



HAL
open science

The Hammer and the Dance: Equilibrium and Optimal Policy during a Pandemic Crisis

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

► **To cite this version:**

Christian Hellwig, Tiziana Assenza, Fabrice Collard, Martial Dupaigne, Patrick Fève, et al.. The Hammer and the Dance: Equilibrium and Optimal Policy during a Pandemic Crisis. 2022. hal-03763122

HAL Id: hal-03763122

<https://hal.science/hal-03763122>

Preprint submitted on 29 Aug 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Hammer and the Dance: Equilibrium and Optimal Policy during a Pandemic Crisis*

Macroeconomics Group[†]
Toulouse School of Economics

May 3, 2020

Preliminary and Incomplete

Abstract

We develop a comprehensive framework for analyzing optimal economic policy during a Pandemic crisis in a dynamic economic model that trades off pandemic-induced mortality costs against the adverse economic impact of policy interventions. We use the comparison between planner problem and dynamic decentralized equilibrium to highlight the margins of policy intervention and describe optimal policy actions. As our main conclusion, we provide a strong and novel economic justification for the current approach to dealing with the Pandemic, which is different from the existing health policy rationales. This justification is based on a simple economic concept, the shadow price of infection risks, which succinctly captures the static and dynamic tradeoffs between economic prosperity and mortality risk as the Pandemic unfolds.

*We are grateful for comments from Aditya Goenka and Martin Hellwig.

[†]Christian Hellwig (corresponding author, christian.hellwig@tse-fr.eu); Tiziana Assenza, Fabrice Collard, Martial Dupaigne, Patrick Fève, Sumudu Kankanamge, Nicolas Werquin, all at Toulouse School of Economics

1 Introduction

The COVID-19 pandemic raises challenging new policy questions: How should governments manage the spread of a contagious disease? How should they weigh containing the pandemic against other policy objectives such as limiting its economic fallout? Ultimately, how should governments ponder economic prosperity and mortality risk?

Most governments have responded to COVID-19 with a combination of lockdown policies that temporarily limit economic activity to control the spread of the pandemic, and economic rescue packages to shield firms and households from negative economic spill-overs of the pandemic crisis. Yet there are important cross-country differences in the lockdown intensities, durations, and in the reliance on public enforcement or agent's private incentives. The path towards recovery raises additional questions: How should optimal de-confinement be structured? How fast should lock-down measures be lifted? How much should policy be guided by economic principles, how much by epidemiological considerations? How much can policy makers rely on private incentives, and how strong are the normative justifications for coercive policy interventions?

In this paper, we develop a framework to clarify the tradeoffs between competing economic and health policy objectives during a pandemic crisis. Two features make COVID-19 especially challenging to control: its fast propagation and the fact that many infections and transmissions are asymptomatic and go undetected. We compare the equilibrium and socially optimal policy responses to a pandemic with these two characteristics.

As our main policy conclusion, we argue that the strategies to control COVID-19 from an epidemiological perspective are also based on sound economic principles: slowing the speed of propagation has economic benefits that go well beyond the medical benefits of reducing congestion in hospitals or gaining time to develop a cure or vaccine. But these prescriptions must be qualified by the distinction between individual and common interests.

We embed a stylized economic interaction game into a dynamic S-I-R model of epidemic propagation.¹ Agents in our model interact on two levels: an economic level that determines

¹Starting with Atkeson (2020) and Eichenbaum, Rebelo, and Trabandt (2020a), a rapidly growing literature has followed this approach to study the impact of COVID-19 in quantitative economic models. We complement these quantitative studies by focusing on theoretical results and quantitative illustrations from a stylized interaction game that keeps the analysis simple, flexible and comprehensive: simple enough to capture the key policy tradeoffs, comprehensive enough to identify principles that can guide policy choices without depending too closely on model specifics, and flexible enough to provide a basis for further extensions and additional margins that are omitted from the present analysis. As a result, we are also able to unify within a single model a number of insights, observations and results that have emerged in this literature.

instantaneous utilities, and a sanitary level that determines the likelihood with which economic interactions result in the transmission of an infectious, potentially lethal disease. The economic interaction is kept deliberately stylized and only assumes that the static economic equilibrium is efficient. The model is thus general enough to encompass many textbook economic models that satisfy sufficient conditions for the 2nd Welfare Theorem, or similar principles developed for other forms of economic interactions.² We allow for multi-dimensional actions, and are thus able to address sectoral differences in equilibrium and policy responses to the pandemic crisis.

We model infections through a "confinement game", which determines individual infection probabilities as a function of individual and aggregate actions in the economic stage game. We assume that the confinement game satisfies an analogue of the static efficiency assumption: there exists an extreme confinement equilibrium that minimizes global infection risks. Any rationale for policy interventions then comes from the tradeoff between competing economic and health care objectives.

We embed this structure into a dynamic S-I-R model: agents are initially at risk (susceptible). They risk being infected by interacting with other infected agents. Once infected, they subsequently either recover or die from the disease. Recovery confers permanent immunity. We assume that agents do not know their own health state (only death is observable).³ This assumption is formally convenient (all agents are symmetric, except those who have died), and it captures the reality of asymptomatic infections in COVID-19.

By defining a *shadow price of infection risks*, we decompose the planner's problem and equilibrium characterization into a sequence of hybrid static interaction games and a recursive characterization of the dynamics of this shadow price from the S-I-R-implied population dynamics. In addition, we re-cast the dynamic problem as a reduced form interaction game with a static tradeoff between infection risk choices and instantaneous utilities.

The shadow price of infection risks summarizes the static tradeoff between current utility and future mortality that the agents or the planner face at a given point in time.⁴ Static policy tradeoffs

²By keeping the nature of economic interactions unspecified, our framework encompasses centralized and decentralized market interactions, non-market interactions in hierarchies and organizations, search and assignment markets etc., always under the assumption that the pre-pandemic equilibrium satisfies a generalized efficiency condition.

³In addition, we abstract from a symptomatic phase with illness prior to death. This is a convenient simplification - adding such a phase is certainly possible and may be of interest for many extensions, especially ones considering the role of the medical sector or the markets for medical equipment, but this is beyond the scope of this paper.

⁴If face masks are used with the sole purpose of reducing infection risks, then the spot price and equilibrium use of face masks serve as a good proxy for the shadow price of infection risks. We explore the economics of face masks, along with testing and contact-tracing, as two natural extensions of our baseline model.

(i.e. which sectors to open and which ones to close) all revolve around aligning private and social marginal rates of substitution to this shadow price. Optimal policy and equilibrium dynamics all evolve around the dynamics of this shadow price in equilibrium and at the planner’s solution. The optimal policy and dynamic equilibrium coincide, if and only if static efficiency conditions are augmented by a dynamic efficiency condition that requires offsetting static and dynamic spill-overs, a condition that is generically violated as the pandemic progresses.

The equilibrium dynamics without policy interventions are characterized by two phases: a strong initial confinement phase that brings new infections under control, and a subsequent phase of gradual deconfinement during which the epidemic slowly progresses until the population reaches a state of herd immunity. During this second phase, the pandemic is neither completely suppressed nor allowed to take off again. In reference to Pueyo (2020), who discusses exactly this approach as a possible deconfinement strategy, we call these phases *the Hammer* and *the Dance*.

The socially optimal policy follows a similar pattern of a strong confinement phase to bring new infections under control, and a subsequent phase of gradual deconfinement with slow progression towards herd immunity. However, the planner’s solution differs from the equilibrium in the timing and intensity of early lockdown measures, as well as the speed of convergence towards a long-term recovery. These differences stem from the interplay of different externalities.

In our benchmark model, we abstract from the prospect of cures and congestion in hospitals. Without these elements, the planner’s best long-term plan is to reach herd immunity quickly, but without infecting more agents than necessary. The planner thus allows the pandemic to peak early, brings it under control and then lets it progress very slowly towards herd immunity, with minimal, but long-lasting economic restrictions. This early peak in infections and mortality is optimal because preventing these infections through harder lockdown would merely postpone, but not avoid them, and prolonging the course of the pandemic carries important economic costs. Agents at equilibrium instead voluntarily opt for strong early confinement to “wait out the storm”: herd immunity is a collective good, and while everyone shares the benefits, no one is eager to help building it by catching an infection. But if everyone holds out until the worst is over, the pandemic progresses much more slowly and lasts longer, which significantly amplifies its economic costs.

What’s more, the externality changes sign during the recovery: 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 exacerbates the death toll in the long run. These dynamic infection externalities can be arbitrarily strong: Since the epidemic’s basic

reproduction coefficient is very close to 1 during deconfinement, each additional infection generates a large number of follow-up infections. Between confining too much too soon and exiting confinement much too fast, the equilibrium generates unnecessary economic hardship and a high number of avoidable deaths.

These differences become particularly striking when we let the speed of epidemic propagation tend to infinity. In this instantaneous propagation limit, the optimal policy brings the population immediately to a state that optimally balances economic prosperity and mortality in the long-run, and then stalls the pandemic at this long-run optimum without ever reaching herd immunity. In this limit, the Hammer happens instantly and the Dance lasts forever: the faster is the pandemic's natural speed of propagation, the longer it takes to reach a full recovery. In contrast, fast propagation reinforces the hold-out motive in equilibrium, resulting in a gradual propagation and long-run convergence that features neither the economic benefits of the high initial peak of infections, nor the mortality benefits of slowing down the pandemic forever.

We then add medical congestion or the prospect of a vaccine or cure to the model. These additions have only small effects on private incentives and equilibrium dynamics, but they substantially alter the planner's incentives by making the initial peak of infections very costly: the immunization externality is offset or overturned by a congestion externality in the medical sector or by the option value to delay infections in hopes for a vaccine. The planner now favors early, decisive interventions as much or even more than agents at equilibrium, and opts for a far more gradual and economically costly path to recovery in hopes of saving more lives for a longer period of time. The path to deconfinement is similar to before and still implies that agents do not sufficiently internalize infection risks for others, resulting in a faster than optimal recovery and higher than necessary mortality.

These results follow from simple cost-benefit analysis of individual and social utility-mortality tradeoffs. Economic restrictions in place today extend or save lives in the future, which explains why it is beneficial to flatten the curve for both the equilibrium and the optimal policy. The two differ by how they treat costs and benefits: in equilibrium, agents only consider the static tradeoff between instantaneous utilities and concurrent infection risks, but take as given future aggregate dynamics: hence they favor early confinement when infection risks are rising and early deconfinement when infection risks are falling. The planner instead internalizes the full dynamic consequences of current actions, from the economic benefits of building herd immunity to the costs of higher subsequent infection rates.

Faster propagation strengthens the benefits of confinement by reducing the temporal distance between instantaneous utility and future mortality. Instantaneous propagation makes this gap

disappear altogether, resulting in a quasi-static tradeoff between economic prosperity and mortality for both the planner and agents at equilibrium. At the limit, the planner's solution stalls the epidemic at precisely the resulting static optimum, while the equilibrium just offsets faster propagation through stronger hold-out incentives.

Our results identify how epidemiological factors and economic incentives jointly shape equilibrium and optimal policy responses. Conceptually our policy design problem maximizes welfare, subject to the constraints imposed by pandemic propagation. One reason why optimal policy stays close to the approach favored by epidemiologists is that even the best economic policy cannot, on its own, escape the reality of epidemic propagation: easing restrictions too fast renews the propagation of the disease, which increases the shadow price of infection risks, pushing optimal policy back towards stronger confinement measures.

While epidemiological constraints determine what the policy maker can hope to achieve, economic incentives determine the impact of these constraints on policy design. Optimal deconfinement policy eases economic restrictions under the constraint that the basic reproduction rate stays close to 1 as the infection makes its way through the population. Interventions that lower infection risk then have very strong substitution effects towards economic activity: health policy measures such as testing or use of face masks that lower the natural progression rate of the epidemic allow the government to ease economic restrictions along this transition path. Likewise, at equilibrium, the pandemic shapes agents' static tradeoffs between instantaneous utilities and infection risks, which are summarized by the equilibrium shadow price. Our analysis thus highlights the role of economic incentives and behavioral responses to the pandemic.

Related Literature: Goenka, Liu and Nguyen (2014, 2019) are among the rare papers to integrate epidemic dynamics into economic models prior to the COVID-19 pandemic, but they have focused more on the long-run consequences for growth, human capital accumulation and health policy.

Many recent papers have integrated economic decisions in an S-I-R framework. Atkeson (2020), Alvarez, Argente, and Lippi (2020), and Gonzalez-Eiras and Niepelt (2020) study the trade-offs between economic activity and infection risks that a social planner faces. Atkeson (2020) frames mitigation policies as reduced-form hump-shaped infection rates and shows that mitigation reduces and delays the peak infection rate.⁵ Alvarez, Argente, and Lippi (2020) and Gonzalez-Eiras and Niepelt (2020) characterize optimal lockdown policies on the intensive (output drop) and extensive

⁵Chang and Velasco (2020) warns against the use of estimated transition probabilities in SIRD models, since forward-looking behavior modifies the actual diffusion of the virus. As in the Lucas critique, elements of the transition matrix are not deep parameters but are function of expected policies.

(duration) margin. In these papers, the social planner internalizes that economic activity increases infection risks, and that current choices determine future dynamics of infection. Piguillem and Shi (2020) study a similar planner’s problem augmented by testing and quarantine measures.

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.⁶ 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. In these models the planner wants to mitigate sooner than private agents. Individual agents’ high activity in the short-run is socially inefficient, which we label a fatalism effect. We complement these quantitative studies with theoretical results on static and dynamic spill-overs. Among other things we highlight that optimality of early lockdowns at the planner solution is the result of medical congestion externalities present in their models, and is easily reversed when immunization spill-overs dominate the short-run policy tradeoffs.⁷

Toxvaerd (2020), Garibaldi, Moen, and Pissarides (2020), and Krueger, Uhlig, and Xie (2020) emphasize private incentives for flattening the infection curve. Toxvaerd (2020) studies the non-cooperative equilibrium of an economy with infection spillovers and shows that costly social distancing may substantially slow down the rate of new infections. Garibaldi, Moen, and Pissarides (2020) discuss social distancing based on insights from search theory and contrast static and dynamic externalities similar to our set-up. They distinguish between infection costs and immunization benefits, but stop short of a full comparison of planner solution and equilibrium. Compared to their ”bottom-up” approach to social distancing based on insights from search theory, we adopt a ”top-down” approach that abstracts from the specifics of a given interaction to focus on general principles, and we characterize the full dynamics of infection and immunization spill-overs. Krueger, Uhlig, and Xie (2020) emphasize the role of static substitution across sectors and sorting by susceptible agents

⁶The flow utility $u(a) = .25[\log a - a]$ and infection probability $\beta(\{a_i\})$ in ? can be directly mapped into the static utility and infection rate in our paper. Farboodi, Jarosch, and Shimer (2020) models the infection externality as a quadratic matching technology. They show that the laissez-faire equilibrium is closer to the optimal dynamics than to the exogenous infection SIR one. Eichenbaum, Rebelo, and Trabandt (2020a) use a consumption tax as the policy instrument. Eichenbaum, Rebelo, and Trabandt (2020b) consider testing and smart (health-status contingent) containment policies.

⁷Other dynamic externalities in the economic model, including Jones, Philippon, and Venkateswaran (2020b)’s learning-by-doing in a mitigation technology (working from home) also modify the planner’s incentives to reduce activities in the short run.

into low risk activities. As in our setup, a planner would subsidize (or tax) certain sectors according to their specific infection risk externalities. Compared to these papers, our analysis highlights that private incentives for social distancing may go too far, if the planner values early immunization.

To conclude we highlight two distinctive features of our work relative to these concurrent COVID-19 papers. First, we complement these primarily quantitative studies by proposing a theoretical framework and results that shed light on different forces in play. Our model unifies a number of contrasting results which follow from specific assumptions about static and dynamic externalities, as described in the preceding paragraphs. Second, we assume that health status is unobservable. Most of the existing studies assume that health status is observable to the individual if not the planner, yet then focus on simple policies that do not condition on this information.⁸ 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, we are able to directly address policies that are designed to deal with this lack of information during the COVID-19 pandemic.

2 Setting the stage

We consider a dynamic game in which a mass Λ_t of agents interacts in each period $t = 1, 2, \dots$ in an *economic stage game* which determines their instantaneous utilities or payoffs each period. But these decisions and economic interactions with their peers also expose agents to the risk of being infected with and potentially dying from an infectious disease.

Hence, we start by juxtaposing the economic stage game with a *confinement stage game*, which summarizes how the agents' decisions determine their risk of infection. In the dynamic game, we will represent the strategic interaction in each period as a hybrid of these two stage games, with a weight on minimizing infection risks that varies with the concurrent prevalence of infections.

Actions: Let $\mathcal{X} \subset \mathbb{R}^K$ be a compact, convex set of feasible economic actions or choices, with non-empty interior $\text{int}(\mathcal{X})$. Let $x \in \mathcal{X}$ denote an individual action, and $X \in \mathcal{X}$ the aggregate choice of the other agents.⁹ 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.

⁸Piguillem and Shi (2020) or Eichenbaum, Rebelo, and Trabandt (2020b) introduce lack of observability but focus mainly on optimal testing policies.

⁹To simplify exposition and notation, we will focus throughout on symmetric pure strategy profiles and equilibria. We will use x and X to draw the distinction between individual choices and aggregate variables.

Allowing for multi-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 drawing a distinction between essential and non-essential sectors, and when discriminating against sectors that increase infection risk, such as mass transport, travel or entertainment events.

Economic stage game: Let $\mathcal{U}(\cdot) : \mathcal{X} \times \mathcal{X} \rightarrow [0, \bar{V}]$ denote the static flow utility of choosing $x \in \mathcal{X}$, when all other agents choose X . $\mathcal{U}(\cdot)$ is continuous, strictly concave and twice continuously differentiable over the interior of $\mathcal{X} \times \mathcal{X}$. Let $\mathcal{V}(\cdot) : \mathcal{X} \rightarrow [0, \bar{V}]$ denote the value of making the same choice as all other agents: $\mathcal{V}(X) = \mathcal{U}(X, X)$. We assume the following about $\mathcal{U}(\cdot)$ and $\mathcal{V}(\cdot)$:

Assumption 1: *There exists X^* , such that $\mathcal{V}(X^*) = \bar{V}$.*

Assumption 1 states that the agent's utility in the static game is maximized at an interior optimum X^* . Moreover, since $\mathcal{U}(x, X^*) \leq \bar{V} = \mathcal{V}(X^*)$ for all $x \in \mathcal{X}$, X^* also represents a symmetric Nash equilibrium of the static game in which each agent chooses $x \in \mathcal{X}$, i.e. the symmetric Nash equilibrium X^* decentralizes the utilitarian planner's solution, which represents the economic best-case scenario. 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 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 R denote the probability with which an agent is infected within a given period, conditional on being susceptible to infection. We assume that R 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 $R(x, X) \cdot \pi(i)$, where $R(\cdot) : \mathcal{X} \times \mathcal{X} \rightarrow [\underline{R}, 1]$ 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 $\mathcal{R}(X) \cdot \pi(i)$, where $\mathcal{R}(X) = R(X, X)$. We make the following additional assumption about $R(\cdot)$ and $\mathcal{R}(\cdot)$:

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

Assumption 2 states that the agent's infection rate R is minimized at an interior optimum $\hat{X} \neq X^*$, and since $R(x, \hat{X}) \geq \mathcal{R}(\hat{X}) = \underline{R}$, 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 $R(x, X)$. We interpret \hat{X} as the "extreme confinement equilibrium",

which is best-case scenario from a health policy perspective.¹⁰ In the dynamic model, it is also the action that maximizes the long-term survival rate within the population. The assumption $\hat{X} \neq X^*$ generates a conflict between maximizing economic well-being $\mathcal{V}(\cdot)$ and minimizing infection risk $\mathcal{R}(\cdot)$ that is at the core of our analysis.

We let $\bar{R} = \mathcal{R}(X^*) > \underline{R}$ denote the infection risk at the economic optimum and $\underline{V} = \mathcal{V}(\hat{X}) \in (0, \bar{V})$ denote the instantaneous utility at the extreme confinement equilibrium. Any collective action with a strictly higher infection risk than \bar{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. that instantaneous utilities $\mathcal{U}(\cdot)$ and infection probabilities $R(\cdot)$ are independent of the mass of participating players Λ . 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.

3 Economic well-being vs. infection risk: Static tradeoffs

Consider now a hybrid stage game in which agents' payoffs are given by

$$\mathcal{U}(x, X) - \phi R(x, X),$$

where ϕ represents a shadow price associated with infection risk that measures its importance in agents' decisions relative to economic well-being: At $\phi = 0$, we recover the economic stage game, and when ϕ converges to ∞ the hybrid stage game converges to the confinement game.

The planner's solution $X^*(\phi)$ within this hybrid stage game satisfies

$$\nabla \mathcal{V}(X) = \phi \nabla \mathcal{R}(X).$$

We recover the standard result: the planner's solution equates the marginal rates of substitution in $\mathcal{V}(\cdot)$ to the marginal rates of substitution in $\mathcal{R}(\cdot)$, i.e. the planner equates the marginal tradeoffs in instantaneous utility and infection rates when substituting between different dimensions of X .

The Nash equilibrium $X^{eq}(\phi)$ of the hybrid game satisfies

$$\nabla_1 \mathcal{U}(X, X) = \phi \nabla_1 R(X, X),$$

¹⁰For 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.

where ∇_1 denotes the gradient with respect to the individual action x . Thus the same margin of substitution between different dimensions operates in the private decisions, but the tradeoffs that individual agents are facing may be different from the ones faced by the planner.

Proposition 1 asks under what condition the planner's solution coincides with the Nash equilibrium of the hybrid game:

Proposition 1 *The solution $X^*(\phi)$ to the planner's problem is a Nash equilibrium of the hybrid game indexed by ϕ , if and only if*

$$\nabla_2 \mathcal{U}(X^*(\phi), X^*(\phi)) = \phi \nabla_2 R(X^*(\phi), X^*(\phi)),$$

where ∇_2 denotes the gradient with respect to the aggregate choice X .

Proposition 1 follows immediately from observing that $\nabla \mathcal{V}(X) = \nabla_1 \mathcal{U}(X, X) + \nabla_2 \mathcal{U}(X, X)$ and $\nabla \mathcal{R}(X) = \nabla_1 R(X, X) + \nabla_2 R(X, X)$. $\nabla_2 \mathcal{U}(X, X)$ measures the utility spill-overs from the aggregate choice towards any individual agent, holding constant that agent's decision. $\nabla_2 R(X, X)$ measures the infection rate spill-overs from the aggregate choice X , or infection externalities from the other agents, scaled by the shadow price ϕ of infection risks. This proposition states that the decentralization result extends from the two stage games to the hybrid game if and only if utility spill-overs of individual choices are exactly offset by infection risk spill-overs at $X^*(\phi)$.

The necessary and sufficient condition in proposition 1 states that

$$\left. \frac{\frac{\partial \mathcal{U}(x, X)}{\partial X_i}}{\frac{\partial R(x, X)}{\partial X_i}} \right|_{x=X=X^*(\phi)} = \phi$$

for any dimension of activity i . This condition says that along all dimensions of X , the marginal rate of substitution between utility and infection risk spill-overs, which we will call the *spill-over marginal rate of substitution*, must be the same, and equal to the shadow price of infection risk: a marginal change of activity in any two sectors leads to the same marginal tradeoff between instantaneous utility and infection risk spill-overs as the one imposed in the planner's objective. Coupled with the first-order conditions for the planner's problem and the equilibrium, we obtain:

$$\left. \frac{\frac{\partial \mathcal{U}(x, X)}{\partial X_i}}{\frac{\partial R(x, X)}{\partial X_i}} \right|_{x=X=X^*(\phi)} = \left. \frac{\frac{\partial \mathcal{U}(x, X)}{\partial x_i}}{\frac{\partial R(x, X)}{\partial x_i}} \right|_{x=X=X^*(\phi)} = \left. \frac{\frac{\partial \mathcal{V}(X)}{\partial X_i}}{\frac{\partial \mathcal{R}(X)}{\partial X_i}} \right|_{x=X=X^*(\phi)} = \phi,$$

i.e. the planner's solution is an equilibrium, if and only if private, spill-over, and social marginal rates of substitution are all equal to the shadow price ϕ . This condition is equivalent to stating that $X^*(\phi)$ is a global maximizer of $\mathcal{U}(x, X) - \phi R(x, X)$ with respect to both of its arguments,

which is the generalization of assumptions 1 and 2 in the economic and infection risk stage games. Proposition 1 thus states that the planner's solution can be decentralized if and only if the spill-over marginal rate of substitution is aligned with private and social marginal rates of substitution at the planner's solution.

We do not need to take a stance on the direction of these 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. Proposition 1 shows that the planner doesn't weigh them in terms of their absolute, but their relative strengths.

Policy Implications: What does this result tell us about optimal policy? Consider a planner who can impose restrictions $\hat{\mathcal{X}} \subset \mathcal{X}$ on the choice sets of agents to bring the equilibrium in line with the planner's solution.

The main policy insight is that such restrictions must serve to equate private and social marginal rates of substitution with each other and with ϕ . The efficiency condition suggests that we should restrict activities in which the spill-over marginal rate of substitution is lower than ϕ , and conversely subsidize or protect activities in which this marginal rate of substitution in spill-overs is higher than ϕ , i.e. economic spill-overs are large relative to infection risk spill-overs. Moreover, the further the spill-over marginal rate of substitution is from ϕ , the larger the intervention should be. Hence, activities in which the marginal infection risk spill-overs are very high, relative to their marginal economic spill-overs, such as socializing, going out to restaurants, entertainment events (large scale concerts or sports events) or inessential long-distance travel, should be the most heavily restricted at any point in time, while activities that generate important economic spill-overs but whose infection risk spill-overs are not too large should be subsidized. This is not a statement about the absolute magnitude of the spill-overs but about their relative magnitudes: there may be sectors, like groceries or healthcare, that have substantial infection risk spill-overs but keeping them open is justified by the important positive economic spill-overs of their activity. A similar argument may apply to public education, if the positive economic spill-overs associated with public education exceed the negative infection risk spill-overs through children at school.

Second, the proposition informs us how these policy restrictions change as we vary ϕ , the shadow price of infection risk. When ϕ increases, the optimality conditions imply that agents are privately and socially more inclined to give up utility to control infection risk. Proposition 1 then states that the equilibrium is efficient if and only if the spill-over marginal rates of substitution scale one

for one with the private and social marginal rates of substitution. Hence, restrictions should be eased or even reversed on activities whose spill-over marginal rate of substitution changes less than proportionally with the private marginal rate of substitution at $X^*(\phi)$, and restrictions should be introduced or tightened on activities whose spill-over MRS changes more than one-for-one with the private MRS at $X^*(\phi)$. Importantly, an increase in ϕ will not automatically result in an across the board tightening of restrictions: since agents already have a private incentive to change to the increase in ϕ , all sectors will simultaneously shift towards higher private marginal rates of substitution, so the key question for tightening or relaxing restrictions on any given activity is whether the spill-over MRS changes more or less than the private MRS.

These insights also inform us about optimal strategies for deconfinement, the periods in which the shadow price of infection risk ϕ converges back to 0: suppose that we can order sectors by their spill-over marginal rate of substitution. Restrictions and subsidies should then be lifted from "the center" to the "extremes", starting with those sectors whose spill-over MRS are the closest to the current shadow price of infection risk, and gradually expanding outwards.¹¹

Third, the absolute magnitude of spill-overs matters for the urgency of intervention in each sector. Certain activities may rank poorly in terms of their spill-over marginal rate of substitution, but because both economic and infection risk spill-overs are small in absolute values, the private and social marginal rates of substitution remain closely aligned with each other.

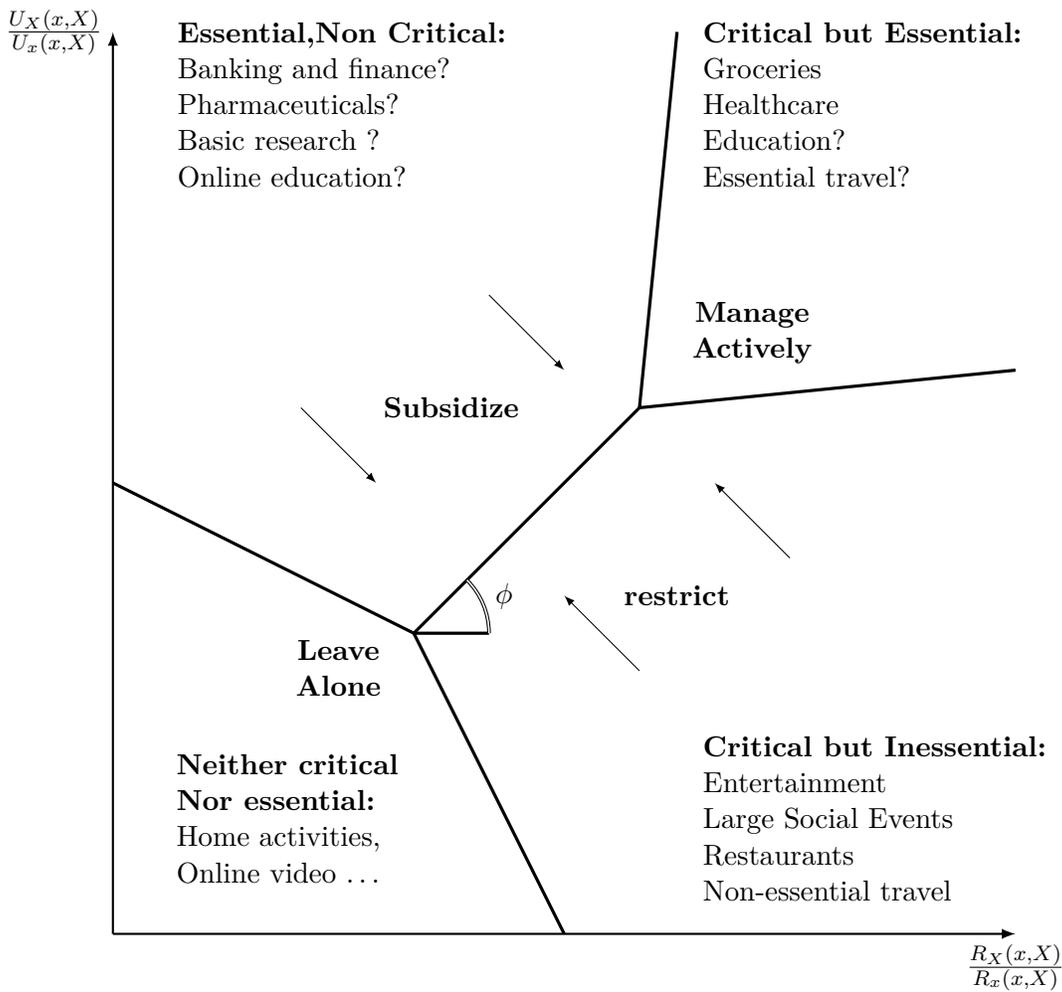
To summarize, simple but sound policy advice consists of the following points which are illustrated by Figure 1.

(i) Restrict activities that generate strong infection risk externalities but weak economic externalities, but protect or subsidize *essential* economic activities that have strong positive economic spill-overs especially if they have weak infection risk externalities,

(ii) Very carefully manage activities that are both economically essential and *critical* from an

¹¹The shadow price of infection risk can also be used to think about re-organizing activities within sectors or even organizations. For example, face-to-face interactions may generate a natural trade-off between productivity and infection risks that can also be captured by the tradeoff between \mathcal{U} and R . Agents then modulate face-to-face interactions with their clients to internalize the private marginal rate of substitution between productivity and infection risks while the planner also internalizes the spill-over marginal rates of substitution. These spill-overs are bound to be important in places that naturally generate additional face-to-face contacts with third parties - such as large workplaces, public spaces like universities, or open public spaces. Whether the resulting private incentives to mitigate infection risks are also sufficient from a social point of view then really comes down to the relative spill-over effects: Our principle that efficiency requires alignment of private and social marginal rates of substitution and the ensuing implications apply very broadly - broadly speaking to any decision margin under the control of private agents that results in a tradeoff between economic outcomes (utility or productivity) and infection risk.

Figure 1: Simple Policy Advice



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 tradeoff is resolved efficiently.

(iii) Interventions should scale with the magnitude of ϕ , the shadow price of infection risk, and be largest in those dimensions of activities whose spill-overs are furthest from ϕ . If ϕ scales with the fraction of infected agents in the population, this also shows that efficient interventions must happen fast, and may require almost day-to-day management during the onset of a pandemic like COVID-19, whose initial growth has been extremely fast.

A Useful Decomposition: In the remainder of this section, we decompose the planner's solution and equilibrium characterization into implementation rules $X^*(\cdot)$ and $X^{eq}(\cdot)$ that determine X^* and X^{eq} as functions of the implemented infection risk R , and a reduced form interaction game over the choice of R . These functions further describe Pareto and equilibrium frontiers $\mathcal{V}^*(R)$ and $\mathcal{V}^{eq}(R)$ between instantaneous utilities and infection risks for the planner and for agents at equilibrium. These decompositions allow us to study dynamic optimal policy and equilibrium through the lens of a reduced-form tradeoff, treating infection risks R as the policy variable. We then provide necessary and sufficient conditions for decentralization of the planner's solution in terms of this reduced-form representation.

As ϕ varies from 0 to ∞ , the planner's solution in the hybrid game traces out a Pareto frontier between $\mathcal{V}(X^*(\phi))$ and $\mathcal{R}(X^*(\phi))$, and inverting $\mathcal{R}(X^*)$, we find the implementation rule $X^*(R)$, the Pareto frontier $\mathcal{V}^*(R) = \mathcal{V}(X^*(R))$, and the shadow price function $\phi^*(R) = \mathcal{V}^{*'}(R)$. $\mathcal{V}^*(R)$ is strictly increasing, concave and satisfies the Inada conditions $\mathcal{V}^*(\bar{R}) = \bar{V}$, $\mathcal{V}^*(\underline{R}) = \underline{V}$, $\mathcal{V}^{*'}(\bar{R}) = 0$ and $\lim_{R \rightarrow \underline{R}} \mathcal{V}^{*'}(R) = \infty$.¹²

Likewise, the equilibrium action $X^{eq}(\phi)$ in the hybrid game determines the equilibrium infection risk $\mathcal{R}(X^{eq}(\phi))$. Inverting the latter, we obtain the equilibrium implementation rule $X^{eq}(R)$, from which we define the equilibrium efficiency frontier $\mathcal{V}^{eq}(R) = \mathcal{V}(X^{eq}(R))$. Notice that $\mathcal{V}^{eq}(\cdot) \leq \mathcal{V}^*(\cdot)$, and they are equal, if and only if $X^{eq}(\cdot) = X^*(\cdot)$, i.e. if and only if the equilibrium implementation rule is efficient (which holds automatically if $K = 1$).

Now, fix R and $X^{eq}(R)$ and consider the agents' problem of maximizing $\mathcal{U}(x, X^{eq}(R))$ subject to an upper bound constraint on $R(x, X^{eq}(R))$:

$$\mathcal{U}^{eq}(r, R) = \max_{x \in \mathcal{X}} \mathcal{U}(x, X^{eq}(R)), \text{ subject to } R(x, X^{eq}(R)) \leq r.$$

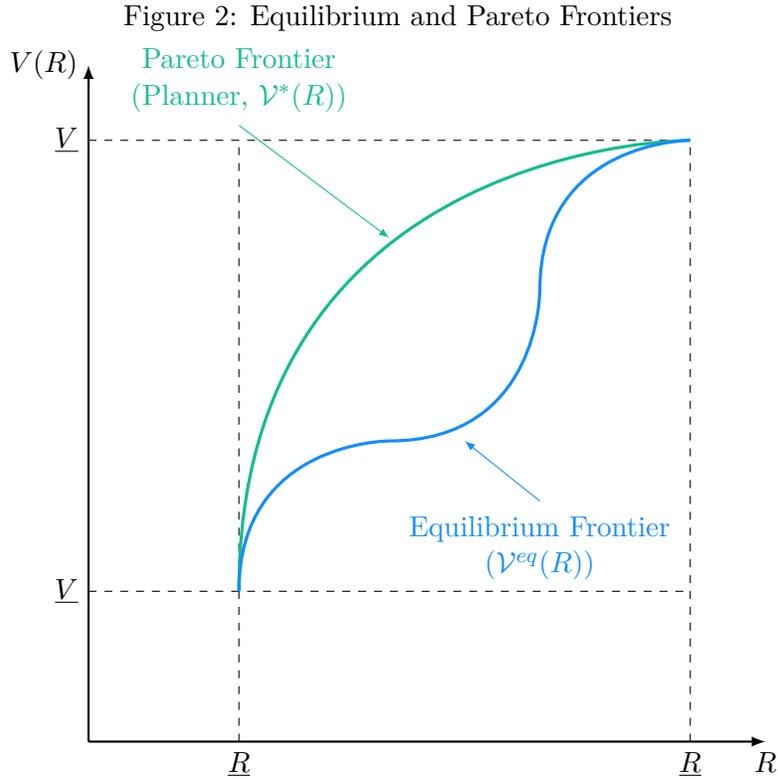
This reduced-form utility function $\mathcal{U}^{eq}(r, R)$ is strictly increasing, concave, and satisfies the Inada conditions at $\bar{R}(R) = \max_{x \in \mathcal{X}} R(x, X^{eq}(R)) \geq \mathcal{R}(X^{eq}(R))$ and $\underline{R}(R) = \min_{x \in \mathcal{X}} R(x, X^{eq}(R)) \leq$

¹²Alternatively, $\mathcal{V}^*(R)$ and $X^*(R)$ are given by the constrained planner's problem $\mathcal{V}^*(R) = \max_{X \in \mathcal{X}, R \geq \mathcal{R}(X)} \mathcal{V}(X)$.

$\mathcal{R}(X^{eq}(R))$ whenever $\arg \min_{x \in \mathcal{X}} R(x, X^{eq}(R)) \in \text{int}(\mathcal{X})$.¹³ Hence if infection rates 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 . This allows us to reduce the hybrid game to a one-dimensional interaction game in the choice of R , with reduced-form payoffs

$$\mathcal{U}^{eq}(r, R) - \phi \cdot r.$$

In addition, $\mathcal{V}^{eq}(R) = \mathcal{U}^{eq}(R, R)$, and the equilibrium shadow price function is $\phi^{eq}(R) = \mathcal{U}_r^{eq}(R, R)$.



The characterization of $\mathcal{V}^*(\cdot)$ (in green) and $\mathcal{V}^{eq}(\cdot)$ (in blue) is summarized in Figure 2. The general shape of $\mathcal{V}^{eq}(\cdot)$ depends on economic and infection risk spill-overs. However, $\mathcal{V}^{eq}(\cdot)$ satisfies the same Inada conditions as $\mathcal{V}^*(R)$ at \bar{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 inside $\mathcal{V}^*(R)$ at its boundaries.

¹³When $K = 1$, the solution to the constraint optimization problem is determined directly from the constraint: $\mathcal{U}(r, R) = \mathcal{U}(x(r, R), X^{eq}(R))$, where $x(r, R)$ solves $r = R(x, X^{eq}(R))$.

Proposition 2 provides necessary and sufficient conditions for decentralizing the planner's solution $X^*(R)$ as an equilibrium of the reduced-form hybrid game. This proposition combines efficient implementation ($X^{eq}(R) = X^*(R)$ and $\mathcal{V}^{eq}(R) = \mathcal{V}^*(R)$) with an additional condition that decentralizes the planner's choice of R .

Proposition 2 *For any $R \in [\underline{R}, \overline{R}]$, $X^*(R)$ is implemented in a Nash equilibrium of the hybrid game with $\mathcal{U}^{eq}(\cdot)$, if and only if (i) $\mathcal{V}^{eq}(R) = \mathcal{V}^*(R)$ and (ii) $\mathcal{U}_R^{eq}(R, R) = \phi^*(R) - \phi^{eq}(R)$.*

This proposition gives two necessary and sufficient conditions for efficiency in the reduced form game: (i) Efficient implementation and (ii) "Offsetting spill-overs": any marginal spill-over from R in the reduced form marginal utility function $\mathcal{U}_R^{eq}(R, R)$ must be matched at equilibrium by an offsetting spill-over in shadow values $\phi^*(R) - \phi^{eq}(R)$. We call $\phi^* - \phi^{eq}$ the *dynamic spill-over*, as opposed to the static spill-over captured by $\mathcal{U}_R^{eq}(R, R)$, since the shadow values of infection risk are derived from the planner's and agents' discounted continuation values in the dynamic model.

The planner's implementation rule $X^*(\cdot)$ can be globally decentralized, if and only the efficient implementation and offsetting spill-overs conditions hold for all $R \in [\underline{R}, \overline{R}]$. Assuming that the offsetting spill-overs condition holds globally is extremely stringent: when ϕ^* and ϕ^{eq} are endogenously determined by the dynamics of infection, this condition requires that dynamic spill-overs are identical for any two states that lead to the same policy choice R .

Condition (ii) is necessary since $\mathcal{V}^{eq'}(R) = \mathcal{U}_R^{eq}(R, R) + \phi^{eq}(R) \neq \phi^*(R) = \mathcal{V}^{*'}(R)$, together with $\mathcal{V}^{eq}(R) = \mathcal{V}^*(R)$ implies that $\mathcal{U}^{eq}(R \pm \varepsilon, R) > \mathcal{V}^*(R)$ for some small perturbation ε , contradicting that R was an equilibrium of the reduced form game. Moreover, Conditions (i) and (ii) together imply that $\phi^{eq}(R) = \mathcal{V}^{*'}(R) - \mathcal{U}_R^{eq}(R, R) = \mathcal{V}^{eq'}(R) - \mathcal{U}_R^{eq}(R, R) = \mathcal{U}_r^{eq}(R, R)$, and hence that R is implemented in a Nash equilibrium of the hybrid game, along with $X^{eq}(R) = X^*(R)$.

To re-cap, this section has accomplished three objectives:

First, we have formulated the tradeoff between economic prosperity and control of an infectious disease as a static interaction game with competing objectives of maximizing contemporaneous utility and minimizing future infection risk. In doing so we have introduced a hybrid game which weighs the competing objectives by a shadow price on infection risk.

Second, we have given a necessary and sufficient condition for efficiency of the equilibrium of the hybrid game. This condition is much more restrictive than our baseline assumptions of efficiency at equilibrium for the two benchmark games.

Third, we have decomposed the equilibrium interaction in the hybrid game into a reduced form interaction game in infection rates along with an implementation rule that determines the

equilibrium or planner's optimal action for a given targeted infection risk R , and then expressed the conditions for efficiency in terms of reduced-form tradeoffs between instantaneous utility and infection risks at the Pareto frontier and in equilibrium.

These results summarize static tradeoffs between economic activity and infection risks, and they also set the stage for the dynamic model. By decomposing the characterization of the planner's problem and the equilibrium in each period into a stationary implementation rule and a reduced form interaction game between instantaneous utility and infection risks, we are able to analyze the dynamic model recursively as a sequence of hybrid stage games with given reduced form payoffs $\mathcal{V}^*(R)$ for the planner's problem and $\mathcal{U}^{eq}(r, R)$ for the equilibrium, augmented by shadow prices ϕ^* and ϕ^{eq} that summarize the planner's and agents' concern about their future and that are endogenously linked to the dynamics of infection risks. We can therefore treat infection risk R as our basic choice variable in the dynamic model and compare planner's problem and dynamic equilibrium through the lens of static reduced form utilities $\mathcal{V}^*(R)$ and $\mathcal{U}^{eq}(r, R)$ and the dynamics of shadow prices ϕ^* and ϕ^{eq} .

4 Economic well-being vs. infection risk: Dynamic tradeoffs

We now consider a dynamic game with an unfolding epidemic. The economic stage game is infinitely repeated among a mass Λ_t of agents who remain alive in period t . The epidemic is summarized by a simple S-I-R structure: 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.

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 assume perfect foresight, i.e. despite idiosyncratic uncertainty about infection incidence, aggregate population shares of the different types are perfectly predictable. We focus on a symmetric equilibrium, in which all agents take the same equilibrium action.

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 $\pi = \left(\pi_t(s) \quad \pi_t(i) \right)'$, and the equilibrium as a Markov-perfect equilibrium in π . The vector π admits the representation

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

where $R_t = R(x_t, X_t) \cdot \pi_t(i)$ denotes the probability with which an agent is infected in period t , as described above for the confinement game. 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 .

An agent's expected discounted utility flow is

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

where $\Lambda_t(x^{t-1}, X^{t-1})$ is the probability that the agent survives to period t , which is a function of the initial distribution π_0 , individual and aggregate choices (x^{t-1}, X^{t-1}) up to period $t - 1$, and $\beta \in (0, 1)$ is the time discount factor. This welfare criterion summarizes the dynamic tradeoff between instantaneous utilities $\mathcal{U}(x_t, X_t)$ and survival probabilities $\Lambda_t(x^{t-1}, X^{t-1})$.

A symmetric Nash equilibrium in the dynamic game is a sequence of 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.

Dynamic planner problem: The utilitarian social planner's objective is

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

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

$$v^*(\pi) = \max_{X \in \mathcal{X}} \{ (1 - \beta) \mathcal{V}(X) + \beta (1 - \delta \pi(i)) v^*(\pi_{+1}) \}$$

$$\text{where } \pi_{+1} = (1 - \delta \pi(i))^{-1} \cdot T(\mathcal{R}(X) \pi(i)) \cdot \pi$$

We let $X^*(\pi)$ denote the corresponding social planner's decision rule.

In the dynamic model, current choices affect instantaneous utilities directly, and continuation values indirectly through their effect on the resulting infection rate $R(x, X)$. Making use of the observations from the previous section, we restate this planner's problem as a choice over R :

$$v^*(\pi) = \max_{R \in [\underline{R}, \bar{R}]} \{(1 - \beta) \mathcal{V}^*(R) + \beta (1 - \delta \pi(i)) v^*(\pi_{+1})\}$$

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

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

Hence we decompose the planner's decision rule $X^*(\pi)$ into a target infection rate $R^*(\pi)$ and the static implementation rule $X^*(R)$ for a given target R that we derived in the previous section.

Since R affects π_{+1} linearly as a one-for-one increase in $\pi(i)$ and reduction in $\pi(s)$, we can represent the planner's optimal choice through the planner's shadow price of infection risk $\Phi^*(\pi)$:

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

The planner's shadow price of infection risk is equal to the discounted marginal social cost of an additional infection, scaled by 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. $\Phi^*(\cdot)$ is a function of the current state π .

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 posterior about her own infection state. The probability that the agent survives until next period is $1 - \delta \pi^k(i)$.¹⁴ Her decision problem is stated as follows

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

$$\text{where } \pi_{+1}^k = (1 - \delta \pi^k(i))^{-1} \cdot T(R(x, X(\pi)) \pi_t(i)) \cdot \pi^k$$

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

An aggregate choice function $X^{eq}(\cdot)$ is a symmetric Markov-perfect equilibrium if given an initial private belief $\pi^k = \pi$, $X^{eq}(\cdot)$ is a best response to itself.

¹⁴Individual survival probabilities evolve recursively according to $\Lambda_{t+1}(x^t, X^t) = \Lambda_t(x^{t-1}, X^{t-1}) \cdot (1 - \delta \pi_t^k(i))$, or $\Lambda(\pi^k) = \gamma / (\gamma + \delta (1 - \pi_t^k(i) - \pi_t^k(s)))$

We similarly decompose the Markov-perfect equilibrium characterization into a static implementation rule $X^{eq}(R)$ that implements R as the equilibrium choice in the static hybrid game, and a reduced form dynamic interaction that determines the equilibrium infection rate $R^{eq}(\pi)$. Restating the agent's decision problem as a choice over $r \in [\underline{R}, \bar{R}]$, we obtain:

$$\begin{aligned} \hat{v}(\pi^k, \pi; R(\cdot)) &= \max_{r \in [\underline{R}, \bar{R}]} \left\{ (1 - \beta) \mathcal{U}^{eq}(r, R) + \beta (1 - \delta \pi^k(i)) \hat{v}(\pi_{+1}^k, \pi_{+1}; R(\cdot)) \right\} \\ \text{where } \pi_{+1}^k &= (1 - \delta \pi^k(i))^{-1} \cdot T(r \pi_t(i)) \cdot \pi^k \\ \pi_{+1} &= (1 - \delta \pi(i))^{-1} \cdot T(R(\pi) \pi(i)) \cdot \pi \\ \text{and } \mathcal{U}^{eq}(r, R) &\equiv \max_{x \in \mathcal{X}, R(x, X^{eq}(R)) \leq r} \mathcal{U}(x, X^{eq}(R)) \end{aligned}$$

The function $\mathcal{U}^{eq}(r, R)$ is the reduced-form indirect utility of choosing r when the aggregate action $X^{eq}(R)$ implements an equilibrium infection rate R . The equilibrium target infection rate $R^{eq}(\cdot)$ is a fixed point to the best response correspondence that is associated with this value function.

Taking first-order conditions, exploiting the linearity of continuation values with respect to r , and evaluating at $\pi^k = \pi$, we obtain the equilibrium shadow price of infection risk $\Phi^{eq}(\cdot)$:

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

The equilibrium shadow price weighs the discounted private marginal cost of being infected by the probability with which the agent privately risks being infected, evaluated at $\pi^k = \pi$. The latter multiplies the aggregate infection rate $\pi(i)$ with the individual probability of being susceptible $\pi^k(s) = \pi(s)$. The difference between the private and social shadow price comes down to the private and social marginal costs of an infection. They differ because at equilibrium the agent does not internalize that becoming infected increases the risks of subsequent infection for other agents.

Efficient Decentralization: By Proposition 2, the planner solution is decentralized as a Markov-perfect equilibrium if and only if static spill-overs exactly offset dynamic spill-overs at $R = R^*(\pi)$, or $\Phi^*(\pi) - \Phi^{eq}(\pi) = \mathcal{U}_R^{eq}(R^*(\pi), R^*(\pi))$ for all π . At equilibrium, agents internalize neither the static spill-overs their choice of r has on their peers, nor the dynamic spill-overs of their choices through higher future infection rates, but if (and only if!) the two just offset each other they have no incidence on the equilibrium outcome. In addition to the static efficiency condition $\mathcal{U}^{eq}(R, R) = \mathcal{V}^*(R)$, we require an equalization of private, social and spill-over marginal rates of substitution between current utility and continuation values at all states:

$$\frac{\mathcal{V}^{*'}(R^*(\pi))}{\Phi^*(\pi)} = \frac{\mathcal{U}_r^{eq}(R^*(\pi), R^*(\pi))}{\Phi^{eq}(\pi)} = \frac{\mathcal{U}_R^{eq}(R^*(\pi), R^*(\pi))}{\Phi^*(\pi) - \Phi^{eq}(\pi)} = 1$$

This offsetting spill-overs condition requires that any two states that deliver the same policy $R^*(\pi)$ also generate exactly the same dynamic spill-overs. With a "hump-shaped" policy that we will show below is the natural response to the pandemic at both the equilibrium and the planner's solution, this offsetting spill-overs condition can hold only if dynamic spill-overs at a given value $R^*(\cdot) = R$ are the same at the onset of the pandemic and during the recovery phase. But that can't happen with the evolution of dynamic spill-overs that we describe below. In addition, the above characterizations suggest that shadow prices scale with $\pi(s)\pi(i)\beta/(1-\beta)$, and hence any dynamic spill-over is bound to swamp the static one when β is close to 1.

Shadow Prices and Continuation Values: We complete this section with a more complete characterization of the shadow prices and continuation values. 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 satisfy

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

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 $\{R_t\}_{t=0}^\infty$. 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 .¹⁵

Denoting by $v_s^*(\pi), v_i^*(\pi)$ and $v_r^*(\pi)$, the life-time utility of a susceptible, infected and recovered agent at the planner's solution, $v^*(\pi)$ admits the following representation

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

which gives us the following expressions:

$$\begin{aligned} -\frac{\partial v^*(\pi)}{\partial \pi(i)} &= v_r^*(\pi) - v_i^*(\pi) - \pi(s)\frac{\partial v_s^*(\pi)}{\partial \pi(i)} - \pi(i)\frac{\partial v_i^*(\pi)}{\partial \pi(i)} - (1-\pi(s)-\pi(i))\frac{\partial v_r^*(\pi)}{\partial \pi(i)} \\ -\frac{\partial v^*(\pi)}{\partial \pi(s)} &= v_r^*(\pi) - v_s^*(\pi) - \pi(s)\frac{\partial v_s^*(\pi)}{\partial \pi(s)} - \pi(i)\frac{\partial v_i^*(\pi)}{\partial \pi(s)} - (1-\pi(s)-\pi(i))\frac{\partial v_r^*(\pi)}{\partial \pi(s)} \end{aligned}$$

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

¹⁵It is easy to prove this step by substituting this guess into the above recursion for V_t .

of the infection rate has on susceptible, infected and recovered agents. These terms, in particular $-\frac{\partial v_s^*(\pi)}{\partial \pi(i)}$, capture 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.

We subtract the marginal value of immunization from the marginal value of recovery to 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.¹⁶

A similar decomposition applies at equilibrium, but the equilibrium doesn't internalize the indirect benefits of recovery and immunity. An agent's equilibrium value function can be represented as a probability-weighted expectation of V_t^s , V_t^i , and V_t^r , using the private beliefs to weight the three different states:

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

With $\hat{v}_s^{eq}(\pi) = \hat{v}_s(\pi, \pi)$, $\hat{v}_i^{eq}(\pi) = \hat{v}_i(\pi, \pi)$, and $\hat{v}_r^{eq}(\pi) = \hat{v}_r(\pi, \pi)$, we obtain

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

for π^k close to π . The RHS represents the expected value to the agent of implementing the same sequence $\{R_t\}_{t=0}^\infty$ as prescribed by the equilibrium at π , which is feasible, though not necessarily optimal, for the agent starting from private belief π^k close to π . Since the RHS is linear in π^k , and equals $\hat{v}^{eq}(\pi)$ when $\pi^k = \pi$, we obtain the private marginal cost of infection:

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

At equilibrium, the private marginal cost of an additional infection corresponds just to the value difference between being susceptible and being infected, but taking as given the values associated

¹⁶This 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.

with each state, which are determined by the dynamics of the aggregate population state. Hence, the equilibrium does not internalize the probability-weighted indirect effects of an additional infection on the continuation values of each type.

Basic value function accounting yields the following additional properties for V_t^s , V_t^i and V_t^r :

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

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

$$V_t^s = V_t^i + (\gamma + \delta) \sum_{s=0}^{\infty} \beta^{s+1} \left(V_{t+s+1}^i - \frac{\gamma}{\gamma + \delta} V_{t+s+1}^r \right) \prod_{k=0}^{s-1} (1 - R_{t+k} \pi_{t+k}(i)) \in [V_t^i, V_t^r]$$

with

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

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. As we will discuss below, the same is not necessarily true for the social marginal cost of infection risks $\frac{\partial v^*(\pi_{+1})}{\partial \pi(s)} - \frac{\partial v^*(\pi_{+1})}{\partial \pi(i)}$.

5 Dynamic Equilibrium and Optimal Policy

S-I-R Dynamics: We now link these recursive equilibrium conditions to the dynamics of π that are generated by the SIR model. Since $1 - \pi_t(s) - \pi_t(i)$ is monotonically increasing and bounded, the population state $\{\pi_t\}$ 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

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

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 , this leads to a hump-shaped profile for $\pi_t(i)$, which is at first increasing, and then decreasing once $R_t \pi_t(s) \leq \gamma + \delta(1 - \pi_t(i))$.

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

$$\begin{aligned} R_t^* &= R^*(\pi_t^*) \text{ and } \pi_{t+1}^* = (1 - \delta\pi_t^*(i))^{-1} \cdot T(R^*(\pi_t^*) \pi_t^*(i)) \cdot \pi_t^* \\ R_t^{eq} &= R^{eq}(\pi_t^{eq}) \text{ and } \pi_{t+1}^{eq} = (1 - \delta\pi_t^{eq}(i))^{-1} \cdot T(R^{eq}(\pi_t^{eq}) \pi_t^{eq}(i)) \cdot \pi_t^{eq}. \end{aligned}$$

Combining the above dynamics with the two first-order conditions yields the following result:

Proposition 3 *Suppose that there is an (arbitrarily) small, positive fraction $\pi_0(i) > 0$ of infected agents in the initial population, and that $\bar{R} > \gamma + \delta$. Then 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 \bar{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 \bar{R} , and the economy returns to the pre-pandemic equilibrium in a state of herd immunity:*

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

We consider the economic and pandemic dynamics with a small initial fraction of infected agents. The assumption that $\bar{R} > \gamma + \delta$ implies that the basic reproductive rate $\mathcal{R}_0 = \bar{R} / (\gamma + \delta)$ at the pre-pandemic equilibrium exceeds 1, and hence the initial infection, however small, can take hold within the population. If in addition $\gamma + \delta > \bar{R}$, then there is the possibility to immediately contain the disease by lowering \mathcal{R}_0 below 1.

Proposition 3 highlights two properties of the economic response to the epidemic which are true at both the equilibrium and at the social planner's solution.

First, both the planner and the agents at equilibrium optimally *flatten the infection curve* by moving away from the utility maximizing action towards the infection risk minimizing one at the onset of the pandemic. This slows the rate at which the pandemic progresses and therefore slows down the rate at which agents are infected and subsequently die. Infections peak later and at a lower level than without a behavioral response.

Importantly, we obtain the rationale for flattening the infection curve 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.¹⁷ Instead 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. This is true both at the planner solution and at the equilibrium.

Second, the proposition shows that there are nevertheless stark limits on the equilibrium and planner’s solution in the long run. Eventually, the epidemic must subside, and both equilibrium and planner’s solution convert 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*. In turn, 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.

Observing a full recovery of the equilibrium to the pre-pandemic steady-state is not overly surprising, since private incentives for confinement disappear when the risk of infection disappears. The result is perhaps more surprising for the planner, who could in principle commit to policies that lower instantaneous utilities and mortality risk in the long run by permanently lowering R . Facing a long-run tradeoff between $\mathcal{V}^*(R)$ and $\pi_\infty^*(s) = (\gamma + \delta)/R$, and noticing that $\mathcal{V}(R)$ is maximized at \bar{R} , there would appear to be long-run welfare gains from a permanent marginal utility distortion that generates a first-order gain in survival probabilities. However, the planner also factors in the delay between the marginal benefit of increasing R today and the marginal cost of higher future mortality. With discounting, this delay explains why the planner wants to slow down the propagation of the pandemic enough to heavily discount the welfare cost of future mortality in the planner’s objective, but not permanently raise the agents’ survival probability.

Both results follow from $\Phi^*(\pi), \Phi^{eq}(\pi) \sim \pi(s)\pi(i)\beta/(1-\beta)$, i.e. the shadow value of infection risk, and hence the marginal utility costs of equilibrium and optimal policy responses, are proportional to the fraction of currently infected agents $\pi(i)$. Coupling this observation with the short-run and long-run properties of the S-I-R dynamics then leads to the above proposition.

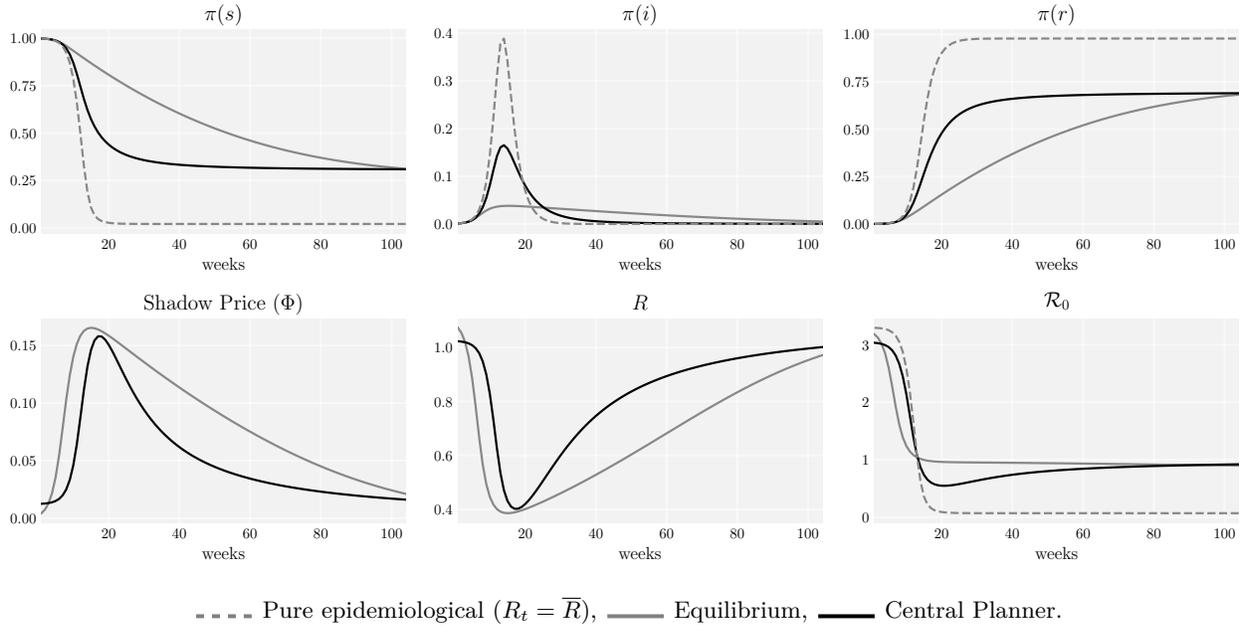
Figure 3 illustrates the economic benefits of flattening the curve in a simulation.¹⁸ The three panels in the top row show the fractions of susceptible agents $\pi_t(s)$, infected agents $\pi_t(i)$ and recovered agents $\pi_t(r) = 1 - \pi_t(s) - \pi_t(i)$ over the course of the pandemic in a purely epidemiological benchmark with $R_t = \bar{R}$, at the equilibrium and at the planner’s solution.¹⁹ The three panels in the

¹⁷We will discuss these channels as quantitative extensions to our baseline model in section 6.1.

¹⁸The parameters are the same as the ones chosen for our benchmark calibration in section ??, except that we have raised the baseline mortality rate $\delta/(\gamma + \delta)$ from 0.5% to 1.5% to better illustrate our main results. Herd immunity is reached when $\pi_t(s) \leq 0.303$.

¹⁹Cumulative mortality is equal to $(\delta/\gamma)\pi_t(r)$ and thus proportional to the fraction of agents who recovered.

Figure 3: The Hammer and the Dance (Proposition 3)



bottom row show the shadow price of infection risks, the equilibrium and optimal policies and the basic reproductive rates \mathcal{R}_0 .

Equilibrium and optimal policy substantially dampen the overall rate of infection early on, the equilibrium more so than the optimal policy. They do not let the infection run its natural course, but seek to reduce the initial peak of infection at a lower level, and thereby substantially reduce the long-run rate of mortality, to near the minimum level necessary to establish herd immunity.

Interestingly, the planner's solution is less restrictive early in the course of the pandemic than the equilibrium, and subsequently recovers faster and with lower long-run mortality than the equilibrium. This points to immunization externalities being more important than infection externalities early in the pandemic, and infection externalities becoming very important later on, while immunization externalities disappear with convergence to herd immunity: early on, the planner internalizes that a recovery requires establishing herd immunity, or in other words, preventing too many infections early on will just postpone them in time and delay the recovery. 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 and therefore the long-run mortality at equilibrium eventually exceeds mortality at the planner's solution, even though it was way lower early on.

We sharpen this dynamic characterization for the cases in which β or \underline{V}/\bar{V} are close to 1. Interpreting $\beta/(1-\beta)$ as the speed of propagation, the limit when $\beta \rightarrow 1$ focuses on a case where the spread is extremely fast. The ratio \underline{V}/\bar{V} instead captures the relative magnitudes of economic surplus vs. mortality costs, and when $\underline{V}/\bar{V} \rightarrow 1$, mortality risk takes priority over economic distortions at all times: The flow utility of keeping an agent alive for another period is bounded below by \underline{V} , while the the cost of economic distortions is bounded above by $\bar{V} - \underline{V}$. If $\underline{V} = \bar{V}$, we have an infinite repetition of the confinement game in which all agents have as their primary objective to maximize their own survival probability. This is achieved at the extreme confinement equilibrium with $R = \underline{R}$ in all periods.

Proposition 4 *For any $\eta > 0$, there exists $\xi > 0$, such that with $\max\{\beta, \underline{V}/\bar{V}\} > 1 - \xi$, equilibrium has the following structure:*

(i) $R^{eq}(\pi_t) < \underline{R} + \eta$ 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 $R^{eq}(\pi_t)$ 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 $R^{eq}(\pi_t)$ are close to $(\gamma + \delta)/\pi_t(s)$ and $\pi_t(s)$ slowly declines at rate less than $R_t\eta$.

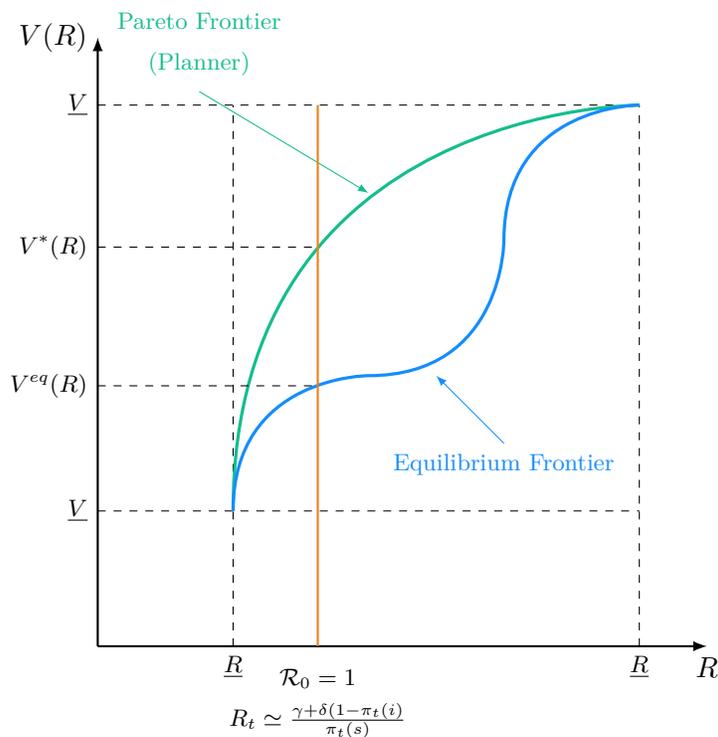
Deconfinement ends when $\pi_t(s)$ reaching the herd immunity threshold $(\gamma + \delta)/\bar{R}$, $\pi_t(i)$ converging to 0, and $R^{eq}(\pi_t)$ converging back to \bar{R} .

Optimal policy $R^*(\cdot)$ follows the same pattern, but the onset of the Hammer phase is delayed until $\frac{\partial v^*(\pi_{+1})}{\partial \pi(s)} - \frac{\partial v^*(\pi_{+1})}{\partial \pi(i)}$ is bounded sufficiently far away from 0.

Proposition 4 describes equilibrium and optimal policy with a fast speed of propagation, or small economic surplus. At equilibrium, the shadow price of infection risks $\Phi^{eq}(\pi)$ then becomes arbitrarily large, relative to the marginal values $\mathcal{V}^{*'}(\cdot)$ and $\mathcal{U}_r^{eq}(\cdot)$ for all $R > \underline{R}$, whenever $\pi(i)$ is sufficiently far from 0 (at least η). When faced with such a situation, agents enact a massive voluntary confinement, close to the extreme confinement equilibrium ("The Hammer") with $R^{eq}(\pi_t)$ arbitrarily close to \underline{R} until the proportion of infected agents is controlled within a narrow band

$\pi(i) \in (0, \eta)$.^{20,21} 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, with a return back to \bar{R} . During the Dance phase, the equilibrium policy $R^{eq}(\pi_t)$ cannot stray far from $(\gamma + \delta(1 - \pi(i))) / \pi_t(s)$, the level that maintains the basic reproduction rate of the infection at 1. The speed of deconfinement is then dictated by the speed at which $\pi_t(s)$ progresses towards herd immunity. Slower deconfinement would trigger a decline in $\pi(i)$, lowering the shadow values of infection risks and removing pressure to keep confinement policies in place. Faster deconfinement instead will increase the infection rate, but this raises shadow values and restores the pressure for stricter confinement policies.²² These dynamics are summarized by Figure 4, with the red line at which the basic reproduction rate, \mathcal{R}_0 , equals 1 gradually shifting to the right over time.

Figure 4: Dynamics



²⁰The labels "The Hammer" and "The Dance" refer to Pueyo (2020) who proposes these phases as a possible strategy for deconfinement.

²¹This Hammer phase is not necessary if the pandemic starts from an initial share of infections below η .

²²The 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.

The optimal policy follows a similar pattern as the equilibrium, but with one important qualification: where as the private cost of infection $v_s^*(\pi) - v_i^*(\pi)$ 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. This reflects the importance of immunization externalities, captured by $\frac{\partial v^*(\pi_{+1})}{\partial \pi(s)}$, early on during the pandemic. However, we show that these immunization externalities diminish as the pandemic progresses, so eventually the social marginal cost of infections becomes sufficiently large so that optimal policy follows the same Hammer and Dance sequencing as the equilibrium.

These results illustrate how the planner's and private shadow values stabilize the optimal and equilibrium path of policy during the deconfinement phase. During the Dance, the path of policy is dictated by the speed at which the epidemic progresses towards herd immunity. The equilibrium and the planner control this speed through the width of the band $(0, \eta)$ at which they find it optimal to stabilize the epidemic, and bringing it back inside this band whenever the infection rate steers too high or too low. And the more patient they are, or the faster the epidemic spreads, i.e. the higher is β , the more the optimal plan tightens the band within which $\pi(i)$ is stabilized.

The instantaneous propagation limit: We now develop further the limit, in which $\beta \rightarrow 1$. As noted above, the planner could, in principle, opt for permanent restriction policies that bound $R^*(\cdot)$ permanently away from \bar{R} to lower the long-run mortality rate. Here, we show that optimization of this long-run tradeoff re-emerges in the limit as $\beta \rightarrow 1$. Define R^* as the long-run optimal policy that maximizes $\Lambda(R^*) \cdot \mathcal{V}^*(R^*)$:

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

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 $\Delta \rightarrow 0$, 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 . In this limit, the infection has the potential to propagate instantaneously in calendar time. We show the following result:

Proposition 5 (*Instantaneous Propagation Limit*): *In the limit as $\Delta \rightarrow 0$:*

(i) **The Dance never ends:** *the planner's optimal choice of policy converges to $R^*(\tau) = R^*$ for all $\tau > 0$. In addition, $\pi(i, \tau)$ converges to 0 and $\pi(s, \tau)$ converges to $(\gamma + \delta) / R^*$ for all $\tau > 0$.*

(ii) **Infinitely strong dynamic spill-overs:** *At the limit of the planner's solution*

$$\lim_{\Delta \rightarrow 0} \frac{\Phi^{eq}}{\Phi^*} = 0 \text{ for all } \tau > 0,$$

and therefore the planner's solution cannot be decentralized as a Markov-Perfect equilibrium, if static spill-overs are bounded.

(iii) **Strong hold-out and infection externalities at equilibrium:** In the limit as $\Delta \rightarrow 0$, the equilibrium policy $\lim_{\Delta \rightarrow 0} R^{eq}(\tau) \pi(s, \tau) = \gamma + \delta$ and $\lim_{\Delta \rightarrow 0} \pi(i, \tau) = 0$ for all $\tau > 0$. In addition, $\pi(s, \tau)$ gradually declines over time, converging to a finite limit $\pi(s, \infty) \in (0, (\gamma + \delta) / \bar{R})$.

Proposition 5 shows that the long-run tradeoff 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 close to a level of infection and recovery associated with the long-run optimum. 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 $\Delta \rightarrow 0$).

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 $\Lambda(\tau)$ converge quasi-instantaneously to the long-run values consistent with a given policy $R^*(\tau)$. But then, at each point during the Dance phase, the planner faces the same quasi-static tradeoff between economic distortion and survival probability, which has a static optimum at R^* , the policy that maximizes $\Lambda(R) \cdot \mathcal{V}^*(R)$. Hence, in the limit as $\Delta \rightarrow 0$, 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 still remains relevant: 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 $R^*(\tau) = R^*$ at $\tau = 0$ the epidemic will overshoot the long-run optimum, so that $\pi(s, \tau) < (\gamma + \delta) / R^*$ for $\tau > 0$.

The fact that the restriction to R^* becomes permanent at the instantaneous propagation limit illustrates another important aspect of optimal policy: the planner responds to optimal policy by slowing down the recovery phase. With discounting or finite speed of propagation, the tradeoff 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 at the planner's solution, the actual speed of progression is inversely related to its potential speed of progression: in response to fast progression of the epidemic, the planner slows or stalls its long-run resolution into the very distant future. This is the only way in which the planner can implement the optimal long-run tradeoff between mortality risk and economic distortions.

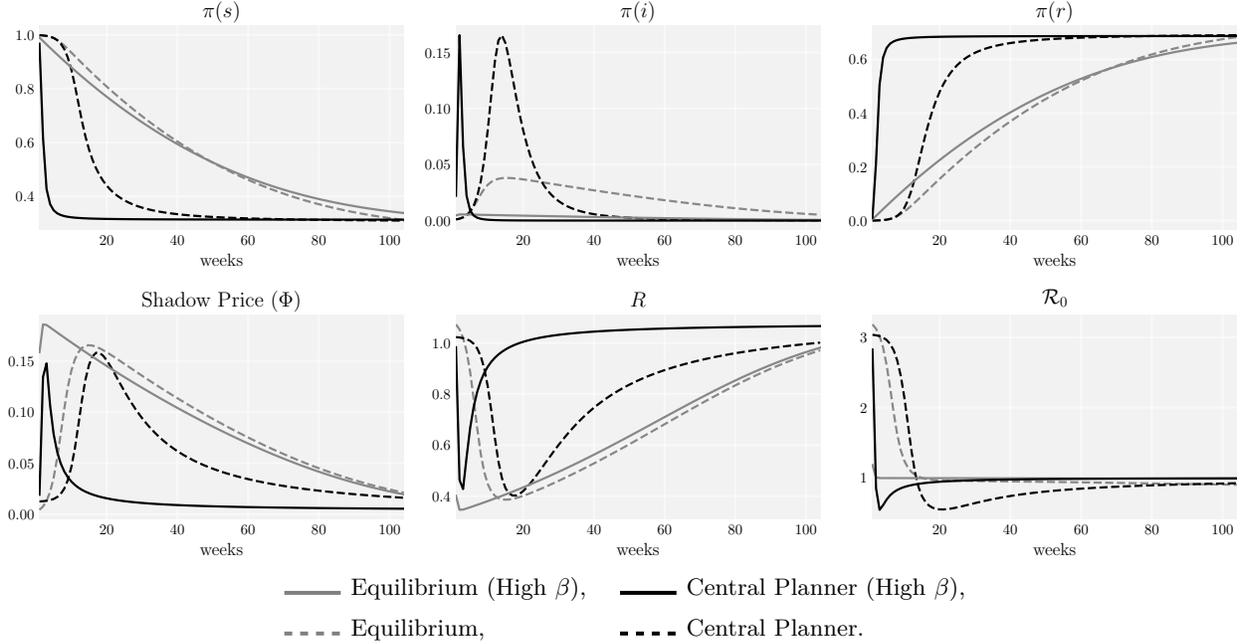
Part (ii) of Proposition 5 shows that dynamic spill-overs are infinitely strong at the planner's

instantaneous propagation limit. At this limit, the private marginal cost of infection equals $\frac{\delta}{\gamma+\delta}\mathcal{V}^*(R^*)$, which weighs an agent's long-run utility $\mathcal{V}^*(R^*)$ by the probability that infection results in death. The social marginal costs of an infection instead multiply $\frac{\delta}{\gamma+\delta}\mathcal{V}^*(R^*)$ by a multiplier that determines the number of additional infections generated by one further infection at the long-run optimum. This multiplier turns out to be infinitely large, i.e. since the basic reproductive rate of the virus equals 1, 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. Any 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. The social costs of an infection can thus be infinitely larger than the private costs.

Part (iii) of Proposition 5 shows that the equilibrium dynamics at the instantaneous propagation limit are very different from the planner's solution. The planner implements an instantaneous convergence to a long-run optimum, which in the limit requires a discrete jump in mortality at the very start of the pandemic, after which everything stabilizes with no further infections and death. But this cannot be part of an equilibrium: if there was a positive jump in mortality at any date, then agents would have a strict incentive to wait out the storm through massive voluntary confinement just before the storm was expected to occur. In other words, at the onset of the pandemic, immunization externalities generate a very strong hold-out motive that substantially slow down the actual speed of propagation, so that even at the instantaneous propagation limit, the pandemic must spread gradually. On the other hand, the long run recovery still requires building herd immunity and thus a higher mortality rate than the planner's long-run optimum.

The comparison between planner's solution and equilibrium at the instantaneous propagation limit highlights the two major sources of inefficiency in response to the pandemic: initially, the immunization externality generates a hold-out motive that inefficiently slows the initial speed of propagation, which amplifies the duration and severity of confinement policies. In the long run, due to the infection externality, agents exit their confinement too fast, relative to the social optimum, which generates higher than necessary long-run mortality. Figure 5 illustrates the insights from propositions 4 and 5. Here, we use the path from Figure 3 as a benchmark and show what happens when we let the pandemic propagate seven times faster: we assume that each infection and recovery takes place over a few days, rather than a few weeks, but we report the results on the same calendar time scale as the benchmark. We see that in the planner's solution the pandemic peaks much faster and is quickly stabilized at level very close to the long-run optimum: the basic reproductive rate \mathcal{R}_0

Figure 5: Proposition 5



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 even more from the beginning, at a rate that just offsets the increase in propagation speed. We clearly see the far slower speed of convergence of the pandemic at the equilibrium, the associated higher economic costs (measured by the shadow price of infection risks), and the higher mortality in the long run.

In this calibration, herd immunity is reached when $\pi(s) \leq 0.303$, and the long-run optimum is located at $\pi(s) = 0.314$. At the benchmark calibration, $\pi(s)$ reaches 0.339 after 40 weeks and 0.305 after 100 weeks at the planner's solution. With faster propagation, $\pi(s)$ instead reaches 0.316 after 20 weeks only, and progresses only to $\pi(s) = 0.313$ after 100 weeks. This shows the convergence to the long-term optimum and extremely slow long-run convergence at the planner's solution. At the equilibrium instead, we obtain $\pi(s) = 0.287$ after 100 weeks at the benchmark and $\pi(s) = 0.301$ after 100 weeks with fast propagation. This illustrates that the equilibrium features inefficiently slow convergence and inefficiently high long-run mortality compared to the planner's solution.

6 Quantitative Results

We now illustrate these results with a simple calibration that disentangles the underlying static and dynamic externalities. We then cover three natural extensions: congestion externalities in the medical sector, the probability of finding a vaccine, and the consequences of delay in the initial response to the pandemic.

Calibration: We assume a utility function of the central planner that takes the form

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

for $R \in [\underline{R}, \bar{R}]$ and $\mathcal{V}^*(\cdot) \in [\underline{V}, \bar{V}]$, or equivalently,

$$\mathcal{V}^*(R) = \underline{V} + (\bar{V} - \underline{V}) \sqrt{1 - \left(\frac{R - \bar{R}}{\bar{R} - \underline{R}}\right)^2}.$$

This elliptic functional form ensures that $\mathcal{V}^*(R)$ satisfies the Inada conditions $\lim_{R \rightarrow \underline{R}} \mathcal{V}^{*'}(R) = \infty$ and $\mathcal{V}^{*'}(\bar{R}) = 0$. We assume moreover that the individual utility function $\mathcal{U}^{eq}(r, R)$ is given by

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

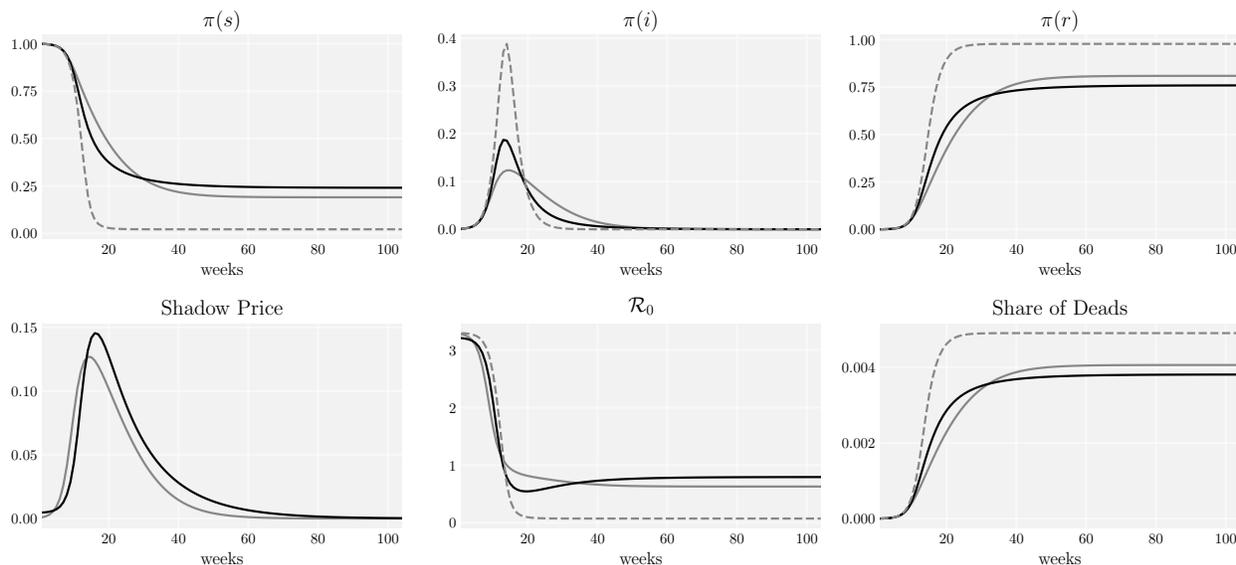
At the symmetric equilibrium, we have $\mathcal{U}^{eq}(R, R) = \mathcal{V}^*(R)$, i.e. the equilibrium coincides with the planner's efficiency frontier (efficient implementation, $X^{eq}(\cdot) = X^*(\cdot)$), and $\mathcal{U}_r^{eq}(R, R) = \alpha \mathcal{V}^{*'}(R)$. If $\alpha = 1$, the individual 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 $u_r(R, R) < \mathcal{V}^{*'}(R)$ so the economic externality dominates the infection externality. Conversely, if $\alpha > 1$, the infection externality dominates.

Our benchmark calibration 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 Jones, Klenow, and Hall (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.²³

Baseline results: Figure 6 shows the dynamics of the epidemic in our baseline calibration. The top three panels give respectively the shares of susceptible, infected, and recovered individuals. The bottom three panels give respectively the shadow price of infections, the reproduction coefficient \mathcal{R}_0 , and the share of deceased individuals.

Figure 6: Benchmark



----- Pure epidemiological ($R_t = \bar{R}$), ——— Equilibrium, ——— Central Planner.

We already discussed the interpretation of these graphs in our discussion of Propositions 3 and 5 above. During the first 15 weeks of the epidemic, the optimal and equilibrium dynamics consist of flattening the curve of infections: in both cases the fraction $\pi(i)$ of infected agents rises less than half as high as in the pure epidemiological model that does not take endogenous economic interactions into account. Correspondingly, the \mathcal{R}_0 coefficient remains lower than in the pure S-I-R model during this period. Delaying infections allows long-run herd immunity to be reached with a smaller fraction of dead individuals. Note that the social and private incentives are roughly aligned during the initial curve-flattening phase: the optimal and equilibrium shadow prices of infection rise in tandem. However, the shadow price is slightly lower for the central planner, implying that the share of

²³In the Appendix we run several robustness checks. We show the sensitivity of our results to the death rate δ , and to the calibration of the utility gap \underline{V}/\bar{V} .

infected agents rises higher, and the \mathcal{R}_0 drops slightly more slowly, than in equilibrium. The second phase of the epidemic, during which the shadow price of infections slowly decreases, corresponds to the controlled deconfinement period. During this phase, the shadow price 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. As before, there is more initial confinement at equilibrium than at the planner's solution, and the subsequent convergence to herd immunity is slower, but the difference is smaller than in the previous graphs which had assumed a higher mortality rate. The relative positions of the shadow prices results from the combination of dynamic externalities that we discuss below.

Figure 7: Utility \mathcal{U}/\bar{V} and Consumption $(\mathcal{U} - \underline{V}) / (\bar{V} - \underline{V})$

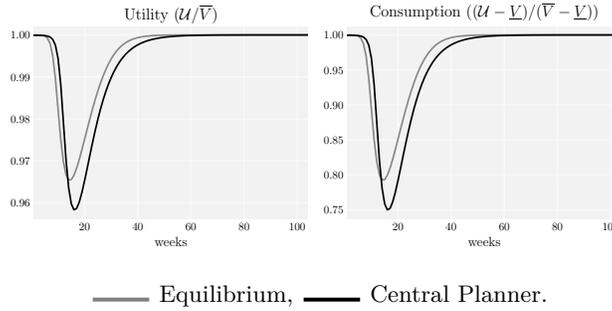


Figure 7 plots the utility and consumption in the social optimum and in equilibrium. Both show a sharp drop during the curve-flattening phase due to the economic lockdown, followed by a slow recovery during the deconfinement. In the long-run, the economy converges back to the optimum with utility \bar{V} . 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, but a faster subsequent recovery.

Dynamic Spill-Overs: Figure 8 unpacks the dynamic externalities that drive the difference between the equilibrium and socially optimal dynamics observed in Figures 6 and 7.

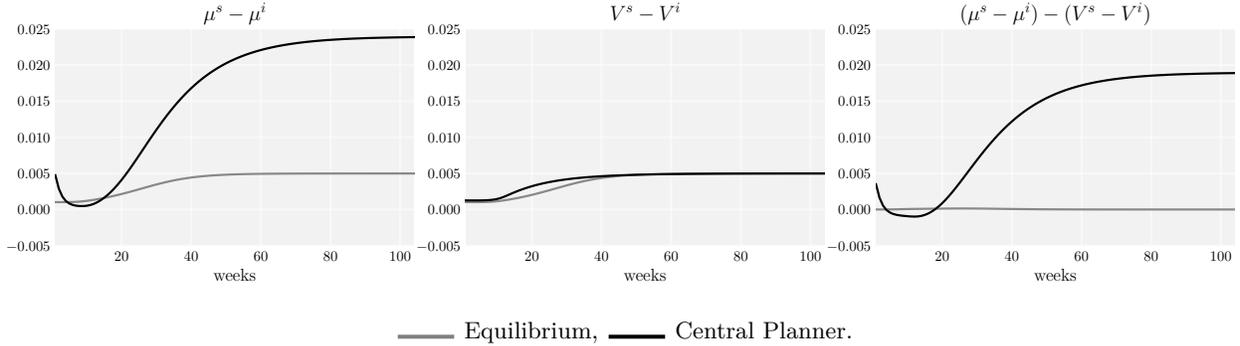
The solid curve in the first panel shows the social marginal cost of an additional infection,

$$\mu^s - \mu^i \equiv \frac{\partial v^*(\pi)}{\partial \pi(s)} - \frac{\partial v^*(\pi)}{\partial \pi(i)}.$$

The solid curve in the second panel gives the private marginal cost of an additional infection,

$$V^s - V^i \equiv v_s^*(\pi) - v_i^*(\pi).$$

Figure 8: Externalities



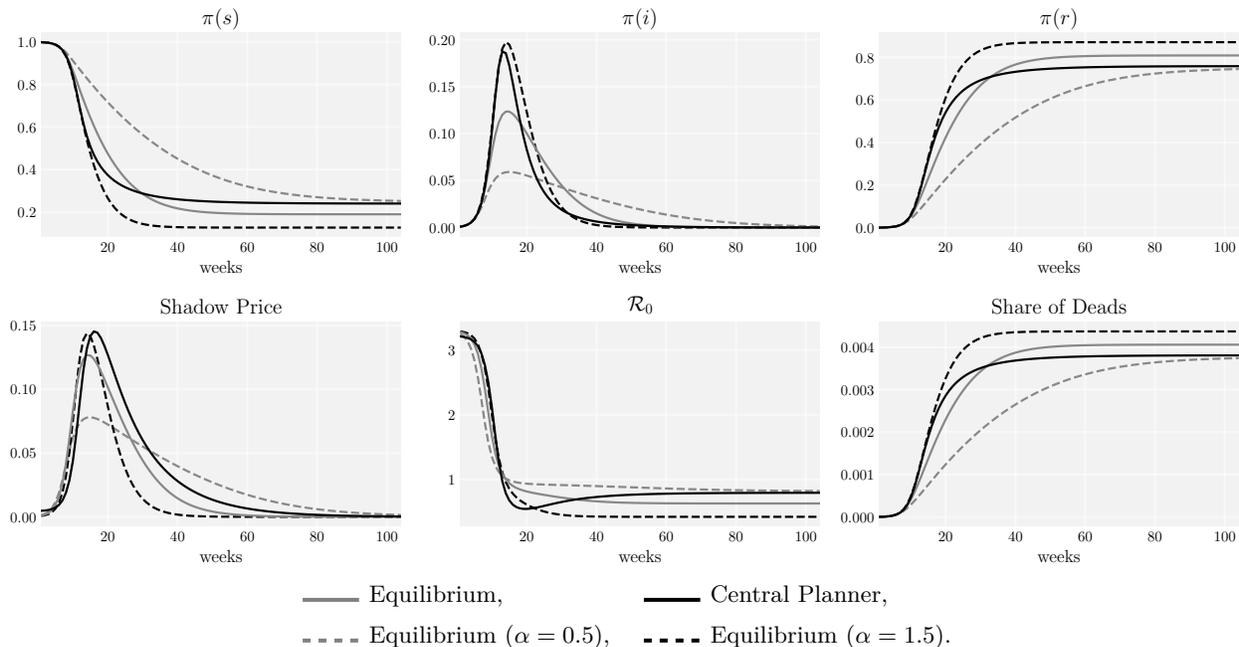
The solid curve in the third panel is the difference between the social cost and the private cost, $(\mu^s - \mu^i) - (V^s - V^i)$. It is therefore a measure of the dynamic externalities at play.

The middle panel shows that the direct cost of a new infection is always positive and increasing. In the long-run, this direct cost is dwarfed by the additional cost created by the dynamic infection externalities: 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: individuals do not internalize that being infected today brings the economy closer to the long-run herd immunity. That is, an infected individual reduces the number of agents who will need to be infected in the future. The optimal path towards herd immunity consists of letting the epidemic spread faster than in equilibrium in the very early stages, which reduces economic costs during the recovery without raising long-run mortality.

Static Externalities: So far we have assumed $\alpha = 1$, meaning that the static externalities offset each other and the social optimum is an equilibrium of the hybrid game. We now explore the sensitivity of our results to these static externalities. Figure 9 plots the results obtained for $\alpha = 0.5$ and $\alpha = 1.5$, and compares them to the planner’s solution which remains unchanged, and to the equilibrium without static externalities ($\alpha = 1$).

When $\alpha = 1.5$, the short-run spill-overs from infection risk dominate: agents do not internalize that their actions raise short-term infection risks for others. We observe offsetting spill-overs at work in the short run: the fact that agents do not fully internalize static infection risks partly offsets the fact that the dynamic immunization externality, resulting in an equilibrium policy that tracks

Figure 9: Varying α



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 excessively fast deconfinement and a higher level of long-run mortality.

When $\alpha = 0.5$, the economic externality dominates, i.e. agents do not internalize the adverse static utility consequences of their actions for other agents. Here, the static economic externality reinforces the immunization externality in the short-run, resulting in a lower peak infection rate and a slower than efficient path of deconfinement. Offsetting spill-overs now appear in the long run: static economic externalities offset the dynamic infection externality, resulting in a slower path to recovery and a lower level of long-run mortality.

6.1 Extensions and applications

Medical Sector Congestion: Here we introduce a congestion externality in the medical sector by letting the death rate be increasing in the fraction of infected agents. For our simulations we follow Jones, Philippon, and Venkateswaran (2020b) and 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 calibrate 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 $\varphi = 0.1682$. The results are presented in Figures 10 and 11.²⁴

Figure 10: Congestion Effects

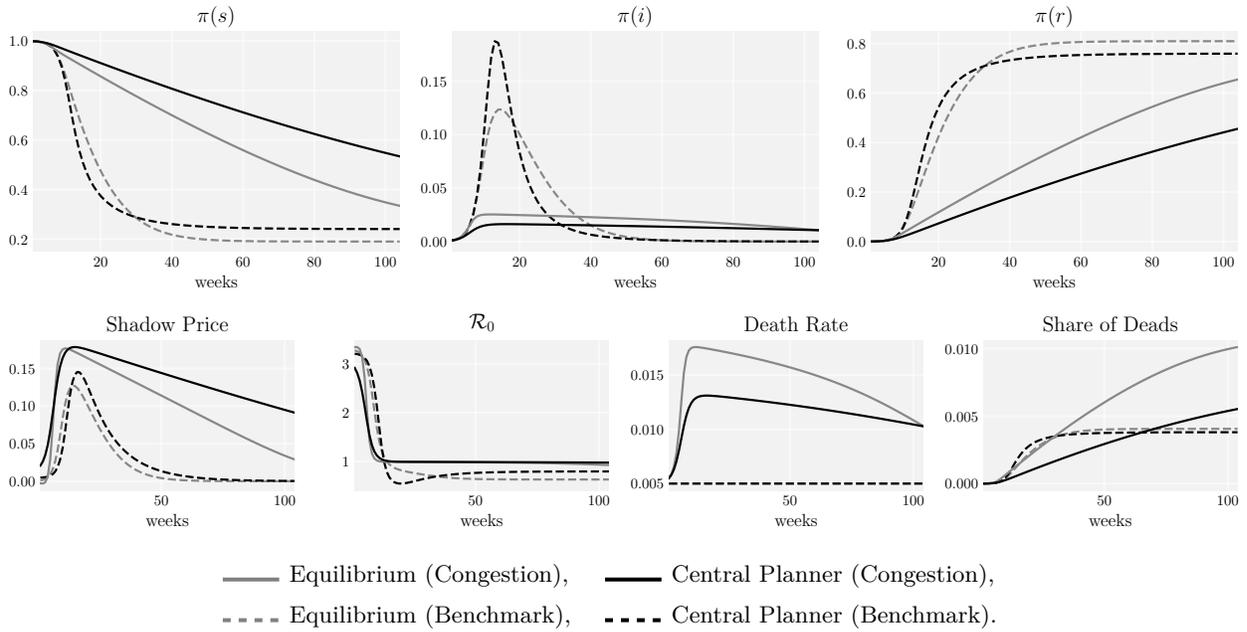
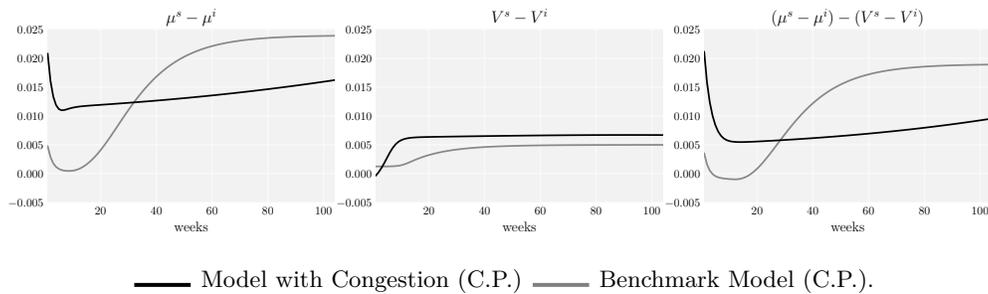


Figure 11: Congestion Effect: externalities



Compared to the baseline model, the shadow 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

²⁴Figures 17 and 18 in the Appendix show the same graphs over a longer horizon.

rate more than linearly due to the congestion externality. The third panel of Figure 11 shows clearly that the dynamic externality is 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 (first panel) 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 externality is not internalized by private agents, the rate of infections, and in turn the death rate, grows much faster in equilibrium than they should under the optimal policy. 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 4 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.²⁵ 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.

Hoping for Vaccines or Cures: Next, we introduce the possibility of a vaccine: in each period, there is a possible probability that all the susceptible individuals move immediately to the recovered state. This vaccine comes too late for agents who are already infected. Like the medical congestion model, vaccines primarily alter the computation of spill-overs, while the baseline FOC only changes marginally, discounting the shadow prices $\Phi^*(\pi)$ and $\Phi^{eq}(\pi)$ by a factor $1 - \xi$, where ξ denotes the weekly probability of discovering a vaccine.

We set ξ to $1/52 = 0.0192$, so that a vaccine is expected to arrive after one year.²⁶ The results are presented in Figures 12 and 13.

The dynamics of equilibrium resemble those of the baseline calibration. Individuals do understand that a vaccine may be found, so they are even more willing to hold out and reduce economic activity

²⁵Formally, the planner’s and agent’s FOCs remain unchanged, but the calculation of dynamic spill-overs changes.

²⁶A 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.

Figure 12: Possibility of a Vaccine

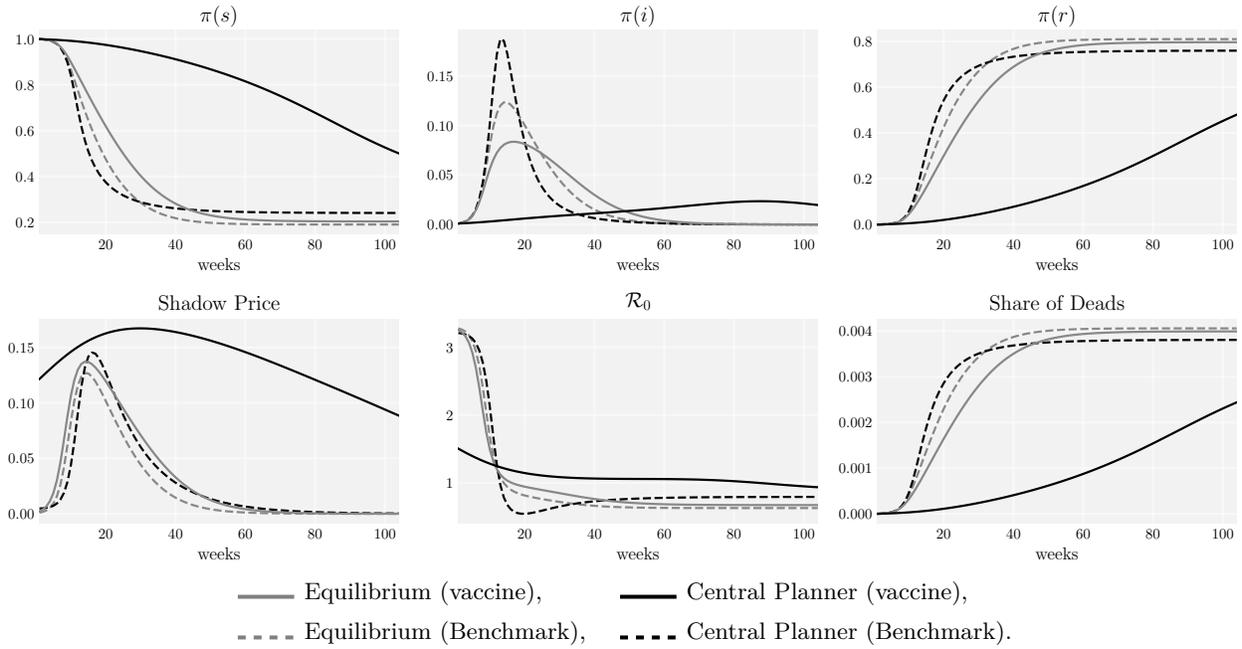
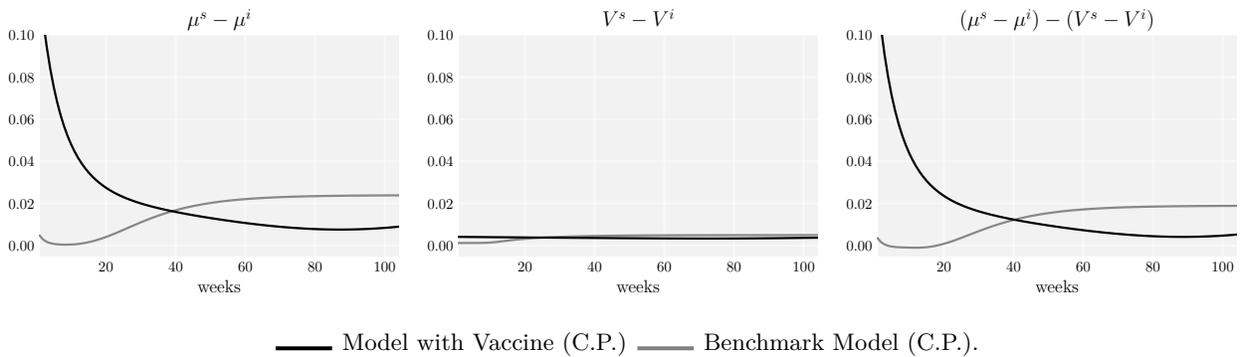


Figure 13: Possibility of a Vaccine: Externalities



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 Figure 6. 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 total social cost of infections represented in Figure 13 is bounded away from zero due to large infection externality that dominates the herd-immunity externality from the start: 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.

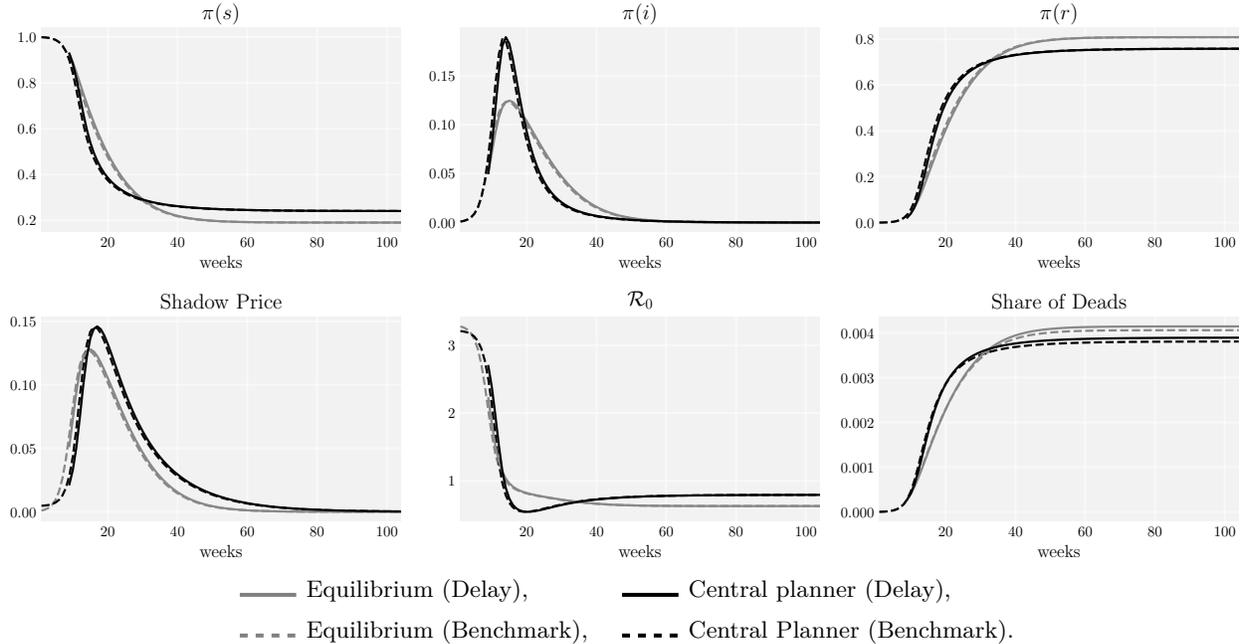
Because the equilibrium doesn’t internalize this value of delaying the pandemic to develop a vaccine, it reaches a much higher level of long-run mortality than the planner’s solution.

As a robustness check, we report the effects of a vaccine that is available at a 2-year horizon ($\xi = 1/104 = 0.0096$) in the appendix. 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 a 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 are 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.

Delaying the Optimal Response: For our last quantitative exercise, we plot the optimal and equilibrium dynamics after the epidemic has been allowed to propagate naturally for two months, as described by the S-I-R dynamics without any adjustment in economic behavior. The results are presented in Figure 14 where we super-impose the baseline dynamics of Figure 6 starting from the same date.

The optimal and equilibrium responses are qualitatively and quantitatively close to those obtained in the baseline calibration. The shadow prices and the peak of infections rise slightly higher, and the long-run share of dead individuals is slightly larger. But overall, the overall cost of not acting as soon as the epidemic starts is extremely low in our benchmark calibration. The reason for this is that in the beginning the delay in policy action is not costly, either at the equilibrium, or for the planner, if it can be compensated by a further tightening of policy in the short-run without affecting long-run outcomes - in other words, in the initial stages of the pandemic, policy is easily

Figure 14: Illness Idle for 2 Months



Note: The economy is first simulated for 2 months using the model without behavior (basic SIR model), then the economy (equilibrium and central planner allocations) is started from the state reached by the SIR model.

substitutable in time, as long as this substitution does not compromise the long-run convergence towards herd immunity.

To see how the long-term consequences matter for the urgency of short-run interventions, we re-run the same experiment in the presence of medical congestion externalities, and find that delay is far more consequential.

6.2 Policy Implications

We conclude this section by summarizing what our theoretical and quantitative results imply for optimal policy design. The main conceptual insight from our analysis is that the dynamics of equilibrium and optimal policy are governed by the dynamics of the shadow values Φ^{eq} and Φ^* . These are shaped by the dynamics of the epidemic, but in turn also determine how the dynamic policy choices shape how the epidemic progresses.

Our theoretical and quantitative results highlight the interplay between dynamic immunization externalities at the start of the pandemic and infection externalities as the pandemic progresses. At our baseline, the planner internalizes the social value of herd immunity and therefore lets the initial outburst run far enough to implement a controlled convergence towards a long-run recovery with a

minimal level of mortality. At the instantaneous propagation limit, this controlled recovery takes the extreme form of instantaneous convergence to the long-run optimum, i.e. small, but long-lasting, if not permanent economic restrictions: the planner responds to faster propagation with a slower recovery.

At equilibrium, instead, the immunization externality takes over in the short run: for obvious reasons, agents are far less eager to participate in the collective effort to build herd immunity than the planner would like to see, i.e. they have strong private incentives for confinement. The resulting hold out slows down the pandemic as well as the progression towards herd immunity - the epidemic peaks later and at a lower level than at the planner's solution, and the economic costs are longer-lasting than at the planner's solution. Moreover, past its peak, infection externalities take over, i.e. agents do not internalize that by risking an infection they risk infecting others later on. Deconfinement occurs too fast in equilibrium, which leads to higher long-run mortality. The equilibrium thus results in an inefficiently slow recovery, and still results in higher long-run mortality than the planner's solution.

Static spill-overs may either reduce the need for short run interventions and increase the need for long-run interventions, or do the opposite – offsetting spill-overs can occur either at the beginning or at the end of the pandemic, but never throughout.

Long-term perspectives determine how much and how fast the planner should intervene at the beginning: the prospect of cures or vaccines reduces the value of herd immunity, and medical sector congestion amplifies infection externalities. This reduces the importance of short-run immunization externalities at the onset, resulting in swifter and far more decisive short run action, in line with the equilibrium. These interventions are optimal because they save lives, rather than merely delaying infections over short horizons. However, they need to be complemented with stronger restrictions along the recovery path, where they amplify the agents' private incentives to exit confinement at a faster rate than optimal.

Finally, getting policy right during the deconfinement period is at least as difficult and as important as during the initial confinement period or at the peak of the pandemic where private and social incentives were strongly aligned on bringing new infections under control. First, how well the policy manages static spill-overs determines whether outcomes are close to the Pareto frontier, which minimizes the static economic costs of pandemic control. Second, the gap between private and public incentives really opens up during deconfinement, and the equilibrium incentive to confine faster than optimal really impacts long-term mortality. The dynamic spill-overs can become arbitrarily large: the closer the policy is to the level at which $\mathcal{R}_0 = 1$, the larger is the multiplier

effect of follow-up infections from each additional infection. Controlling the pandemic by keeping \mathcal{R}_0 below 1 is akin to an upper bound constraint on policy $R_t \leq (\gamma + \delta - \delta\pi_t(i)) / \pi_t(s)$, limiting how far economic restrictions can be relaxed without letting the epidemic grow again. Managing static inefficiencies is paramount to implementing these restrictions with the least amount of economic sacrifice.

We translate these observations into simple policy priorities:

(i) The top policy priority must be to be able to control new infections at all times.

(ii) Early, decisive action ("The Hammer") is warranted if it helps to save lives in the long run. Otherwise, if it merely delays infections in the short run, it just lengthens the recovery and inflicts higher economic costs on the population. Likewise, being slow to act is costly only once it causes deaths that could have been avoided.

(iii) Optimal deconfinement ("The Dance") keeps \mathcal{R}_0 below 1 and relaxes economic restrictions as much as possible within this bound. Policy must control the epidemic, not the other way around.

(iv) Beware of infection risk spill-overs during deconfinement: do not pay for faster recovery with higher mortality, and do not count on private agents to fully understand the risks they pose to others.

(v) The social value of herd immunity really depends on whether there exist better long-term perspectives, and at what horizon.

In many respects, these policy prescriptions are not too far from the ones suggested by epidemiologists. This is not surprising because the dynamics of optimal policy are dictated by the dynamics of the epidemic. They are nevertheless well grounded in basic economic principles: the shadow value of infection risks, the long-run tradeoff between mortality and economic prosperity, dynamic immunization and infection risk spill-overs, and the economics of discounting the future.

7 Face masks, testing and contact-tracing

We complete the paper with two extensions that strike us as particularly relevant for optimal deconfinement policies: the use of face masks, and testing along with contact-tracing.

Face masks: Here we 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 $\lim_{m \rightarrow 0} f_m(m, M) = \infty$, 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.

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

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

The corresponding FOCs yield

$$\begin{aligned} \nabla \mathcal{V}(X) &= \phi f(M, M) \nabla \mathcal{R}(X) \\ C'(M) &= -(f_m(M, M) + f_M(M, M)) \phi \mathcal{R}(X) \end{aligned}$$

Therefore, the use of masks directly equates the social marginal rate of substitution between instantaneous utility and effective infection risks $f\mathcal{R}(X)$ 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 $-(f_m(M, M) + f_M(M, M)) \phi \mathcal{R}(X)$.²⁷ Due to a behavioral response towards utility maximization, the effective infection risk $f\mathcal{R}(X)$ does not decline one-for-one with face masks, but at a rate

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

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

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

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

²⁷We can translate the use of face masks into a reduced form game over the choice of f , given a cost function $C(f)$. 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.

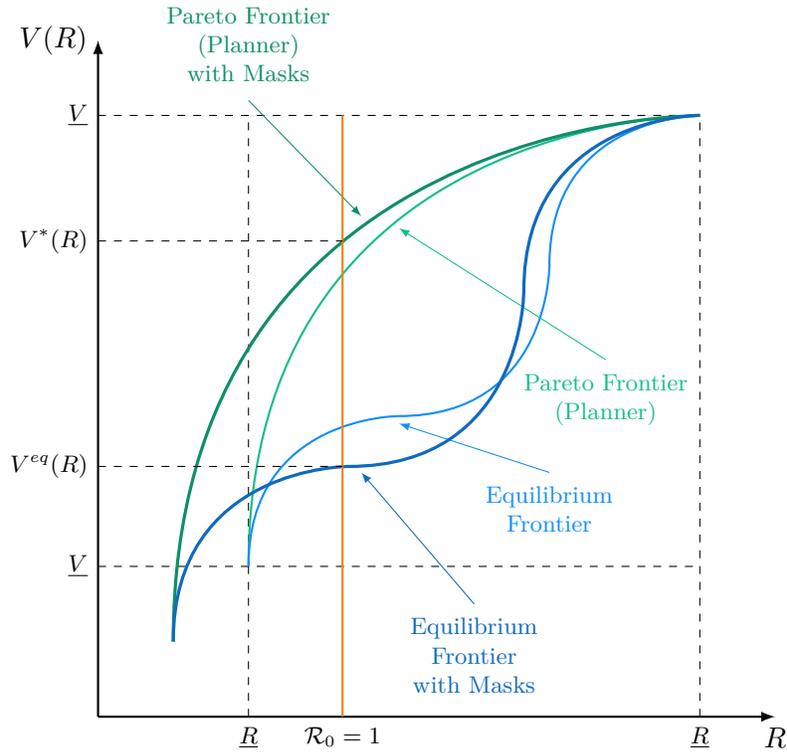
which yields the following first-order conditions for a symmetric equilibrium:

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

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

As before, we can compute the equilibrium frontier between instantaneous utility and effective infection risks $\tilde{R} = fR$. Because of externalities, the new equilibrium frontier is not guaranteed to be strictly higher than the old one. Figure 15 illustrates how the introduction of face masks changes the planner's and equilibrium frontiers.

Figure 15: Introduction of Face Masks



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

mask usage for others:

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

Furthermore, we can represent the shadow price ϕ as

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

i.e. the shadow price of infection risks is a function of the equilibrium price P and the quantity M of face masks. Holding supply of masks constant at $M = \bar{M}$, the shadow price ϕ is proportional to $P/\mathcal{R}(\phi f)$, which yields an elasticity of ϕ w.r.t. P of $\mathcal{E}_{\phi, P}|_{M=\bar{M}} = 1/(1 - \mathcal{E}_R(\phi f)) > 1$. 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.

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, 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 \bar{R} .

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

$$\frac{\tilde{\mathcal{V}}'(\tilde{R}) \tilde{R}}{\tilde{\mathcal{V}}(\tilde{R})} = \frac{\delta}{\tilde{R} - \delta}.$$

Since $\tilde{\mathcal{V}}'(R)/\tilde{\mathcal{V}}(R) < \mathcal{V}'(R)/\mathcal{V}^*(R)$ for all R , the long run optimum with face masks satisfies $\tilde{R} < R^*$. 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.

Testing and contact-tracing: Our second extension concerns 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

$$\hat{\pi}(i) = \pi(i) (1 - \Pr(\text{test}|i)),$$

where $\Pr(\text{test}|i)$ 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 additional state variable $\hat{\pi}(i)$ to keep track of the fraction of agents in quarantine, $\pi(i) - \hat{\pi}(i)$.

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.²⁸ 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 $\pi(i) \rightarrow 0$.

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

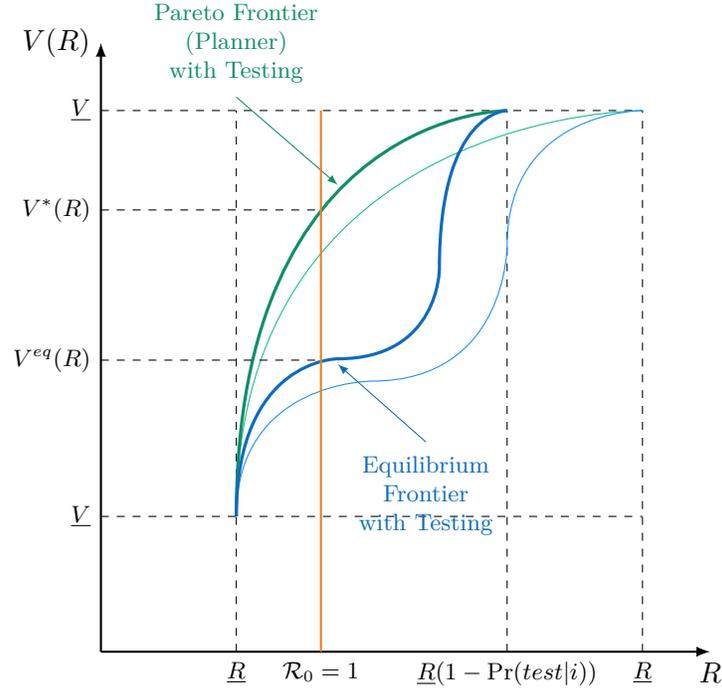
$$\begin{aligned}\pi_{t+1}(i) &= \frac{1 - \gamma - \delta}{1 - \delta\pi_t(i)}\pi_t(i) + \frac{\hat{\pi}_t(i)\pi_t(s)R_t}{1 - \delta\pi_t(i)} \\ &= \frac{1 - \gamma - \delta}{1 - \delta\pi_t(i)}\pi_t(i) + \frac{(1 - \Pr(\text{test}|i))\pi_t(s)R_t}{1 - \delta\pi_t(i)}\pi(i),\end{aligned}$$

augmented by the law of motion for infected agents currently in quarantine, $\pi(i)\Pr(\text{test}|i)$. Therefore, if testing reduces the fraction of infected agents in circulation by a factor $1 - \Pr(\text{test}|i)$, this allows the planner to sustain the same effective infection risk \tilde{R} with R increased by a factor an offsetting factor $1/(1 - \Pr(\text{test}|i)) > 0$. As with face masks, this amounts to a shift in the efficiency frontiers from $\mathcal{V}^*(R)$ to $\mathcal{V}^*(\tilde{R})$ and from $\mathcal{V}^{eq}(R)$ to $\mathcal{V}^{eq}(\tilde{R})$, where \tilde{R} denotes effective infection risks $\tilde{R} = R/(1 - \Pr(\text{test}|i))$. \tilde{R} reaches \bar{R} at $R = \bar{R}(1 - \Pr(\text{test}|i))$, 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 16.

The key is thus to raise $\Pr(\text{test}|i)$, i.e. to test and catch agents once they are infected. By

²⁸Berger, Mongey and Herkenhoff (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.

Figure 16: Frontiers with Testing



Bayes' Rule, we express $\Pr(\text{test}|i)$ as

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

Here $\Pr(\text{test}) \in (0, 1)$ represents the fraction of the population that can be tested within a period, which we take as a parameter proportional to $\hat{\pi}(i)$, and hence small - think of the ratio $\Pr(\text{test})/\hat{\pi}(i)$ the testing capacity relative to ongoing undetected infections. The parameter ρ 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, $\Pr(\text{test}|i) = \hat{\pi}_t(i)$ and $\Pr(\text{test}|i) = \rho \Pr(\text{test})$, and they identify only a small fraction $\rho \Pr(\text{test})$ 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 $\delta \hat{\pi}_t(i)$ 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 $K(\delta \hat{\pi}_t(i) + \mu_t) / (1 - \delta \pi_t(i))$

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 $\hat{\pi}_t(i)$ 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.²⁹ This conditional infection rate corresponds to the unconditional probability of catching an infection $\hat{\pi}_t(i) \pi_t(s) R_t$, divided by the probability of being in contact with an infected person $K\hat{\pi}_t(i)$. Hence they have a probability $\rho \Pr(i|test)$ of returning a positive test result, where

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

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

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

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 $\hat{\pi}_t(i) \rightarrow 0$:

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

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, $K\hat{\pi}_t(i)$, 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.³⁰ This in turn will lead the planner to control the pandemic

²⁹This 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.

³⁰See Pollinger (2020) for a more detailed and complete analysis of the interaction of testing and confinement.

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.³¹

8 Discussion

Our paper provides a parsimonious framework for analyzing economic responses during a pandemic crisis. Here, we discuss the key assumptions that simplified the analysis and suggest potentially interesting directions for future work.

The key step has been to decompose the analysis into a static and a dynamic part, and treat the dynamic part as a repetition of the static game with evolving shadow price on infection risks. This has allowed us to separately identify static spill-overs from infection risk and economic activity, and dynamic propagation spill-overs. Such a static-dynamic decomposition is possible because continuation values are linear in the entries of the transition matrix. It naturally generalizes to other links from stage game decisions to population dynamics, through for example a "recovery game" that measures the impact of individual actions on the probability of recovering from the disease γ , or a "survival game" that summarizes how actions affect mortality risk δ .

The economic interaction embeds a wide range of economic models that satisfy static efficiency at equilibrium (assumption 1). The main limitation of our structure is that focusing on static economic interactions precludes any dynamic adjustment processes such as capital adjustment, dynamic labor market flows, sectoral reallocation with adjustment costs and time to build, or dynamic economic spill-overs, for example through net worth constraints and balance sheet multipliers.

A further simplification has been to focus on a model with symmetric agents. This leaves aside demographic factors, cross-sectional heterogeneity in infection and mortality risks, or differential access to treatment. It would be possible to introduce those using state-of-the-art heterogeneous agents methods, but subject to the challenge of significantly enlarging the state space that needs to be tracked.

Pollinger 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 infection level.

³¹Here, 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.

Infection risks were modeled through a second stage game that also satisfies static efficiency at equilibrium (assumption 2). The two static efficiency assumptions imply that any rationale for policy interventions in the model comes from the tradeoff between economics and epidemic. Most of the qualitative predictions of our model would not change without this assumption, but the case for interventionist policy during a pandemic would only be strengthened if there were direct infection externalities at the equilibrium of the confinement stage game. A strong case for intervention exists even without those direct infection externalities, if the private and social marginal rates of substitution between utility and infection risks aren't well aligned.

The S-I-R model was also kept deliberately simple. We have abstracted from symptomatic transmissions and managing the flow of patients. Extending our approach in that direction or otherwise enriching the epidemiological dynamics would be conceptually easy, but the practical challenge comes from the need to keep track of extra states in the transition matrix, and hence extra state variables in the population dynamics.³² Perhaps the simplest modification to our current structure may be to assume that immunity is not permanent, so that agents transit with positive probability from immunity back to being susceptible.

The assumption that only death is observable allowed us to treat all agents as identical. This assumption mirrors the reality of asymptomatic transmissions with COVID-19, but it also abstracts from a certain number of issues. First and foremost, our model illustrates a clear conflict of interest between agents who have recovered and who would like to return to the pre-pandemic equilibrium as fast as possible, those who are infected and who thus also have nothing to gain from further confinement, and those who are still susceptible to infection who benefit from confinement policies. In our model, the veil of ignorance equalizes these three types, but adding information about infection and recovery will generate a stark conflict of interest between those who have nothing to gain from further confinement and those who do.

Perfect foresight adds further simplification, since knowledge of the current population state π allows the policy maker to manage further propagation perfectly. In reality, governments have to rely on noisy estimates of the true infection rate, based on test results that are biased for example by the aim to target and quarantine those who are infected. When π is unknown and must be estimated from medical test results, a new tradeoff emerges between using tests to catch infections and using tests to estimate π (i) to make sure the infections stays under control. Simple intuition suggests that the costs of straying from the equilibrium infection rate are asymmetric during deconfinement, and

³²Presumably adding a symptomatic state this would be similar to the model with testing and quarantines by allowing exogenous entries into quarantine when an infection becomes symptomatic.

therefore when π is unknown the policy maker should stray on the side of caution, especially when propagation happens sufficiently fast and actual infection rates are very sensitive to policy choices.

Finally, we have adopted a utilitarian welfare criterion, which again seems adequate from the veil of ignorance perspective. It aligns private and social preferences over symmetric utility and mortality outcomes, which also gives the equilibrium the best shot at decentralizing the planner's solution. One could equally well adopt other criteria. For example a Rawlsian criterion that maximizes welfare for the worst off agents might place a higher weight on survival probabilities, and favor policy interventions that are more decisive early on and longer-lasting. Strong myopia or hyperbolic discounting by policy makers or economic agents, perhaps a stand-in for the political economy of upcoming elections, would have the opposite effects of placing too strong a concern on immediate economic prosperity, especially when immediate infection risks are small. In general, a policy maker who does not fully internalize mortality risk will shift the policy towards higher economic prosperity, while a policy maker who overweighs concerns about mortality risks shifts policies towards stronger confinement. However, the tight link between the rate of deconfinement and the rate of new infections limits how far policy makers can go to pursue their own objectives before the epidemic catches up with them.

9 Conclusion

We have analyzed static and dynamic policy design during a pandemic crisis. At the core of our model is a tradeoff between economic activity and dynamic infection risks: by engaging in market activities, agents expose themselves and others to the risk of a potentially lethal infection. Infections in our model go undetected, and they spread fast, in line with the challenges posed by the ongoing COVID-19 pandemic. Static efficiency conditions follow the usual principle of aligning private and social marginal rates of substitution between utility and infection risks. This implies a differential regulation of economic activities, depending on whether infection risk externalities are more important than economic spill-overs or vice versa. Dynamic efficiency conditions are determined by the interplay between immunization and infection externalities in our benchmark model, along with medical congestion effects and the option value of waiting for a vaccine in extensions. They determine whether the equilibrium responds too much too soon early on during the pandemic, resulting in higher than necessary economic costs, or too little, too late, resulting in higher than necessary mortality. And once the peak of infections has passed, agents are tempted to exit confinement too fast and too soon, increasing mortality more than is optimal.

These results offer a novel economic perspective on optimal lockdown policies. They also highlight the challenges that lie ahead during the deconfinement phase when private and social costs of infection start to diverge. The stakes for optimal deconfinement policy are high, and they will ultimately determine whether the initial lockdown truly saved lives, rather than postponing deaths to a later phase of the pandemic.

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>.

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>.

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>.

Chang, Roberto and Andrés Velasco. (2020). “Economic Policy Incentives to Preserve Lives and Livelihoods.” Working Paper 27020, National Bureau of Economic Research. URL <http://www.nber.org/papers/w27020>.

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 during epidemics.” Tech. rep. URL https://sites.google.com/site/mathiastrabandt/home/downloads/EichenbaumRebeloTrabandt_Testing

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>.

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 ‘lock-down’ during an epidemic.” *Covid Economics* 7 (20 April):68–87. URL <https://cepr.org/sites/default/files/news/CovidEconomics7.pdf>.

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>.

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>.

Toxvaerd, Flavio. (2020). “Equilibrium Social Distancing.” Tech. Rep. 2021, University of Cambridge Working Papers in Economics.

10 Appendix: Proofs

Propositions 1, 2 and 3: The proofs follow immediately from the arguments given in the text.

Proposition 4: Fix K and K' such that $K' \geq \hat{v}_s^{eq}(\pi) - \hat{v}_i^{eq}(\pi) \geq \frac{\bar{R}}{\gamma + \delta} K$. This implies

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

Therefore, the equilibrium policy satisfies $\frac{1 - \beta}{\beta} \mathcal{U}_r^{eq}(R, R) \geq \frac{\bar{R}}{\gamma + \delta} \pi(s) K \pi(i)$, or $\frac{1 - \beta}{\beta} \mathcal{U}_r^{eq}(R, R) \geq K \pi(i)$, whenever $\pi(s) > (\gamma + \delta) / \bar{R}$. Now fix $\kappa < \frac{1 - \delta}{1 - \gamma - \delta + \bar{R}} < 1$ and $\eta > 0$. There exists $\xi > 0$, such that whenever $\max\{\beta, \underline{V}/\bar{V}\} > 1 - \xi$ and $R > \underline{R} + \eta$, $\frac{1 - \beta}{\beta} \mathcal{U}_r^{eq}(R, R) < K \kappa \eta$, and therefore $\mathcal{U}_r^{eq}(R, R) < \Phi^{eq}(\pi)$, 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

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

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

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

Therefore $\mathcal{U}_r^{eq}(R, R) > \Phi^{eq}(\pi)$ 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 constants K and K' such that $K' \geq \frac{\partial v^*(\pi)}{\partial \pi(s)} - \frac{\partial v^*(\pi)}{\partial \pi(i)} \geq \frac{\bar{R}}{\gamma + \delta} K$. For given $\pi_0(i)$, we can construct an upper bound K' 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

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

given that $\frac{\partial v^*(\pi_{t+s+1})}{\partial \pi(i)} < 0$, $\beta < 1$ and $v^*(\pi_{t+s+1}) < \bar{V}$. 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.

Proposition 5: Written in calendar time, the planner's problem is

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

subject to the law of motion for π . Proposition 4 implies that there exists $N > 0$, such that $\lim_{\Delta \rightarrow 0} \pi(i, \Delta n) = 0$ and $\lim_{\Delta \rightarrow 0} \left(R(n\Delta) - \frac{\gamma + \delta}{\pi(s, \Delta n)} \right) = 0$ for all $n > N$. It then follows that

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

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 $R(\tau) = R^*$. Abstracting therefore from the Law of Motion for $\pi(s, \Delta n)$, the optimal policy thus sets $R(\tau) = R^*$ and $\pi(s, \tau) = \frac{\gamma + \delta}{R^*}$ for any $\tau > 0$.

We complete the proof of part (i) by showing that there exists a policy path $\{R_{n\Delta}\}$ that displays fast (geometric) convergence to R^* , and that 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 $\Lambda(\Delta(n+1))\pi(i, \Delta(n+1)) = \Lambda(\Delta n)\pi(i, \Delta n)(1 - \delta)$, $\Lambda(\Delta n)\pi(i, \Delta n)$ converges geometrically with a rate $1 - \delta < 1$ towards 0. With $\Lambda(0) = 1$, the total measure of surviving agents converges to $1 - \pi(i, 0)$ as $n \rightarrow \infty$, and the total measure of agents who have recovered converges to $(\gamma/\delta) \cdot \pi(i, 0)$, and the proportion of susceptible agents converges to

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

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

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

If $\pi(i, 0) \leq \bar{\pi}(i)$, there exists a path that converges geometrically to a limit with $\lim_{n \rightarrow \infty} \pi(s, \Delta n) = \frac{\gamma + \delta}{R^*}$. If instead $\pi(i, 0) > \bar{\pi}(i)$, then the planner reaches $\pi(i, n\Delta) \leq \bar{\pi}(i)$ 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.

For part (ii), we observe that $\lim_{\Delta \rightarrow 0} \Phi^* = \mathcal{V}^{*'}(R^*) > 0$ for all $\tau > 0$. We thus show that $\lim_{\Delta \rightarrow 0} \Phi^{eq} = 0$. First, note that at the instantaneous propagation limit with $R(\tau) = R^*$ for $\tau > 0$, $V^s(\tau) - V^i(\tau) = \frac{\delta}{\gamma + \delta} \mathcal{V}^*(R^*)$. This result follows from the fact that $V^s(\tau) = V^r(\tau) = \mathcal{V}^*(R^*)$ when $\pi(s, \tau)$ is stabilized, and $V^i(\tau) = \frac{\gamma}{\gamma + \delta} V^r(\tau)$, when $\beta \rightarrow 1$ (recovery and mortality are resolved instantaneously).

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

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

where $\eta(i, \tau + \Delta k) = \pi(i, \tau + \Delta k) / \Delta$. Taking limits as $\Delta \rightarrow 0$ on both sides with $n = \tau' / \Delta$, and noting that for $\tau > 0$, $\Lambda(\tau + \Delta k) \rightarrow \Lambda(\tau) > 0$, $\pi(s, \tau + \Delta n) \rightarrow \pi(s, \tau) > 0$, and $R(\tau + \Delta k) \rightarrow R^* > 0$, we obtain

$$0 = \lim_{\Delta \rightarrow 0} \sum_{k=0}^{\tau'/\Delta} \Delta \eta(i, \tau + \Delta k) = \lim_{\Delta \rightarrow 0} \int_0^{\tau' - \tau} \eta(i, \tau + s) ds$$

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

11 Additional Figures

Figure 17: Congestion Effects (Long Horizon)

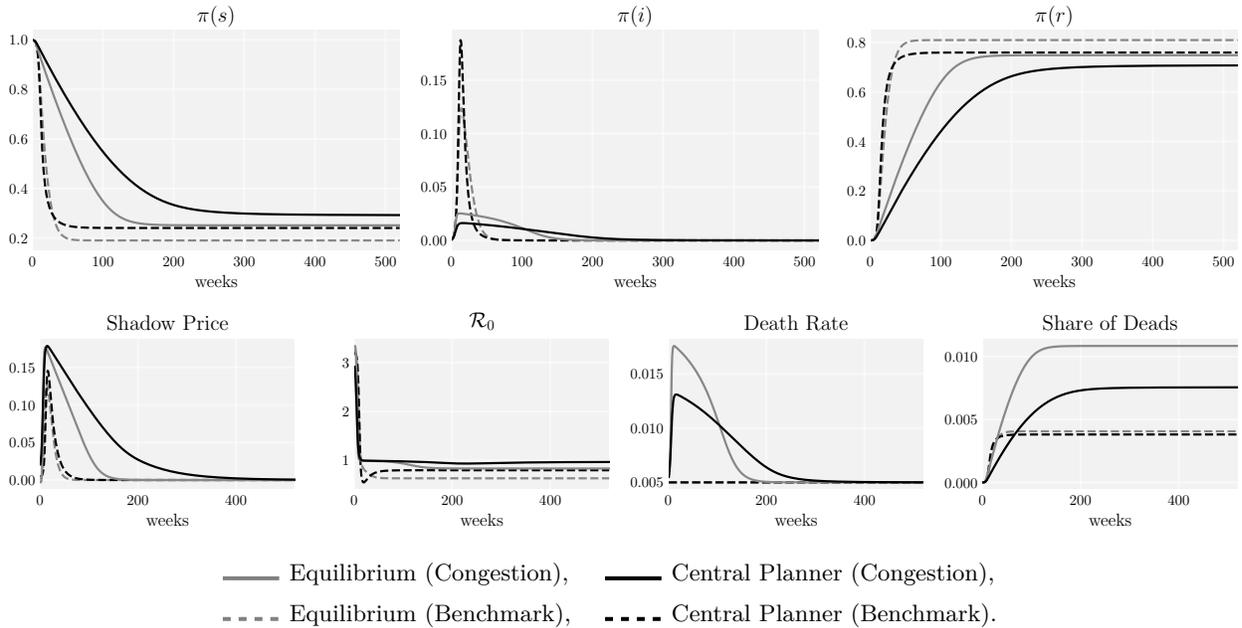


Figure 18: Congestion Effect: Externalities (Long Horizon)

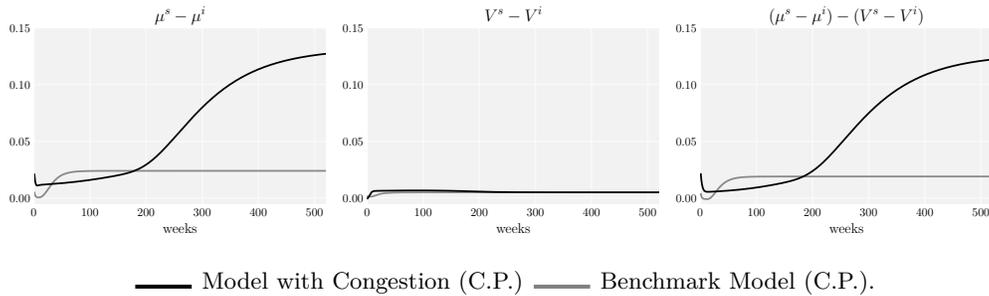


Figure 19: Possibility of a Vaccine (2 year average waiting time)

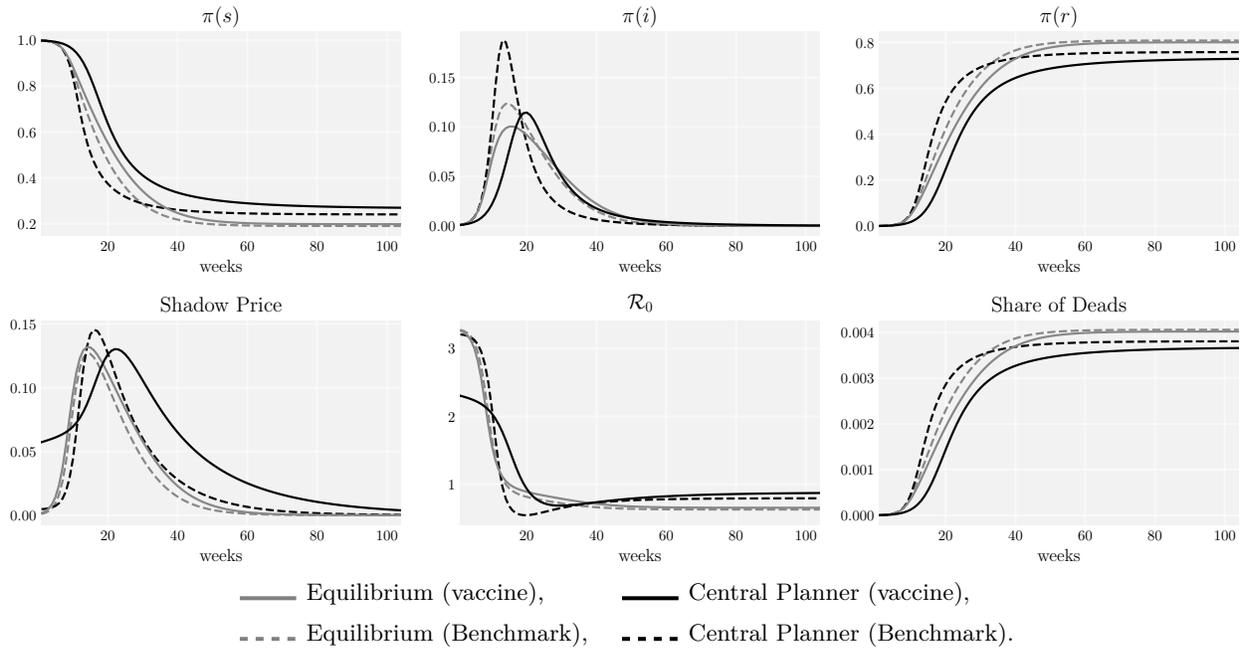


Figure 20: Possibility of a Vaccine: Externalities (2 year average waiting time)

