection risks, face masks not only relax economic restrictions, but also reduce the scope for harmful dynamic spill-overs, and the need for other regulatory measures.

Finally, the analysis reveals a close link between the shadow price of infection risks and the price and quantity of face masks. The market for face masks may thus provide a useful market signal for tracking the shadow price of infection risks.

We now characterize the set of FOCs in the central planner allocation and in an equilibrium.

Central Planner:

$$
\begin{aligned}
v(\pi_t) = \max_{R_T \in [\underline{R}, \overline{R}]} & (1-\beta)(\mathcal{V}^*(R_t) - C(M_t)) + \beta(1-\delta\pi_t(i))[(1-\xi)v(\pi_{t+1}) + \xi\overline{V}(1-\psi\pi_t(i))] \\
s.t. \quad & \pi_{t+1}(s) = \pi_t(s)\left(\frac{1 - R_t f(M_t, M_t)\pi_t(i)}{1-\delta\pi_t(i)}\right) \\
& \pi_{t+1}(i) = \frac{R_t f(M_t, M_t)\pi_t(s)\pi_t(i) + (1-\delta-\gamma)\pi_t(i)}{1-\delta\pi_t(i)}
\end{aligned}
$$

The set of first order conditions is then given by

$$
\begin{aligned}
\mathcal{V}^{*\prime}(R_t) &= \frac{\mu_t(s) - \mu_t(i)}{1-\beta}\frac{f(M_t, M_t)\pi_t(s)\pi_t(i)}{1-\delta\pi_t(i)} \\
C'(M_t) &= -\frac{\mu_t(s) - \mu_t(i)}{1-\beta}\frac{R_t\pi_t(s)\pi_t(i)}{1-\delta\pi_t(i)})(f_1(M_t, M_t) + f_2(M_t, M_t)) \\
\mu_t(s) &= \beta(1-\xi)(1-\delta\pi_t(i))\Bigg[\mu_{t+1}(s)\frac{(1 - R_{t+1}f(M_{t+1}, M_{t+1})\pi_{t+1}(i))}{1-\delta\pi_{t+1}(i)} \\
&\qquad + \mu_{t+1}(i)\frac{R_{t+1}f(M_{t+1}, M_{t+1})\pi_{t+1}(i)}{1-\delta\pi_{t+1}(i)}\Bigg] \\
\mu_t(i) &= \beta(1-\xi)(1-\delta\pi_t(i))\Bigg[\mu_{t+1}(s)\pi_{t+1}(s)\frac{(\delta - R_{t+1}f(M_{t+1}, M_{t+1}))}{(1-\delta\pi_{t+1}(i))^2} \\
&\qquad + \mu_{t+1}(i)\frac{R_{t+1}f(M_{t+1}, M_{t+1})\pi_{t+1}(s) + 1 - \gamma - \delta}{(1-\delta\pi_{t+1}(i))^2} \\
&\qquad - \beta(1-\xi)\delta v(\pi_{t+2}) - \beta\delta\xi\overline{V}\,(1 + \psi/\delta - 2\psi\pi_{t+1}(i))\Bigg]
\end{aligned}
$$

where $f_i(M, M)$ denotes the partial derivative of $f(\cdot,\cdot)$ with respect to its i-th argument. The quantities $\mu(s)$ and $\mu(i)$ denote respectively the discounted marginal continuation value of $\pi(s)$ and $\pi(i)$:

$$
\begin{aligned}
\mu_t(s) &= \beta(1 - \delta(\pi_t(i))\pi(i))v_{\pi(s)}(\pi_{t+1}) \\
\mu_t(i) &= \beta(1 - \delta(\pi_t(i))\pi(i))v_{\pi(i)}(\pi_{t+1})
\end{aligned}
$$

Equilibrium: The equilibrium writes

$$
\begin{aligned}
v(\pi_t^k, \pi_t) = \max_{r_t \in [\underline{R}, \overline{R}]} & (1-\beta)(\mathcal{V}^{eq}(r_t, R_t) - P_t M_t) + \beta(1-\delta\pi_t^k(i))v(\pi_{t+1}^k \\
s.t. \quad & \pi_{t+1}^k(s) = \pi_t^k(s)\left(\frac{1 - r_t f(m_t, M_t)\pi_t(i)}{1-\delta\pi_t^k(i)}\right) \\
& \pi_{t+1}^k(i) = \frac{r_t f(m_t, M_t)\pi_t^k(s)\pi_t(i) + (1-\delta-\gamma)\pi_t^k(i)}{1-\delta\pi_t^k(i)} \\
& \pi_{t+1}(s) = \pi_t(s)\left(\frac{1 - R_t f(M_t, M_t)\pi_t(i)}{1-\delta\pi_t(i)}\right) \\
& \pi_{t+1}(i) = \frac{R_t f(M_t, M_t)\pi_t(s)\pi_t(i) + (1-\delta-\gamma)\pi_t(i)}{1-\delta\pi_t(i)}
\end{aligned}
$$

The set of first order conditions, at a symmetric equilibrium ($r_t = R_t$, $m_t = M_t$, $\pi_t^k(s) = \pi_t(s)$, $\pi_t^k(i) = \pi_t(i)$),

is given by

$$\begin{aligned}
\mathcal{V}_r^{eq}(r_t, R_t) &= \frac{\mu_t(s) - \mu_t(i)}{1-\beta} \frac{f(M_t, M_t)\pi_t(s)\pi_t(i)}{1-\delta\pi_t(i)} \\
p_t &= -\frac{\mu_t(s) - \mu_t(i)}{1-\beta} \frac{R_t\pi_t(s)\pi_t(i)}{1-\delta\pi_t(i)} f_1(M_t, M_t) \\
\mu_t(s) &= \beta(1-\delta\pi_t(i)) \Bigg[\mu_{t+1}(s) \frac{(1 - R_{t+1} f(M_{t+1}, M_{t+1})\pi_{t+1}(i))}{1-\delta\pi_{t+1}(i)} \\
&\qquad + \mu_{t+1}(i) \frac{R_{t+1} f(M_{t+1}, M_{t+1})\pi_{t+1}(i)}{1-\delta\pi_{t+1}(i)} \Bigg] \\
\mu_t(i) &= \beta(1-\delta\pi_t(i)) \Bigg[\mu_{t+1}(s)\pi_{t+1}(s) \frac{\delta(1 - R_{t+1} f(M_{t+1}, M_{t+1})\pi_{t+1}(i)))}{(1-\delta\pi_{t+1}(i))^2} \\
&\qquad + \mu_{t+1}(i) \frac{1 - \delta - \gamma + \delta R_{t+1} f(M_{t+1}, M_{t+1})\pi_{t+1}(s)\pi_{t+1}(i)}{(1-\delta\pi_{t+1}(i))^2} \\
&\qquad - \beta\delta v(\pi_{t+2}, \pi_{t+2}) \Bigg]
\end{aligned}$$

where $\mu(s)$ and $\mu(i)$ denote respectively the discounted marginal continuation value of $\pi(s)$ and $\pi(i)$:

$$\begin{aligned}
\mu_t(s) &= \beta(1 - \delta(\pi_t(i))\pi(i)) v_{\pi(s)}(\pi_{t+1}) \\
\mu_t(i) &= \beta(1 - \delta(\pi_t(i))\pi(i)) v_{\pi(i)}(\pi_{t+1})
\end{aligned}$$

In the decentralized equilibrium, the price is given by $P_t = (1-\mu_t)C'(M_t)$ where $\mu_t = 0$ in a standard equilibrium, while $\mu_t = f_2(M_t, M_t)/(f_1(M_t, M_t) + f_2(M_t, M_t))$ in the equilibrium that covers the positive externalities from face mask usage for others.

E.7 Testing and Contact-tracing

Here we consider the economic effects of testing and contact-tracing. By testing and quarantining anyone with a positive test result, one can reduce the number of undetected infections to

$$\widehat{\pi}(i) = \pi(i)\left(1 - \Pr(test|i)\right),$$

where $\Pr(test|i)$ denotes the fraction of infected agents that have had a positive test result and are thus in quarantine, which we interpret as a temporary exit from the game.

Adding testing and quarantines into the model comes with two challenges. First, we need to add an addition state variable $\widehat{\pi}(i)$ to keep track of the fraction of agents in quarantine, $\pi(i) - \widehat{\pi}(i)$. Second, testing alters agents' beliefs about their own health status, if they are informed of a negative test result. Hence, we need to keep track of heterogeneity across agents according to their test history. By focusing on the instantaneous propagation limit, we can side-step those two issues.[43] In this limit the infection rate converges to 0 along the path to deconfinement, and if the fraction of agents being tested does the same, then the fraction of agents who are tested is negligible and doesn't affect aggregate population dynamics. Furthermore, instantaneous propagation implies that the resolution of the quarantine phases and belief differences from past test results are very short lived, and it allows us to simplify the short-run analysis on the impact of testing on optimal policy by focusing on the policy that stabilizes the infection rate. Here, we develop the analysis by assuming that $\pi(i)$ is arbitrarily small, as in proposition 4, and then present the results for the limit in which $\pi(i) \to 0$.

[43] Berger, Herkenhoff, and Mongey (2020) show how to include additional state variables in an SIR model to capture the information generated through testing. Piguillem and Shi (2020) integrate such a structure into a simple dynamic planner's problem with capacity constraints in the medical sector, but focus on simple testing and quarantine policies. Eichenbaum, Rebelo, and Trabandt (2020b) extend their baseline model to allow for testing. Like us, these papers emphasize the potential for testing to relax untargeted quarantine measures. However, they do not analyze such measures from an optimal policy design perspective, and they do not combine testing with contact-tracing, which is key to maximize the containment potential from testing and quarantine policies.

With testing, the law of motion for $\pi(i)$ becomes

$$\begin{aligned}\pi_{t+1}(i) &= \frac{1-\gamma-\delta}{1-\delta\pi_t(i)}\pi_t(i) + \frac{\widehat{\pi}_t(i)\,\pi_t(s)\,R_t}{1-\delta\pi_t(i)} \\ &= \frac{1-\gamma-\delta}{1-\delta\pi_t(i)}\pi_t(i) + \frac{(1-\Pr(test|i))\,\pi_t(s)\,R_t}{1-\delta\pi_t(i)}\pi(i),\end{aligned}$$

augmented by the law of motion for infected agents currently in quarantine, $\pi(i)\Pr(test|i)$. Therefore, if testing reduces the fraction of infected agents in circulation by a factor $1-\Pr(test|i)$, this allows the planner to sustain the same effective infection risk $\widetilde{R}$ with R increased by a factor an offsetting factor $1/(1-\Pr(test|i)) > 0$. As with face masks, this amounts to a shift in the efficiency frontiers from $\mathcal{V}^*(R)$ to $\mathcal{V}^*\left(\widetilde{R}\right)$ and from $\mathcal{V}^{eq}(R)$ to $\mathcal{V}^{eq}\left(\widetilde{R}\right)$, where $\widetilde{R}$ denotes effective infection risks $\widetilde{R} = R/(1-\Pr(test|i))$. $\widetilde{R}$ reaches $\overline{R}$ at $R = \overline{R}(1-\Pr(test|i))$, strictly to the left of the original threshold, so testing lowers the threshold for a full recovery both at the equilibrium and the planner's solution. Testing also improves economic welfare and lowers mortality at the long-run optimum. In summary, the short-run and long-run substitution effects of testing are similar to the ones discussed above for face masks., and summarized in figure 17 of Section 5.1 in the main text.

The key is thus to raise $\Pr(test|i)$, i.e. to test and catch agents once they are infected. By Bayes' Rule, we express $\Pr(test|i)$ as

$$\Pr(test|i) = \frac{\rho\Pr(i|test)\Pr(test)}{\widehat{\pi}_t(i)}$$

Here $\Pr(test) \in (0,1)$ represents the fraction of the population that can be tested within a period, which we take as a parameter proportional to $\widehat{\pi}(i)$, and hence small - think of the ratio $\Pr(test)/\widehat{\pi}(i)$ the testing capacity relative to ongoing undetected infections. The parameterρ represents the probability of returning a positive test result from an infected agent, $1-\rho$ is the proportion of false negative test outcomes.

If tests are completely random, $\Pr(i|test) = \widehat{\pi}_t(i)$ and $\Pr(test|i) = \rho\Pr(test)$, and they identify only a small fraction $\rho\Pr(test)$ of agents who are actually infected. Testing is effective if it concentrates on "probable cases" that are most likely to return positive test results. This requires some form of tracing agents who have come in with other infected agents.

To be specific, suppose that tests can be directed towards the contacts of the most recent set of identified infections: each period, each agent randomly interacts with a finite number K of other agents and these contacts can be traced into the next period. By identifying contacts of (i) the fraction $\delta\widehat{\pi}_t(i)$ of agents who passed away most recently without being in quarantine, and (ii) a measure μ_t of agents who tested positive in the last period, we have a fraction $K(\delta\widehat{\pi}_t(i) + \mu_t)/(1-\delta\pi_t(i))$ of the population as potential test subjects. We assume that this pool exhausts the test capacity $\Pr(test)$. Each of these test candidates had a probability $\widehat{\pi}_t(i)$ of being infected before meeting one of the prior positive cases, and in turn has a probability $\pi_t(s)\,R_t/K$ of being infected at the meeting.[44] This conditional infection rate corresponds to the unconditional probability of catching an infection $\widehat{\pi}_t(i)\,\pi_t(s)\,R_t$, divided by the probability of being in contact with an infected person $K\widehat{\pi}_t(i)$. Hence they have a probability $\rho\Pr(i|test)$ of returning a positive test result, where

$$\Pr(i|test) = \widehat{\pi}_t(i)(1-\gamma-\delta) + \pi_t(s)\,R_t/K.$$

Substituting $\Pr(i|test)$ into $\Pr(test|i)$ and $\Pr(test|i)$into the law of motion for $\pi_{t+1}(i)$, we obtain the modified law of motion for the model with testing.

Now recall that at fast propagation limit, optimal policy stabilizes the proportion of infected agents. The policy that stabilizes $\pi_t(i)$ satisfies

$$\gamma + \delta - \delta\widehat{\pi}_t(i) = (1-\Pr(test|i))\,\pi_t(s)\,R_t$$

[44] This is abstracting from the possibility that these agents may simultaneously catch the infection from other sources. When $\pi_t(i)$ goes to zero, the probability that any agent incurs multiple simultaneous infections goes to 0.

Substituting $\Pr(i|test)$ into $\Pr(test|i)$ and then into this equation, we obtain a quadratic equation for $\frac{\pi_t(s)R_t}{\gamma+\delta}$, which has as a solution as $\widehat{\pi}_t(i) \to 0$:

$$\frac{\pi_t(s)R_t}{\gamma+\delta} = \frac{1-\sqrt{1-4C}}{2C}, \text{ where } C = \frac{\rho \Pr(test)}{K\widehat{\pi}_t(i)}(\gamma+\delta)$$

measures the capacity of tests to detect new infections: it is the ratio of the maximum number of positive test results, $\rho \Pr(test)$, to the undetected number of potential new cases, $K\widehat{\pi}_t(i)$, and their expected duration, $(\gamma+\delta)^{-1}$. Notice that $1 \leq \frac{\pi_t(s)R_t}{\gamma+\delta} \leq 2$, and $\frac{\pi_t(s)R_t}{\gamma+\delta}$ is increasing in C, and reaches 2 when $C = 1/4$. At that point, $\Pr(test|i) = 1/2$, so half of infected individuals are in quarantine. Testing thus improves upon the test-free stabilizing policy, and can reduce new infections by up to 50% through quarantine in this stylized example.

Testing is similar to face masks in that they both improve the static tradeoff between utility and infection risks. It will therefore also affect dynamics in a similar way: strong static substitution towards economic activity during deconfinement, weaker substitution at the long-run optimum. There is one major difference, however: testing lowers the threshold for herd immunity in a full recovery, and hence reduces long-run mortality, by lowering the threshold value of $\pi_t(s)$ at which the epidemic reaches herd immunity. This in turn will lead the planner to control the pandemic faster in the beginning. The key difference between testing and face masks is that face masks are untargeted. But as the risks of new infections subside, so do the private and social benefits from wearing masks. In contrast, testing is targeted and affects a much smaller number of agents that are traced from prior identified infections. This policy remains effective even as one approaches the limit of herd immunity.[45]

[45] Here, we have abstracted from the costs of implementing tests, but if the aggregate testing capacity has high returns at low volumes of testing, then these costs also vanish during the return to a long-run steady-state.