Submitted to *Quantitative Economics*

Integrated epi-econ assessment: quantitative theory

TIMO BOPPART
Institute for International Economic Studies, Stockholm University and Department of Economics, University of Zurich

KARL HARMENBERG
Department of Economics, University of Oslo

JOHN HASSLER
Institute for International Economic Studies, Stockholm University

PER KRUSELL
Institute for International Economic Studies, Stockholm University

JONNA OLSSON
Department of Economics, NHH Norwegian School of Economics

Aimed at pandemic preparedness, we construct a framework for integrated epi-econ assessment that we believe would be useful for policymakers, especially at the early stages of a pandemic outbreak. We offer theory, calibration to micro-, macro-, and epi-data, and numerical methods for quantitative policy evaluation. The model has an explicit microeconomic, market-based structure. It highlights trade-offs, within period and over time, associated with activities that involve both valuable social interaction and harmful disease transmission. We compare market solutions with socially optimal allocations. Our calibration to covid-19 implies that households shift their leisure and work activities away from social interactions. This is especially true for older individuals, who are more vulnerable to disease. The optimal allocation may or may not involve lockdown and changes the time allocations significantly across age groups. In this trade-off, people's social leisure time becomes an important factor, aside from deaths and GDP. We finally compare optimal responses to different viruses (SARS, seasonal flu) and argue that, going forward, economic analysis ought to be an integral element behind epidemiological policy.

KEYWORDS. integrated assessment, epidemiology, covid-19, time-use data.

JEL CLASSIFICATION. C6, E6, I1.

Timo Boppart: timo.boppart@iies.su.se
Karl Harmenberg: karl.harmenberg@econ.uio.no
John Hassler: john.hassler@iies.su.se
Per Krusell: per.krusell@iies.su.se
Jonna Olsson: jonna.olsson@nhh.no
Harmenberg gratefully acknowledges funding from the Research Council of Norway (project 316301) and Riksbankens Jubileumsfond (program M23-0019).

1. Introduction

May 2023, the World Health Organization announced the end of the emergency phase of covid-19, with close to seven million recorded deaths and with the actual death toll believed to be even higher. The pandemic as well as the policy responses across the world affected virtually all facets of life for a large share of the world's population. Policies previously unimagined such as sustained "lockdowns" were quickly implemented. Were these policies good? More importantly, what is an appropriate policy response to the next large epidemic? And how should governments improve their pandemic preparedness? Echoing Lucas (1988), "the consequences for human welfare involved in questions like these are simply staggering."

A key part of pandemic preparedness is to have robust, off-the-shelf models of the interaction between the economy and the epidemiological spread of a disease that can be used for comparing different policy options: *integrated epi-econ assessment* models. At the beginning of the pandemic, we clearly did not have such models. Equally clearly, a large literature with this aim quickly emerged. The purpose of the present paper is to propose a framework that can quickly be deployed in the future when facing new pandemic threats. The framework was designed with covid-19 in mind and shares many features with papers in the literature that emerged. However, it also offers sufficient generality that comparisons can be made between different viruses; we make such comparisons in the paper. In particular, it makes clear how policy should vary, qualitatively and quantitatively, with the specific features of the virus and its consequences. Thus, the goal is not to make a post-mortem analysis of covid-19 and possible policy errors made; such an evaluation would require a wealth of microeconomic and macroeconomic data that may be available now but certainly was not available at the onset of the pandemic. Rather, the usefulness of our setting is from the ex-ante perspective, as a policymaker's guide facing the next pandemic, without access to the data required for a full evaluation.

The task of designing good epidemic policy is difficult for at least two reasons which motivate the need for integrated epi-econ assessment models. First, the understanding of the evolution of the pandemic—its roots, its dynamic evolution and how different interventions would affect these dynamics—may be quite incomplete. From our perspective, this incomplete knowledge only partly involves epidemiological and related natural-science aspects: since a virus spreads through human interaction, to a large extent we also need knowledge of how people behave and react to the pandemic. Here, economic analysis is at the heart of the matter: the spread of a virus, and the precautions taken to avoid it, are fundamentally decisions made by people in typical *economic* contexts: they involve trade-offs in decisions about how much and what to consume and about how much and where to work; for firms a myriad of related decisions appear as well. We thus view it as an important goal for economics to adapt our frameworks to incorporate epidemics.

The second difficulty is (i) to identify trade-offs between, on the one hand, the mitigation of the public-health effects of the pandemic and, on the other, *economic values*, such as output and consumption; and (ii), when a trade-off is unavoidable, use policy to strike an appropriate balance. Here again, quantitative economic analysis appears central. Moreover, in our view it is not only helpful but also necessary to develop coherent quantitative theoretical frameworks for such evaluations. In particular, by making clear—based on a solid understanding of the mechanisms involved (the first difficulty)—how the balance is struck, policymakers can also hope to be more successful in communicating the reasons behind their interventions and recommendations and thereby be more effective in achieving the goals they set out.

To address these challenges, we introduce epidemiology into a setting with microeconomic foundations, thus both allowing positive and normative analysis to be carried out explicitly. We offer an analysis of the covid-19 pandemic, but an even more important goal of our paper is the development of a core setting on which further analysis—e.g., more complex, or different, viruses, a richer sociological setting, more demographic detail—can be added seamlessly. Because of this more general purpose, we have imposed a number of requirements, which we now briefly list and motivate.

First, we build on *explicit microfoundations:* by clearly describing actors and their decision problems in an explicit equilibrium context, we produce a setting that can provide rich answers to counterfactual experiments. The setting also allows comparison of the planner's optimum with market allocations (where the market allocation may incorporate policy variables). We say "the" optimum since our framework presumes full insurance and a dynastic structure with a representative family consisting of members of different age groups. Our framework can be extended to incorporate lack of insurance as well but we regard the representative-agent setting as a useful starting benchmark.

Second, we incorporate elements of *sociology*. Obviously, social interactions are key for the spread of viruses, so specifying the degree of such interactions in different economic activities—and disciplining the model parameters with associated microeconomic data—is one of our main goals. No less important, social interaction is central for any welfare evaluation in this area. An example is the changed nature of leisure that materializes during times of restricted social activities. In short, the costs of restricting behavior—as for example in lock-downs—need to include effects on leisure and its nature. In our core setting, we tie all social interactions to economic activities and group these activities by their degree of social contact.[1]

Third, our paper aims to provide *quantitative* predictions: even though new, interesting mechanisms can surface when we explore and develop this area further, we see it as secondary to answering questions about magnitudes. To be clear, this area contains important qualitative dimensions; for example, is it better to lock down, or to target herd immunity while protecting the healthcare system? These are, qualitatively, radically different policies, and—as we show in this paper—they can both be the optimal choice. However, the central issue is to find out which one applies in any given situation and this is a wholly quantitative question that can only be settled after careful calibration of the model's central parameters.

[1] In our approach to time use and sociological aspects of consumption, we combine and build further on the approaches in Becker (1965, 1974). Our model of social interactions abstracts from, e.g., different notions of altruism or network theory. Clearly, these features would be interesting to add to our setting.

Fourth, this new area requires numerical methods that can properly handle high-dimensional nonlinear dynamic systems. We do not develop radically new techniques here but provide a solution method that is designed for extensions. It has two parts. One is static: it solves for within-period prices and quantities given some boundary conditions and can accommodate rich heterogeneity and a complex market economy. The other, a dynamic part, connects the boundary conditions and solves for the evolution of these using backward and forward iteration together with dynamic-programming envelope conditions; the boundary conditions consist of state variables and shadow values of additional infections.

Equipped with these elements, we use our *epi-econ integrated assessment model* to address covid-19 and other viruses: the seasonal flu and SARS.[2]

The epi-econ model The general abstract framework we propose consists of a microeconomic model for production, consumption, and time allocation combined with a description of the infection risks associated with the different ways to spend time as well as a state-dependent law of motion for the epidemic. In this framework, individuals are heterogeneous but each individual belongs to a large family within which there is full insurance against any idiosyncratic shocks; at the same time, there are many such families. For this framework, which describes a broad class of models, we characterize equilibrium conditions for the competitive market equilibrium both with "myopic" expectations, where people are unaware of the pandemic, and "perfect foresight" (i.e., rational expectations), where people are fully and correctly informed of the pandemic. We also study the social planner's problem and compare its key equations and implied allocation to the market equations and outcomes. In addition, we provide a robust solution algorithm for all cases.

We put a concrete model structure on this abstract framework as follows. Utility is a function not only of consumption of different goods and services, but also the leisure time individuals spend with the goods and services. In our benchmark model, for simplicity, we dichotomize: goods are either consumed in public or in private, where in-private consumption does not involve social interaction and thus no risk of spreading the disease. Similarly, workers can contribute to production either from home or by working "in the office", with only the latter involving any risk of getting infected. We calibrate utility functions over different goods and leisure activities that involve nontrivial substitution and, similarly, there is nontrivial substitution between work at home and working in the office, all constrained to match available microeconomic data.

Each family has nontrivial demographics among its members. Again, we dichotomize in our main model to "young" and "old" individuals; aside from differing in typical ways (life expectancy, productivity), these two types also differ in their covid-19 vulnerability. Given the strong age gradient in vulnerability to the most negative consequences of covid-19 this is the most obvious starting point for evaluation of the current epidemic, but other or additional heterogeneity could be incorporated.

[2] Though playing out on an entirely different time scale, in important ways the goal in this paper parallels that in the area of climate change, where Nordhaus pioneered a merging of natural-science models of the climate and the carbon cycle with a standard neoclassical economic growth framework. See Nordhaus and Boyer (2000) and, for an explicit market-economy version, Golosov et al. (2014).

Key results and the nature of optimal policy For the baseline, dichotomous model, we find a significant "fear factor": households reallocate their time substantially as a response to the epidemic—if they are fully informed, also in absence of policy interventions. In fact, *rational expectations reduces the death toll in the covid-19 case by 80 percent*, relative to a myopic scenario without behavioral adjustments.

Still, a comparison of the rational-expectations competitive equilibrium and the social planner's optimal scenario turns out to be very interesting. Specifically, we find that *the quality-of-life component is a first-order driver of policy* and that GDP is a poor proxy for welfare of living individuals since it neglects the value of quality leisure. Thus, as the fear factor already deals with the trade-off between the economy and life quite well, the across-group externalities come into focus. In particular, *a social planner distributes the burden of behavioral adjustment more efficiently in the population* by making the young and less vulnerable adjust their behavior beyond what is in their self interest. In the competitive equilibrium, the old are voluntarily staying home to protect their lives, but this comes at a large utility cost since their quality of life plummets. As a result, *the planner reduces output substantially but flow utility during the pandemic's peak increases as the old and vulnerable can enjoy more social leisure*.

The qualitative features of the optimal policy are not a foregone conclusion—our framework is rich enough to rationalize a range of optimal policies, depending on the calibration of the model. In our baseline scenario, no cure is expected to arrive. Here, the social planner optimum is best described as a *"protect the healthcare system"* strategy. The number of infected is kept low enough so that the health system is never overburdened, so that "overshooting" in terms of the number of infected before herd immunity is reached is minimized.

By contrast, if a cure for the disease (or a cheap, perfect vaccine) is expected to arrive soon enough, the social planner chooses a strategy best described as *"suppress"*: keeping the number of infections low by lowering social activity (and thus output) such that the epidemic never takes off (i.e., the reproduction rate is kept, weakly, below one). Further, if the cost of overburdening the healthcare system is low, a qualitatively different optimal allocation transpires: the best response now is to *speed up the spread of the virus*, even more so than under the laissez-faire equilibrium. This way, the welfare costs from isolating and protecting the old fall, as they are incurred over a shorter time span. The number of infected at the peak of the epidemic is high, though the number of deaths still remains at a rather low level as those affected are the young/less vulnerable.

Cross-epidemic restrictions on the value of a statistical life A key input into the calibration of the epi-econ model is the *value of a statistical life*. In our baseline model, we follow Hall et al. (2020) and calibrate the model so that the willingness to pay for an additional day of life is six times daily consumption expenditure. This is a reasonable value, but the literature features a large range of values (see, e.g., Viscusi and Aldy (2003)). We therefore also consider a higher value of a statistical life (a willingness to pay of 11.4 times consumption expenditure, following Glover et al. (2022)). We show that optimal policy can and for reasonable scenarios does shift from a "protect the health care system" strategy to a "suppress" strategy when using the higher value of a statistical life.

We consider the behavioral response to the seasonal flu a *cross-epidemic restriction* on our calibrated model. A value of a statistical life in the upper end of the range implies a counterfactually large responses to the regular seasonal flu, which leads us to conclude that our target for the value of a statistical life is in line with the observed willingness to pay for avoiding the regular flu.

Related literature Upon the onset of the covid-19 epidemic, many economists started analyzing epidemiology models.[3] As a result, there is a sizable such literature and our review of this research here will be limited to the papers that are the most closely related to our work.

Epidemiological models had been used in economics prior to covid-19 in analyses of other viruses.[4] However, Eichenbaum, Rebelo, and Trabandt (2021) (henceforth ERT) is as far as we can tell the first application of epidemiology that is also macroeconomic in the sense that it describes a whole economy of forward-looking agents and market determination of prices. ERT does offer a wholly microeconomic structure (including rational consumers and firms with objectives explicitly described) and the interventions they consider involve fully described policy instruments and comparisons between laissez-faire and fully optimal policy. It is also quantitative in that the model's economic parameters are selected to match (standard) characteristics of macroeconomic data—and the epidemiological parameters are chosen to match known estimates pertaining to the specific features of covid-19. The features of ERT just described are prerequisites for us, and we thus follow ERT in many ways. Our main conceptual addition, in terms of our modeling and quantitative approach, is to move toward explicitly describing time use and to incorporate a sociological element into it. In so doing, we are taking steps toward a development of a socio-economic framework describing how people interact and how they derive utility from it; we use their observed time-use choices, moreover, to construct utility functions that describe these valuations.

Similar papers to ours, that follow ERT in conducting a full, micro-founded epi-econ analysis, include Bethune and Korinek (2020), Chang and Velasco (2020), Farboodi et al. (2021), Garibaldi et al. (2020), Glover et al. (2023), Jones et al. (2021), Kapička and Rupert (2022), Krueger et al. (2022), and van Vlokhoven (2020).[5] Compared to this literature, our aim is entirely quantitative and comprehensive: it is not to prove theorems or

[3]Atkeson (2020) describes the core model developed by epidemiologists (Kermack and McKendrick (1927)) and we use it here.

[4]See, e.g., Geoffard and Philipson (1996), Kremer (1996), Adda (2007), Chan, Hamilton, and Papageorge (2016), and Greenwood, Kircher, Santos, and Tertilt (2019).

[5]There are many papers that are not fully micro-founded or that do not provide a complete planner-vs.-markets comparison, but that are very interesting and valuable in other ways and that have features common to those entertained here. Alvarez et al. (2021) and Giannitsarou et al. (2021) set up and analyze tractable planning problems. Our focus on heterogeneity, both among people and sectors, is shared, either fully or partially with Acemoglu et al. (2021), Acemoglu et al. (2024), Aum et al. (2021), Bodenstein et al. (2022), Brotherhood et al. (2021) and Giagheddu and Papetti (2023). Kaplan et al. (2020) have important insights as regards inequality and the epidemic, which our representative-family framework abstracts from. Eichenbaum et al. (2022b) and Piguillem and Shi (2022) explore the potential for testing policies in epi-econ frameworks. The short paper Boppart et al. (2022) builds directly and explicitly on the present setup in order to analyze the value of vaccination; it does not conduct any model analysis that overlaps with the present work. Bognanni et al. (2020) and Eichenbaum et al. (2024) compare model outcomes with data.

make qualitative points but to develop a setting for practical, quantitative use. In that, we believe we add significantly to the epi-econ literature. An important part of our paper is therefore our approach for solving the model numerically; going forward, computational feasibility can prove to be a hurdle. We do not yet offer a Dynare-like package but our Julia code, available to the reader, is already fast, user-friendly, and easily adaptable. As a proof of concept, we illustrate the power of our methods in Section 7 where we look at richer epidemics and demographics.

Finally, our paper is not just about covid-19 but offers a setting that can be used to study different viruses, including those experienced in the past (we look at the seasonal flu and SARS). We do believe that forward-looking health agencies and governments need to stand ready when the next epidemic surfaces.

Roadmap In the next section, we describe the integrated epi-econ assessment framework compactly, to focus on the general interplay between the economic and the epidemiological sides of the model. Thereafter, in Section 3, we outline the concrete parametrization of the model implemented to evaluate covid-19. In Section 4 we describe how the model is calibrated to time-use data and other data moments. Sections 5 and 6 compare the results from the market allocation to optimal policy, and show how optimal policy qualitatively as well as quantitatively depends on assumptions about key parameters. In Section 7 we enrich the covid-19 model along a number of dimensions, while in Section 8 we test the model by calibrating it to two other diseases: the common seasonal flu and SARS. Finally, Section 9 concludes with a discussion of potential extensions that could easily be incorporated within our framework.

Additional analysis and further details are provided in the Supplemental Appendix (Boppart et al., 2024).

2. General model framework

In this section, we describe the model framework compactly and abstractly. We then characterize equilibrium/optimality conditions and compare the competitive equilibrium with the social planner's solution. In the next section, we introduce the more detailed model structure and calibrate the model.

2.1 *Competitive equilibrium*

We distinguish between the representative family's choices and states (e.g., x) and aggregate choices and states, denoted with bars (e.g., $\overline{x}$). In equilibrium, the representative family's choices and states coincide with the aggregate choices and states (e.g., $x = \overline{x}$).

We use the following notation: for $f(x, y) : \mathbb{R}^{m_1} \times \mathbb{R}^{m_2} \to \mathbb{R}^n$, $\mathrm{D}_x f(x_0, y_0)$ denotes the $n \times m_1$ Jacobian of f with respect to the vector input x, evaluated at (x_0, y_0),

$$\mathrm{D}_x f(x_0, y_0) = \begin{pmatrix} \frac{\partial f_1}{\partial x_1}(x_0, y_0) & \cdots & \frac{\partial f_1}{\partial x_{m_1}}(x_0, y_0) \\ \vdots & \ddots & \vdots \\ \frac{\partial f_n}{\partial x_1}(x_0, y_0) & \cdots & \frac{\partial f_n}{\partial x_{m_1}}(x_0, y_0) \end{pmatrix}.$$

2.1.1 *The household problem* Given the short-run nature of epidemics, we consider an environment without economic state variables (e.g., the capital stock is fixed). The household problem of the representative family can then be decomposed into a dynamic problem of managing the epidemic state and a static problem of optimizing period welfare given a path for the family's epidemic state.

Dynamic problem The dynamic problem that the representative family faces is

$$V_t(\mathcal{S}) = \max_{T,\mathcal{S}'} v_t(T,\mathcal{S}) + \beta V_{t+1}(\mathcal{S}') \qquad \text{s.t.} \qquad \mathcal{S}' = H(\mathcal{S},T,\overline{\mathcal{S}}_t).$$

The representative family has an epidemic state vector $\mathcal{S}$ of dimension $n_{\mathcal{S}}$. The choices of the representative family can be indirectly described as choosing the vector of new infections/transmissions T of dimension n_T. The choice yields an indirect period welfare $v_t(T,\mathcal{S})$ and a future epidemic state $\mathcal{S}' = H(\mathcal{S},T,\overline{\mathcal{S}})$ which depends not only on the family's epidemic state and the transmissions of the family, but also on the aggregate epidemic state at the time, $\overline{\mathcal{S}}_t$.

Static problem The indirect period-welfare function v_t is given by

$$\begin{aligned} v_t(T,\mathcal{S}) = \max_{x,c} \mathfrak{u}(x,c,\mathcal{S}) \qquad \text{s.t.} \qquad p_t c &= w_t x + \Pi_t, && [\mu_t^F] \\ G(x,\mathcal{S},\overline{x}_t,\overline{\mathcal{S}}_t) &= T, && [\mu_t^G] \\ 0 &= C_{\text{eq}}(x,\mathcal{S}), && [\mu_t^{\text{eq}}] \\ 0 &\le C_{\text{ineq}}(x). && [\mu_t^{\text{ineq}}] \end{aligned}$$

The representative family's indirect objective is to maximize period utility $\mathfrak{u}$ given a certain vector of transmissions T. If T is such that it is impossible to satisfy all constraints, then $v_t(T,\mathcal{S}) = -\infty$. The fraktur u, $\mathfrak{u}$, is used to denote the family's period utility, which, when we unpack the formalism in Section 3, depends on the period utility levels of all the individuals in the family. It also depends on how many family members are alive, which is why the epidemic state $\mathcal{S}$ enters the family's period utility.

In our benchmark, the representative family is comprised of two types: young and old. The family chooses a vector of time allocations x of dimension n_x and a vector of consumption levels c of dimension n_c. The vector of time allocations consists of both the time allocations of the young and the time allocations of the old. The family chooses between different activities to spend time on for the young and the old (e.g., work or leisure), with each activity affecting the family's utility, earned income, and infections. In this abstract formulation, leisure time is simply an activity which earns a zero wage. The scalar Π_t is profits accrued from the representative firm. These profits can be interpreted as returns on fixed sector-specific capital stocks. The row vector p_t is the n_c-dimensional vector of consumption-good prices, the row vector w_t is the n_x-dimensional vector of wages from the different activities, and the function G maps the behavior and state of the family (x and $\mathcal{S}$) together with the aggregate behavior and aggregate state ($\overline{x}_t$ and $\overline{\mathcal{S}}_t$) to a vector of new infections/transmissions T for the family.

Finally, the representative family faces equality constraints on time allocations as described by C_{eq} and inequality constraints as described by C_{ineq}. The Lagrange multipliers on the constraints are indicated by the bracketed variables.

The firm side in the economy is represented by an aggregate production function F yielding a vector of output $F(\overline{x})$ of dimension n_c. Wages are determined by marginal product and profits are given by output net of wage payments.

2.1.2 *Equilibrium definition* A competitive equilibrium in this economy is a sequence of prices w_t, p_t, profits Π_t, time allocations x_t, consumption quantities c_t, epidemiological states $\mathcal{S}_t$ and transmissions T_t for all t such that:

1. The sequences of x_t, c_t, T_t, and $\mathcal{S}_t$ solve the representative family's problem.
2. w_t and Π_t satisfy:
$$w_t = p_t \, \mathrm{D}_x F(\overline{x}_t) \quad \text{and} \quad \Pi_t = p_t F(\overline{x}_t) - w_t \overline{x}_t.$$
3. The time allocation and epidemiological state of the representative family coincide with the aggregate time allocations and the aggregate epidemiological state: $x_t = \overline{x}_t$, $\mathcal{S}_t = \overline{\mathcal{S}}_t$.
4. The goods markets clear: $c_t = F(\overline{x}_t)$.
5. The aggregate law of motion for the epidemic is given by:
$$\overline{\mathcal{S}}_{t+1} = H(\mathcal{S}_t, T_t, \overline{\mathcal{S}}_t).$$

2.1.3 *Equilibrium characterization* The competitive equilibrium comprises the first-order conditions and the envelope condition for the dynamic problem as well as the first-order conditions, envelope condition, and general-equilibrium conditions for the static problem; these conditions are all necessary.

Dynamic equilibrium conditions From the dynamic problem, we obtain the following equilibrium conditions:

$$\mathrm{D}_T v_t(T_t, \mathcal{S}_t) = -\beta \, \mathrm{D}_{\mathcal{S}} V_{t+1}(\mathcal{S}_{t+1}) \, \mathrm{D}_T H(\mathcal{S}_t, T_t, \mathcal{S}_t), \tag{1}$$
$$\mathrm{D}_{\mathcal{S}} V_t(\mathcal{S}_t) = \mathrm{D}_{\mathcal{S}} v_t(T_t, \mathcal{S}_t) + \beta \, \mathrm{D}_{\mathcal{S}} V_{t+1}(\mathcal{S}_{t+1}) \, \mathrm{D}_{\mathcal{S}} H(\mathcal{S}_t, T_t, \mathcal{S}_t), \tag{2}$$

where the first equation is the first-order condition and the second equation is an application of the envelope theorem. The derivative $\mathrm{D}_{\mathcal{S}} H$ is with respect to the first argument of $H(\mathcal{S}, T, \overline{\mathcal{S}})$. We also have the law of motion for the epidemic,

$$\mathcal{S}_{t+1} = H(\mathcal{S}_t, T_t, \mathcal{S}_t). \tag{3}$$

Static equilibrium conditions The static-problem equilibrium can be characterized by six equations. First, the first-order condition combined with equilibrium prices and an application of the envelope theorem with respect to T_t yield

$$\begin{aligned}\mathrm{D}_x \mathfrak{u}(x_t, c_t, \mathcal{S}_t) + \mathrm{D}_c \mathfrak{u}(x_t, c_t, \mathcal{S}_t) \, \mathrm{D}_x F(x_t) = {} & \mathrm{D}_T v_t(T_t, \mathcal{S}_t) \, \mathrm{D}_x G(x_t, \mathcal{S}_t, x_t, \mathcal{S}_t) \\ & + \mu_t^{\text{eq}} \, \mathrm{D}_x C_{\text{eq}}(x_t, \mathcal{S}_t) + \mu_t^{\text{ineq}} \, \mathrm{D}_x C_{\text{ineq}}(x_t)\end{aligned} \tag{4}$$

where $\mathrm{D}_x G$ refers to the derivative with respect to the first argument. Second, an application of the envelope theorem with respect to $\mathcal{S}_t$ yields

$$\mathrm{D}_{\mathcal{S}} v_t(T_t, \mathcal{S}_t) = \mathrm{D}_{\mathcal{S}} \mathfrak{u}(x_t, \mathcal{S}_t) + \mathrm{D}_T v(T_t, \mathcal{S}_t)\, \mathrm{D}_{\mathcal{S}} G(x_t, \mathcal{S}_t, x_t, \mathcal{S}_t) + \mu_t^{\text{eq}}\, \mathrm{D}_{\mathcal{S}} C_{\text{eq}}(x_t, \mathcal{S}_t), \tag{5}$$

where $\mathrm{D}_{\mathcal{S}} G(x_t, \mathcal{S}_t, x_t, \mathcal{S}_t)$ refers to the derivative with respect to the second argument. Finally, the resource constraint, infection constraint, equality constraint and inequality constraint yield

$$c_t = F(x_t), \tag{6}$$

$$G(x_t, \mathcal{S}_t) x_t = T_t, \tag{7}$$

$$0 = C_{\text{eq}}(x_t, \mathcal{S}_t), \tag{8}$$

$$0 = \mu_t^{\text{ineq}} \odot C_{\text{ineq}}(x_t, \mathcal{S}_t). \tag{9}$$

The Hadamard product $\odot$ indicates pointwise multiplication of the vectors. Additionally, the Kuhn-Tucker conditions require $\mu_t^{\text{ineq}} \geq 0$.

2.1.4 *Myopic equilibrium* Later, we also consider a "myopic equilibrium", where the representative family is unaware that its behavior affects its epidemic state. The myopic equilibrium conditions are identical, except that we set the (believed) effect of behavior on infections to zero in equation (4), $\mathrm{D}_x G = 0$.[6]

2.2 *Planner formulation*

The dynamic problem that the planner faces is

$$V_t^{\text{SP}}(\mathcal{S}) = \max_{T, \mathcal{S}'} v^{\text{SP}}(T, \mathcal{S}) + \beta V_{t+1}^{\text{SP}}(\mathcal{S}') \qquad \text{s.t.} \qquad \mathcal{S}' = H(\mathcal{S}, T, \mathcal{S})$$

where the indirect period-welfare function v is given by

$$\begin{aligned} v^{\text{SP}}(T, \mathcal{S}) = \max_{x, c} \mathfrak{u}(x, c, \mathcal{S}) \qquad \text{s.t.} \qquad c &= F(x), && [\mu_t^{F,\text{SP}}] \\ G(x, \mathcal{S}, x, \mathcal{S}) &= T, && [\mu_t^{G,\text{SP}}] \\ 0 &= C_{\text{eq}}(x, \mathcal{S}), && [\mu_t^{\text{eq,SP}}] \\ 0 &\leq C_{\text{ineq}}(x, \mathcal{S}). && [\mu_t^{\text{ineq,SP}}] \end{aligned}$$

2.2.1 *Optimality conditions* The optimal allocation delivers the necessary first-order conditions and the envelope conditions for the dynamic problem as well as the first-order conditions and envelope conditions for the static problem. The optimality conditions are identical to equations (1)–(9) except for the adjustments to equations (2), (4), and (5) that internalize the epidemiological externalities.

[6]Thus, in terms of epidemic outcomes the myopic equilibrium is for all intents and purposes equivalent to a naive model with no behavioral responses to the epidemic (a standard non-economic SIR model). Technically there is a slight difference due to deaths, which slightly shifts the capital-to-labor ratio and consequently the wage rate in the economy, which affect the agent's time allocations, but this effect is so small it is negligible.

Dynamic externality The application of the envelope theorem in the dynamic problem now yields

$$\begin{aligned} \mathrm{D}_{\mathcal{S}} V_t^{\mathrm{SP}}(\mathcal{S}_t) = {} & \mathrm{D}_{\mathcal{S}} v^{\mathrm{SP}}(T_t, \mathcal{S}_t) \\ & + \beta\, \mathrm{D}_{\mathcal{S}} V_{t+1}^{\mathrm{SP}}(\mathcal{S}_{t+1}) \left(\mathrm{D}_{\mathcal{S}} H(\mathcal{S}_t, T_t, \mathcal{S}_t) + \mathrm{D}_{\overline{\mathcal{S}}} H(\mathcal{S}_t, T_t, \mathcal{S}_t) \right) \end{aligned} \tag{2'}$$

where the externality $\mathrm{D}_{\overline{\mathcal{S}}} H(\mathcal{S}_t, T_t, \mathcal{S}_t)$ is the derivative with respect to the third argument, capturing how the aggregate state $\overline{\mathcal{S}}$ affects the law of motion for the family's state, for example through overcrowding of the hospital system.

Static externalities The first-order condition combined with an application of the envelope theorem with respect to T_t yield

$$\begin{aligned} \mathrm{D}_x \mathfrak{u}(x_t, c_t, \mathcal{S}_t) + \mathrm{D}_c \mathfrak{u}(x_t, c_t, \mathcal{S}_t)\, \mathrm{D}_x F(x_t) = {} & \mathrm{D}_T v^{\mathrm{SP}}(T_t, \mathcal{S}_t) (\mathrm{D}_x G(x_t, \mathcal{S}_t, x_t, \mathcal{S}_t) \\ & + \mathrm{D}_{\overline{x}} G(x_t, \mathcal{S}_t, x_t, \mathcal{S}_t)) \\ & + \mu_t^{\text{eq, SP}}\, \mathrm{D}_x C_{\text{eq}}(x_t, \mathcal{S}_t) + \mu_t^{\text{ineq, SP}}\, \mathrm{D}_x C_{\text{ineq}}(x_t, \mathcal{S}_t) \end{aligned} \tag{4'}$$

where the externality $\mathrm{D}_{\overline{x}} G(x_t, \mathcal{S}_t, x_t, \mathcal{S}_t)$ is the derivative with respect to the third argument, capturing that an individual's behavior not only affects her own risk, but also the risk of others, of becoming infected.

Finally, an application of the envelope theorem with respect to $\mathcal{S}_t$ yields

$$\begin{aligned} \mathrm{D}_{\mathcal{S}} v_t^{\mathrm{SP}}(T_t, \mathcal{S}_t) = {} & \mathrm{D}_{\mathcal{S}} \mathfrak{u}(x_t, \mathcal{S}_t) + \mathrm{D}_T v^{\mathrm{SP}}(T_t, \mathcal{S}_t) (\mathrm{D}_{\mathcal{S}} G(x_t, \mathcal{S}_t, x_t, \mathcal{S}_t) \\ & + \mathrm{D}_{\overline{\mathcal{S}}} G(x_t, \mathcal{S}_t, x_t, \mathcal{S}_t)) \\ & + \mu_t^{\text{eq, SP}}\, \mathrm{D}_{\mathcal{S}} C_{\text{eq,SP}}(x_t, \mathcal{S}_t) + \mu_t^{\text{ineq, SP}}\, \mathrm{D}_{\mathcal{S}} C_{\text{ineq,SP}}(x_t, \mathcal{S}_t) \end{aligned} \tag{5'}$$

where the externality $\mathrm{D}_{\overline{\mathcal{S}}} G(x_t, \mathcal{S}_t, x_t, \mathcal{S}_t))$ is the derivative with respect to the fourth argument, capturing that the within-period indirect value of an individual for the planner also depends on whether she is spreading the epidemic to others.

2.3 *Solution algorithm*

The state $\mathcal{S}$ is generally high-dimensional. In our model, described in the next section, the state space is given by seven state variables, of which six are continuous. Thus, solving the model with value function iteration can be challenging and time-consuming. Moreover, the epidemiological dynamics are distinctly non-linear. This makes any type of linearization of the problem unappealing.

We take an approach based on first-order and envelope conditions in solving the model, using the equilibrium conditions described above. The model has two conceptual blocks: (a) a static economy that takes as given a population structure (that of course is generated by the past) and a vector of permissible transmissions, and (b) a dynamic dimension, in which periods are connected by the number of transmissions and the resulting new population structure. The key equilibrium object to solve for is therefore the path of shadow values of infections as we show below.

The following algorithm conjectures a path of shadow values of infections, and then computes an implied path for them. A weighted average of the guess and the implied path provides a new guess, and we find the equilibrium/optimal path of shadow values of infections by repeating the algorithm until convergence. The same algorithm is used to solve both the competitive equilibrium and the planner's problem (the latter is solved by exchanging equations (2), (4), and (5) for (2'), (4'), and (5')). We identify candidates solving all necessary conditions first and, lastly, ascertain sufficiency.

0. **Initial conditions**

   Set a final time period $\bar{t}$ and either conjecture that the epidemic is over at this point or assume that a cure instantaneously arrives a time $\bar{t}$. Set an initial condition for the epidemiological state $\mathcal{S}_0$.

1. **Guess path**

   Guess on a path of shadow values of infections, $\{(\mathrm{D}_T v_t)^{\text{guess}}\}_{t=0,\ldots,\bar{t}}$

2. **Roll the epidemic forward**

   For each $t \le \bar{t}$ and a given $\mathcal{S}_t$, compute within-period objects, including a $\mathrm{D}_{\mathcal{S}} v_t$, and $\mathcal{S}_{t+1}$. At each t, starting with $t = 0$, thus:

   - given $\mathcal{S}_t$ and $\mathrm{D}_T v_t$, solve the static equilibrium conditions, equations (4)–(9), which is a system of equations in $x_t, c_t, T_t, \mu_t^{\text{eq}}, \mu_t^{\text{ineq}}$, and $\mathrm{D}_{\mathcal{S}} v_t$;
   - given $\mathcal{S}_t$ and the newly computed T_t, the epidemic is rolled forward—$\mathcal{S}_{t+1}$ is computed—using the law of motion (3).
   - Continue until $t = \bar{t}$.

3. **Envelope values backwards**

   Equipped with paths for $\mathcal{S}_t, T_t$, and $\mathrm{D}_{\mathcal{S}} v_t$ from the previous step, now solve backwards from the terminal period $\bar{t}$, given a $\mathrm{D}_{\mathcal{S}} V_{\bar{t}+1}$, to obtain $\mathrm{D}_T v_t$ for each t. This works as follows.

   - Use equation (1) to compute the implied $\mathrm{D}_T v_t$ as a function of $\mathrm{D}_{\mathcal{S}} V_{t+1}$, $\mathcal{S}_t$ and T_t.
   - Given $\mathrm{D}_{\mathcal{S}} V_{t+1}, \mathcal{S}_t$, T_t, and $\mathrm{D}_{\mathcal{S}} v_t$, equation (2) can be solved for $\mathrm{D}_{\mathcal{S}} V_t$.
   - Continue until $t = 0$.

   The result of this step is an implied path of shadow values of infection, denoted $\{(\mathrm{D}_T v_t)^{\text{implied}}\}_{t=0,\ldots,\bar{t}}$.

4. **Compare and update guess**

   Compare the implied path, $\{(\mathrm{D}_T v_t)^{\text{implied}}\}_{t=0,\ldots,\bar{t}}$, to the current guess $\{(\mathrm{D}_T v_t)^{\text{guess}}\}_{t=0,\ldots,\bar{t}}$.

   - If the implied path is close enough to the current guess, we have found a solution to the necessary conditions.

- Otherwise, update guess to a convex combination of previous guess and the implied path. Go back to point 2.

Given a solution to the necessary conditions, it is important to carefully evaluate whether there are other solutions. For the household in the market equilibrium and for the planner, this would amount to multiple (local/global) optima; for the market economy, it would amount to multiple equilibria. In the case of our covid-19 economy below, we have conducted search and found that the planning problem can have multiple local optima; the solutions reported below are global. For the corresponding market economy, we did not find multiple competitive equilibria.[7]

3. Model structure

We now present our covid-19 economy. We use x' to denote the transpose of x.

3.1 *Structure of the dynamic problem*

Epidemic state, $\mathcal{S}_t$ The representative family consists of a continuum of young and old (denoted by superscripts y and o), who are either susceptible, infected, or recovered. The family does not know which individuals are susceptible, infected, or recovered, but updates its belief in a Bayesian fashion so that it knows the size of each group. The epidemic state is thus

$$\mathcal{S}_t = \begin{pmatrix} S_t^y \; I_t^y \; R_t^y \; S_t^o \; I_t^o \; R_t^o \end{pmatrix}'.$$

The transmission vector T_t *and the law of motion for the epidemic,* H The vector of transmissions/new infections is the vector of new infections for young and old,

$$T_t = \begin{pmatrix} T_t^y \; T_t^o \end{pmatrix}'.$$

The law of motion for the epidemic is given by

$$H(\mathcal{S}_t, T_t, \overline{\mathcal{S}}_t) = \begin{pmatrix} S_t^y - T_t^y \\ (1-\pi_r - \pi_d^y(\overline{I}_t^y + \overline{I}_t^o))I_t^y + T_t^y \\ R_t^y + \pi_r I_t^y \\ \hdashline S_t^o - T_t^o \\ (1-\pi_r - \pi_d^o(\overline{I}_t^y + \overline{I}_t^o))I_t^o + T_t^o \\ R_t^o + \pi_r I_t^o \end{pmatrix}$$

for $i = y, o$ where the increasing death-rate functions π_d^i $(i = y, o)$ capture that the healthcare system potentially becomes overwhelmed when too many are infected at the same

[7] In general, we did not find multiple local optima. However, when performing comparative statics (e.g., with respect to the value of a statistical life), as the optimal policy qualitatively shifts from, e.g., “protect the healthcare system” to “suppress”, there can be some parameter values where both qualitative policies are local optima.

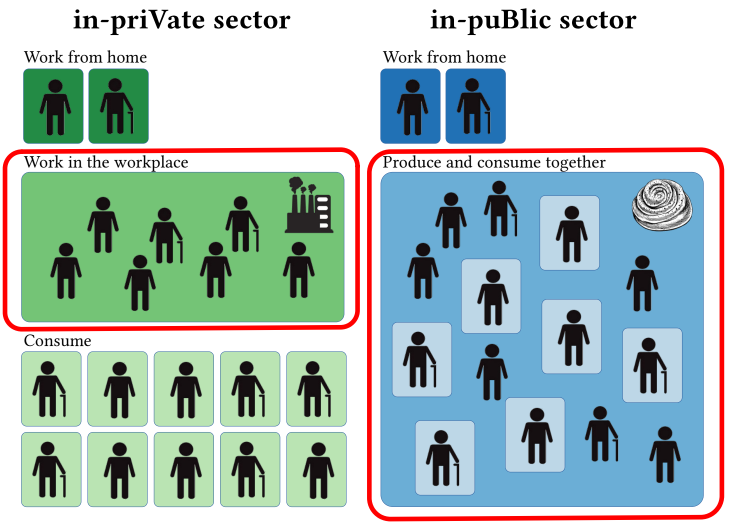

The areas marked in red are where the virus spreads.

FIGURE 1. Illustration of the model.

time. The entries above the dashed line describe the law of motion for the family's young individuals, and the entries below the dashed line describe the law of motion for old individuals.

3.2 *Structure of the static problem*

Overview of production, consumption, and leisure Throughout in our benchmark—as with our demographic young/old structure—we dichotomize so as to introduce minimal heterogeneity without losing key dimensions. In particular, we separate activities with no social interaction from those with full interaction. In Section 7 we introduce more heterogeneity so as to illustrate the power of the approach. Figure 1 schematically illustrates the economic model we have in mind.

There are two sectors in the economy: the in-priVate (V) sector and the in-puBlic (B) sector. When individuals spend leisure in the in-priVate sector (illustrated by the figures in the bottom left in the figure) they are at home (e.g., watching television), and there is no risk of getting infected. The goods and services used for the in-priVate leisure are produced in the workplace or at home (illustrated by the upper half of the left side of the figure, where people are, e.g., in the studio producing a Netflix show or in the Amazon warehouse shipping a new TV set). When people work in the workplace, they interact with their colleagues, and there is a risk of spreading the virus.

The right-hand side of the figure illustrates the in-puBlic sector. In this sector, consumption and work take place jointly, and the virus can spread between those enjoying leisure (e.g., customers in the restaurant) and those working in the sector (e.g., waiters in the restaurant). However, even in this sector there is a possibility (for at least some employees, e.g., the restaurant's accountant) to work from home without physically interacting with others. We now formalize this model.

Time allocations x_t Both the young and the old can spend time on working in the B sector, either in the workplace (n^i_{Bw}) or from home (n_{Bh}), on working in the V sector, either in the workplace (n^i_{Vw}) or from home (n^i_{Vh}), and on leisure, either B leisure (h^i_B) or V leisure (h^i_V). The time-allocation vector of the representative family is

$$x = \begin{pmatrix} x^y \\ \hdashline x^o \end{pmatrix} = \left(n^y_{Bw} \; n^y_{Bh} \; n^y_{Vw} \; n^y_{Vh} \; h^y_B \; h^y_V \;\vdots\; n^o_{Bw} \; n^o_{Bh} \; n^o_{Vw} \; n^o_{Vh} \; h^o_B \; h^o_V \right)'$$

where the entries denote total hours spent on the activities by the young and the old (i.e., *not* per-capita hours spent on the activities). The dashed line separates the variables for the young and the old. Note that this formulation of the time allocations assumes homogeneity within groups: all young/old spend their time equally.

Time constraints C_{eq} *and* C_{ineq} Each individual in the family has one unit of time at their disposal per time period, and cannot spend negative hours on any activity. These constraints are summarized below, with the notation $\text{Pop}^i = S^i + I^i + R^i$,

$$\begin{aligned} C_{\text{eq}}(x, \mathcal{S}) &= \begin{pmatrix} n^y_{Bw} + n^y_{Bh} + n^y_{Vw} + n^y_{Vh} + h^y_B + h^y_V - \text{Pop}^y \\ n^o_{Bw} + n^o_{Bh} + n^o_{Vw} + n^o_{Vh} + h^o_B + h^o_V - \text{Pop}^o \end{pmatrix}, \\ C_{\text{ineq}}(x) &= x. \end{aligned}$$

Production function F Production of the two consumption goods is determined by the function

$$F(x) = \begin{pmatrix} y^B(x) \\ y^V(x) \end{pmatrix}$$

where production in both the B and the V sector is given by a Cobb-Douglas function of CES aggregates,

$$\begin{aligned} y^j(x) &= k_j^\alpha \tilde{n}_{jh}^\nu \tilde{n}_{jw}^{1-\alpha-\nu}, \\ \tilde{n}_{jh} &= \text{CES}(n^y_{jh}, n^o_{jh}; \lambda_n, \varepsilon_n), \\ \tilde{n}_{jw} &= \text{CES}(n^y_{jw}, n^o_{jw}; \lambda_n, \varepsilon_n), \end{aligned}$$

for $j = B, V$ and the constant-elasticity aggregator CES($\bullet$) is defined by

$$\text{CES}(x_1, x_2; \lambda, \varepsilon) = \left(\lambda x_1^{\frac{\varepsilon-1}{\varepsilon}} + (1-\lambda) x_2^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}}.$$

The sector-specific capital stocks k_B and k_V are fixed (and constant over time).

Family utility $\mathfrak{u}$ The family's period utility function is a weighted average of the utility of the young and the old,

$$\mathfrak{u}(x,c,\mathcal{S}) = \max_{c^y,c^o} \sum_{i=y,o} \text{Pop}^i u(x^i/\text{Pop}^i, c^i/\text{Pop}^i) \qquad \text{s.t.} \qquad c^y + c^o = c$$

where consumption is optimally allocated between the two groups.

Let hatted variables denote per-capita consumption and time allocations (e.g., $\hat{x}^i = x^i/\text{Pop}^i$). The flow utility for an individual member of the family is given by

$$\begin{aligned} u(\hat{x}^i, \hat{c}^i) &= \log \text{CES}(\tilde{c}^i_B, \tilde{c}^i_V; \lambda, \varepsilon) + \underline{u}, \\ \tilde{c}_B &= \text{CES}(\hat{c}^i_B, \hat{h}^i_B; \lambda_B, \varepsilon_B), \\ \tilde{c}_V &= \text{CES}(\hat{c}^i_V, \hat{h}^i_V; \lambda_V, \varepsilon_V), \end{aligned}$$

where the constant $\underline{u}$ is added to calibrate the value of a statistical life.

The nested CES structure captures the idea that the consumer needs to spend leisure time to derive utility from a good or service. To derive utility from non-social in-priVate goods (the television set, the streaming service subscription, the groceries) agents need to spend time with it (watch TV, have dinner at home). Likewise, to derive utility from social in-puBlic goods (movie tickets or a restaurant meal) agents need to spend time on in-puBlic leisure (the time spent in the movie theater or restaurant).

Infection rate function G The infection risks in the B and the V workplaces are given by a constant-returns-to-scale (i.e., linear) random meeting technology,

$$\pi^B(\overline{x}, \overline{\mathcal{S}}) = \kappa_B \frac{\sum_{i=y,o} (\overline{n}^i_{Bw} + \overline{h}^i_B)\overline{I}^i/\overline{\text{Pop}}^i}{\sum_{i=y,o} (\overline{n}^i_{Bw} + \overline{h}^i_B)},$$

$$\pi^V(\overline{x}, \overline{\mathcal{S}}) = \kappa_V \frac{\sum_{i=y,o} \overline{n}^i_{Vw}\overline{I}^i/\overline{\text{Pop}}^i}{\sum_{i=y,o} \overline{n}^i_{Vw}}.$$

The vector of infections is given by

$$G(x, \mathcal{S}, \overline{x}, \overline{\mathcal{S}}) = \underbrace{\begin{pmatrix} S^y/\text{Pop}^y & 0 \\ 0 & S^o/\text{Pop}^o \end{pmatrix}}_{\text{Probability of being susceptible}} \underbrace{\left(\begin{array}{cccccc|cccccc} \pi^B & 0 & \pi^V & 0 & \pi^B & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 & 0 & 0 & \pi^B & 0 & \pi^V & 0 & \pi^B & 0 \end{array}\right)}_{\text{Infection risk per time unit, if susceptible}} \underbrace{x}_{\text{Time alloc.}}$$

where the arguments of π^B and π^V are suppressed. The dashed vertical line separates the time allocations of the young and the old while the dashed horizontal line separates the infection risks of the young and the old. The lower-left and upper-right blocks of zeros in the infection-risk matrix capture the fact that the family's infections for the old do not depend on the activities of the family's young, and vice versa.

3.3 *The terminal condition*

To compute the equilibrium, we need a terminal condition for the gradient of the value function after the final period $\bar{t}$, $\mathrm{D}_{\mathcal{S}}V_{\bar{t}+1}$. We assume that the epidemic is over at $t = \bar{t}$ and thus the values of additional susceptible, infected, and recovered young individuals are all equal for $t = \bar{t} + 1$ (and likewise for the old). The value of a single time period is $\tilde{u}^y$ and $\tilde{u}^o$ for a young and an old individual, respectively. The old live for $\mathcal{T}^o$ periods and the value of an additional old person after the final period $\bar{t}$ is simply the discounted sum of the period utility $\tilde{u}^o$ for their remaining life expectancy. The young live for $\mathcal{T}^y$ periods and are assumed to first have the period utility of the young, and for the last $\mathcal{T}^o$ periods of their life the period utility of the old. In sum, the marginal values of additional individuals are given by

$$\mathrm{D}_{\mathcal{S}}V_{\bar{t}+1} = \left(\mathcal{U}^y \; \mathcal{U}^y \; \mathcal{U}^y \,\vdots\, \mathcal{U}^o \; \mathcal{U}^o \; \mathcal{U}^o \right)'$$

with

$$\mathcal{U}^y = \tilde{u}^y \frac{1 - \beta^{\mathcal{T}^y - \mathcal{T}^o - \bar{t}}}{1 - \beta} + \tilde{u}^o \beta^{\mathcal{T}^y - \mathcal{T}^o - \bar{t}} \frac{1 - \beta^{\mathcal{T}^o}}{1 - \beta},$$
$$\mathcal{U}^o = \tilde{u}^o \frac{1 - \beta^{\mathcal{T}^o - \bar{t}}}{1 - \beta}.$$

4. Calibration

In this section we describe the calibration of the model. Table 3 in the end of the section summarizes all parameter values.

4.1 *Calibration of the dynamic problem*

Period length and the discount factor β A period in the model corresponds to a day. We set the discount factor β such that $\beta^{365} = 0.96$.

Initial condition, $\mathcal{S}_0$ We define young as individuals aged 15-60, while old refers to individuals above 60. The initial population size is normalized to one. Thus, calibrating to U.S. demography, we set Pop_0^y, the young population share at time 0, to 0.73 and $\mathrm{Pop}_0^o = 1 - \mathrm{Pop}_0^y$. There is no population growth. At time 0, a share 0.001 of the population, both young and old, is infected,

$$\mathcal{S}_0 = \left(0.999 \cdot \mathrm{Pop}_0^y \; 0.001 \cdot \mathrm{Pop}_0^y \; 0 \,\vdots\, 0.999 \cdot \mathrm{Pop}_0^o \; 0.001 \cdot \mathrm{Pop}_0^o \; 0 \right)'.$$

The law of motion for the epidemic H The law of motion for the epidemic is described by the recovery rate π_r and the mortality rates π_d^y and π_d^o (which are functions of the number of currently infected).

In line with Atkeson (2020) and Eichenbaum et al. (2021) we set the average time from infection to recovery to be 18 days. This time also corresponds to the time from symptom to recovery in Glover et al. (2023). Since our model is daily, π_r, the recovery rate, is set to $1/18$.

We assume that the death rate of the illness is an increasing function of the number of infected. The marginal death rate is given by a logistic function for which the midpoint of the logistic curve occurs at the point where the hospitals are getting overcrowded, which is assumed to occur when the fraction of infected in the population reaches $\hat{I}$.[8] The average death rate in any time period is thus given by:

$$\pi_d^i(I) = \left(\int_0^I \pi_{d,low}^i + \frac{\pi_{d,high}^i - \pi_{d,low}^i}{1 + e^{-k(z-\hat{I})}} \, \mathrm{d}z \right) \Big/ I$$

with I denoting the sum of the young and the old infected. Based on U.S. data, there were 29.4 intensive care units (ICUs) per 100,000 people at the onset of the covid-19 crisis so we assume one ICU per 3,400 people.[9] Further, we assume that three percent of the infected individuals require hospitalization, and, based on estimates for Sweden, that 29% of the hospitalized are in need of intensive care.[10] Taken together, this gives us an $\hat{I} = 1/(0.03 \times 0.29 \times 3,400) \approx 0.034$. In other words, we assume that the death rate will quickly increase when the number of infected reaches 3.4% of the population.

The probability of dying (on a given day) conditional on being infected, when there is no over-crowding in the hospitals, is set to $0.001 \times 1/18$ for the young and $0.025 \times 1/18$ for the old, following Glover et al. (2023). The average infection fatality rate in the population, if the young and old were infected at the same rate, would thus be 0.7%. When the healthcare system is completely overburdened, the probabilities are assumed to treble, to $0.003 \times 1/18$ for the young and $0.075 \times 1/18$ for the old. The steepness of the logistic curve, k, is set to 1,000.

4.2 *Calibration of the static problem*

The calibration of the static problem is more involved. The calibration proceeds in two steps. First, the economic parameters (the production function and preferences) are jointly calibrated to time-use data and long-run growth facts. Second, the epidemiological parameters are calibrated to evidence on the contagiousness of covid-19.

4.2.1 *Calibration targets*

Calibration targets from the ATUS For the calibration targets for the allocation of time, we turn to the American Time Use Survey (ATUS), which provides nationally representative estimates of how and where Americans spend their time. Importantly, it includes data on the full range of nonmarket activities, from relaxing at home to restaurant visits and attending sports events.

We divide the 24 hours a day into three mutually exclusive and complementary exhaustive broad categories: sleep, work, and leisure. Sleep is defined as the time spent

[8]With our calibration, this formulation is effectively a step function, but continuously differentiable.

[9]Based on information from Society of Critical Care Medicine, downloaded June 24, 2020 (link).

[10]Glover et al (2020) assume a hospitalization rate of 2% for the young and 12.5% for the old, which with our population shares would give a weighted average of 4.9%.

TABLE 1. Average minutes per day spent in different activities.

	Young		**Old**	
Activity	Minutes	Fraction of time awake	Minutes	Fraction of time awake
Sleep	527		534	
Time awake	913	1	906	1
Work	413	0.45	236	0.26
Leisure	500	0.55	670	0.74
in-puBlic	*163*	*0.18*	*150*	*0.17*
in-priVate	*337*	*0.37*	*520*	*0.57*

Source: American Time Use Survey (U.S. Bureau of Labor Statistics, 2020). Activities include associated traveling.

either sleeping or experiencing sleeplessness. We define work as the sum of the following activities: market work, core housework (meal preparation and cleanup, doing laundry, ironing, dusting, cleaning, etc.), other home production (home maintenance, outdoor cleaning, vehicle repair etc.), necessity shopping (grocery shopping, going to the bank, etc.), and time spent in education. We also add all travel time associated with any of those activities. Leisure, lastly, is defined as the sum of the following activities: entertainment/social activities/relaxing, child care and caring for other adults, gardening, time spent with pet, personal care, eating and drinking, recreational shopping, civic and religious activities, and own medical care. Again, all travel time associated with any of those activities is added to the total.[11] With these definitions, we compute the share of sleep, work, and leisure for the young and the old respectively. For the calibration, we ignore sleep and focus on the work–leisure trade off.

With this classification, we also want to know the share of leisure time spent on socially intensive activities (h_B vs. h_V) and the share of work time at the workplace (n_{Bw} and n_{Vw} vs. n_{Bh} and n_{Vh}). We define socially intensive activities as activities spent outside the home and correspondingly activities as *not* socially intense if they take place in the respondent's home or yard. We prefer this classification to the alternative "with whom" criterion, since we consider, e.g., the activity of going to the mall for recreational shopping to be a socially intense activity, even though the individual may go there on his/her own. Table 1 shows the time spent in different activities by young and old. As can be seen, despite spending much more time on leisure in total, the old spend approximately the same amount of time on socially intense leisure, i.e., leisure outside their home, as the young.[12]

Our definition of work includes market work, household work, core housework and home production, necessity shopping, and time spent in education. In the same way

[11]This definition of leisure is close to leisure "Measure 4" used by Aguiar and Hurst (2007). Compared to that definition, our leisure concept adds recreational shopping, gardening and time spent with pet, but excludes sleeping and education. In the category leisure shopping we include "Shopping, except groceries, food, and gas", "Comparison shopping", and "Researching purchases, n.e.c.".

[12]For more details about how people spend their time in in-puBlic and in-priVate leisure, see Supplemental Appendix A.2.

as for leisure, we classify all work activities according to where they were performed: in the home or outside home. Of the average working day, 31 percent of the time working is spent at home, and 69 percent outside home, mainly at the workplace. More information can be found in Supplemental Appendix A.1.

Calibration target for the size of the socially intense sector We use employment statistics on the 4-digit NAICS level from BLS to classify sectors in the US. The classification is based on if the sector is assumed to provide goods/services to the socially intense consumption-leisure bundle (and consequently if the workforce interact with customers). The extent to which the sector can be classified as socially intense can be fully (100%), to a high extent (75%), to a somewhat smaller extent (50%) or not at all (0%). We then sum up the affected workforce, and get that out of the total workforce (161,037,700 workers), 20% work producing for the socially active bundle. The lion's share (43%) of the workforce working in the socially intense sector is working in the accommodation and food services provision, followed by "all other retail" (18%) and "non-agricultural self-employed" (14%). An example of how the classification is done is given in Supplemental Appendix A.3.

4.2.2 *Calibration of economic parameters*

Production function, F The production function has as parameters the capital stocks k_B and k_V, the Cobb-Douglas parameters α and ν, and the CES aggregation parameters λ_n and ε_n determining the relative productivity of young/old and the substitutability between them. We set $\alpha = 1/3$ and the substitutability between young and old to $\varepsilon_n = 10$, thus mimicking close to perfect substitution while avoiding bang-bang solutions.

In the data, working hours are on average distributed between work from home (n_h) vs. in the workplace (n_w) such that $n_w/n_h = 2.3$. In the absence of an epidemic, the marginal products of n_h and n_w are equal, delivering $(1 - \alpha - \nu)n_h = \nu n_w$. Hence we obtain that $\nu \approx 0.202$.

We can use these values to assess how much production would be lost if n_w/n_h were forced to fall from 2.3 to, say, 1. The loss would be

$$\frac{3.3^{0.465} 3.3^{0.202}}{4.6^{0.465} 2^{0.202}} \approx 0.95,$$

i.e., output would fall by 5 percent. This is sizable, though not enormous. If n_w/n_h falls to $1/3$ (2 hours worked at the workplace plus 6 hours worked from home out of an 8 hours workday) the output loss is 25 percent. We find these losses reasonable in magnitude, supporting a choice of $\nu = 0.202$. Note also that the choice of production function implies that production cannot take place without at least some work in the workplace. Thus, even if the agents or the planner would want everyone to work from home (to suppress the spread of the virus) it is not feasible, and consequently there will always be some transmissions in the workplace during the course of the epidemic.

The capital stocks k_B and k_V are internally calibrated so that the return on capital is equalized across sectors in the pre-pandemic world, and such that the marginal product of capital net of depreciation is equal to an annualized interest rate of 4 percent. The relative productivity of young and old, determined by λ_n, is internally calibrated together with the preference parameters.

Family utility $\mathfrak{u}$ The family preferences are parametrized by the CES parameters $\lambda, \varepsilon, \lambda_B, \varepsilon_B, \lambda_V, \varepsilon_V$ and the "value of being alive" $\underline{u}$.

Given preference elasticities $\varepsilon, \varepsilon_B$ and ε_V, the other preference parameters are calibrated as follows. The outer CES weight λ is set to match the output share of the in-puBlic sector of 0.2. The inner CES weights λ_B and λ_V together with the productivity weight λ_n are jointly calibrated to match the time spent on leisure for the young (0.55), for the old (0.74), and the ratio of in-puBlic leisure and in-priVate leisure for the young (0.48).

The outer elasticity, ε, controls the elasticity between in-puBlic consumption and in-priVate consumption, but it also controls the elasticity between in-puBlic consumption and in-priVate leisure as well as between in-puBlic leisure and in-priVate consumption. How should the marginal utility of in-puBlic leisure be affected if I buy a new TV? We think a reasonable benchmark is not at all, which motivates the benchmark $\varepsilon = 1$, yielding additive separability between the two consumption-leisure bundles.[13]

To pick ε_B and ε_V, we put additional restrictions on the utility function. First, we require that the income effect dominates the substitution effect in a realistic way: we require that if the economy grows by 2%, hours worked should fall by approximately 0.4% (Boppart and Krusell, 2020). Second, we require that the young should spend a larger fraction of their leisure in the socially intense B activity than the old do, since that is what the data tells us. These two restrictions narrow down the set of permissible ε_B-ε_V combinations substantially. As our benchmark calibration, we use $\varepsilon_B = 0.41$ and $\varepsilon_V = 0.80$.[14]

As a sanity check of our calibration of the utility function we examine the implied Frisch elasticity for the young, which turns out to be 1.1. This might at first sound rather high, but given that the model includes also the very young (our definition of young starts already at the age of 15) and that the Frisch elasticity should correspond to not only the intensive margin elasticity but the aggregate elasticity including also the extensive margin, we think a value of 1.1 is reasonable.

The "value of being alive" $\underline{u}$ is calibrated to match estimates of the value of a statistical life. The value of a single time period, in our case a day, *VSTP*, is for a simple univariate utility function $v(c)$ given by the formula

$$\frac{VSTP}{c} = \frac{v(c)}{cv'(c)}. \tag{10}$$

where *VSTP*, the value of a single time period, is expressed in period-0 units of consumption, see, e.g., Conley (1976) and Shepard and Zeckhauser (1984).

[13] We acknowledge that this is not an obvious conclusion, and therefore perform robustness checks with respect to the value of ε, found in Supplemental Appendix A.4. Even though the details of the reallocation of time in the event of an epidemic change, the substantive conclusions remain. Moreover, we argue that the reallocations with the benchmark $\varepsilon = 1$ seem plausible.

[14] See Supplemental Appendix A.5 for more details.

TABLE 2. Calibration targets for the economic parameters in the static problem.

Target	**Parameter**
Marginal product of capital equal across sectors Marginal product of capital net of depreciation = interest rate	$\}k_V, k_B$
Output share of public sector = 0.2	λ
Leisure young = 0.55 Leisure old = 0.74 in-puBlic leisure/in-priVate leisure for young = 0.48	$\}\lambda_V, \lambda_B, \lambda_n$
Hours worked in the workplace vs. from home ($n_w/n_h = 2.3$)	ν
Realistic income vs. substitution effects $h_B^y/h_V^y > h_B^o/h_V^o$	$\}\varepsilon_B, \varepsilon_V$
Chosen value for *VSTP*	$\underline{u}$
Exogenously picked values (see text)	$\alpha, \varepsilon_n, \varepsilon$

The parameters associated with each target are the parameters which primarily determine the target.

The intuition for the equation is straightforward: the utility value of an additional period of life is $v(c)$, the per-period flow utility. The utility value of an additional unit of consumption is $v'(c)$. The marginal value of an additional period of life, in terms of consumption, is thus $\frac{v(c)}{v'(c)}$. *VSTP*$/c$ is the value of an additional time period, expressed as a multiple of per-period consumption.

We use the young generation's equilibrium allocations of goods and time and adjust $\underline{u}$ so that the above equation is satisfied for a given estimate for *VSTP*.[15]

A range of different values for a time period has been used in the literature (see, e.g., Viscusi and Aldy (2003) for an overview). For our baseline scenario we assume that a period of life is worth 6 times period consumption, following Hall et al. (2020) who base their number on data from The Environmental Protection Agency (EPA). As a robustness and to evaluate the importance of this assumption, we will also use a higher number; for this we will use 11.4, following Glover et al. (2023) who also base their number on data from EPA and the U.S. Department of Transportation.

A summary of the targets used for the calibration of the static economic problem is given in Table 2.

4.2.3 *Calibration of epidemiological parameters* We calibrate based on information available at the outset of the pandemic. In Section 7 we consider a richer structure.

Infection rate function G The infection rate function has as parameters κ_B and κ_V. We consider all interactions equally contagious and set $\kappa_B = \kappa_V$ and calibrate them to match a chosen value for R_0. The estimates of R_0 for covid-19 are uncertain and range between at least 1.4 and 3.9; we use 2.0 in our benchmark simulations.

[15]Given that the young and the old have different time and goods allocations, in theory it matters which type is selected. In practice, however, the difference between $\underline{u}$ based on the allocations of the young or the allocations of the old is small.

We simulate the simplest possible SIR model with a homogeneous population given this estimate of R_0. This gives us a measure of the final number of recovered (which is 78%) if the epidemic were to play out unhindered. Thereafter we use the steady-state time allocations in our calibrated model and find the $\kappa_B = \kappa_V$ yielding the same final number of recovered as in the economy without endogenous behavioral responses.[16]

4.3 *Calibration of the terminal condition*

Life expectancy, $\mathcal{T}^i$ The young and old live up to period $\mathcal{T}^y$ and $\mathcal{T}^o$, respectively. We calibrate $\mathcal{T}^y$ and $\mathcal{T}^o$ to match the remaining life expectancies based on the group definition. Note that a perfect estimate of this would take into account the mortality profile by age within each age group, and weight the conditional life expectancy by that. Such an estimate would require further estimates, so instead we set the average age of a deceased in the young group to 50 years, and the corresponding age in the old group to 80. This implies a remaining life expectancy of 31.6 years for the young, and 9.2 years for the old (Arias and Xu, 2019).

Flow utility after the terminal period, $\tilde{u}$ We set the flow utilities after the terminal period, $\tilde{u}^y$ and $\tilde{u}^o$, equal to the flow utilities of the young and the old absent the epidemic.

5. Results from the baseline model

In this section we contrast how the epidemic plays out in a planner allocation scenario compared to in a competitive equilibrium. In the competitive equilibrium, the representative family either has full information about the epidemic and acts accordingly ("rational expectations") or do not understand the link between its own actions and the risk of getting the disease ("myopic").

We can consider two scenarios for the end of the epidemic: either the epidemic ends endogenously and gradually when the population eventually reaches the herd immunity threshold so that the effective reproduction number goes below one, or there is an exogenous end to it, via the arrival of a vaccine or cure. In this section we present results from a baseline scenario in which the epidemic ends endogenously. This is not because we think such a scenario is more realistic, rather, we prefer this as a baseline scenario since it highlights the trade-offs in a transparent way and facilitates the understanding of the model. In the next section we will explore how the arrival of a cure changes the dynamics.

5.1 *Evolution of the epidemic*

Figure 2 shows the evolution of the epidemic under the myopic competitive equilibrium, the rational-expectations competitive equilibrium, and the planner's allocation under the assumption that no cure or vaccine arrives.

[16]The resulting epidemiological spread is very close but not exactly the same as the SIR model with homogeneous population simulated initially, since young and old have different time allocations and different death rates.

TABLE 3. Summary of calibrated parameters.

Parameter	Description	Value
Preference parameters		
β	Discount factor	$0.96^{1/365}$
λ	Weight on $\tilde{c}_B$	0.26
λ_B	Weight on c_B	0.92
λ_V	Weight on c_V	0.71
ε	Elasticity between $\tilde{c}_B$ and $\tilde{c}_V$	1.0
ε_B	Elasticity between c_B and h_B	0.41
ε_V	Elasticity between c_V and h_V	0.80
VSTP	Value of a statistical time period (as a multiple of period consumption)	6.0
$\underline{u}$	Intrinsic value of life per period (in utils)	2.5
Technology		
α	Capital share	1/3
ν	Home work labor share	0.202
ε_n	Elasticity of substitution between young and old	10
λ_n	Production weight on young	0.62
k_B/k_V	Relative capital stock	0.25
Demographics		
Pop_0^y	Fraction young	0.73
$\mathcal{T}^y$	Remaining life time young	$31.6 \cdot 365$
$\mathcal{T}^o$	Remaining life time old	$9.2 \cdot 365$
Epidemic variables		
R_0	Spread factor standard SIR model	2.0
$\kappa_B = \kappa_V$	Spread factor economic model	0.24
π_r	Recovery rate	1/18
$\pi^i_{d,low}$	Death rate (before overcrowding) [young, old]	[0.001, 0.025] · 1/18
$\pi^i_{d,high}$	Death rate (when overcrowded) [young, old]	[0.003, 0.075] · 1/18
Healthcare system		
ι_h	Fraction of infected in need of hospitalization	0.03
ι_i	Fraction of hospitalized in need of ICU	0.29
ι_b	Inhabitants per ICU bed	3,400
$\hat{I}$	Midpoint logistic function (fraction infected)	1 / $(\iota_h \cdot \iota_i \cdot \iota_b)$
k	Steepness parameter	1,000

See text for description of sources and methodology.

The epidemic in the myopic competitive equilibrium is close to standard SIR dynamics. The health system is overloaded, many young and old get infected, and the epidemic is essentially over because of herd immunity after 300 days.

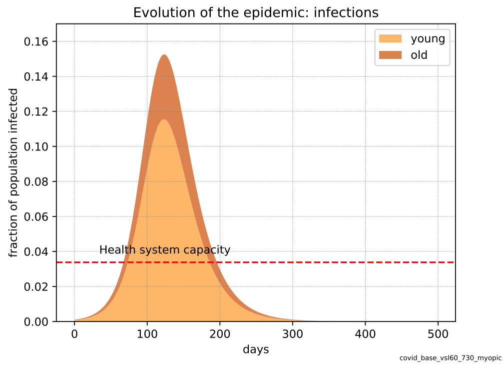

(a) Myopic competitive equilibrium

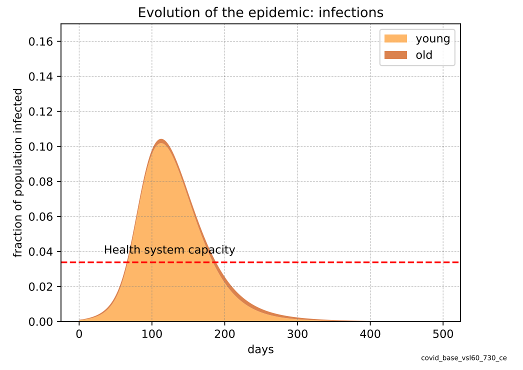

(b) Rational-expectations competitive equilibrium

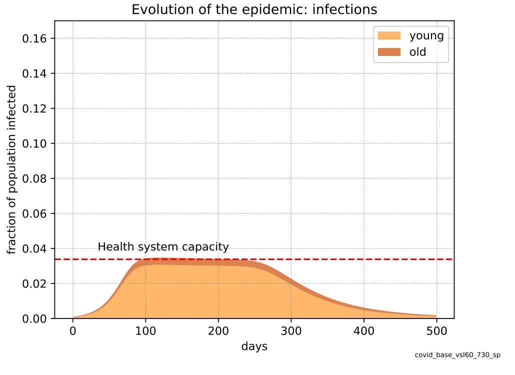

(c) Planner's solution

FIGURE 2. The evolution of the epidemic (fraction of population currently infected) in competitive equilibrium with two different assumptions about information (myopic vs. rational expectations) and the planner's allocation.

The evolution of the epidemic in the rational-expectations competitive equilibrium may at first pass seem similar. The epidemic is essentially over after 300 days. However, with rational expectations, few old become infected. The epidemic is primarily a risk for the old and with rational expectations they shift their behavior away from activities associated with infection risk. The young also do so, but to a much lesser extent, both because their risk is lower and because the labor-market wages give a compensating differential to the young. That the health system is overloaded does not affect the behavior of the young to any larger extent; they perceive their death rate as sufficiently low, even with overcrowded hospitals, to continue being socially active.

In the planner's solution, the evolution of the epidemic is qualitatively different. The planner "protects the healthcare system" and prolongs the epidemic. The planner internalizes the effects of an overloaded health system and keeps infections below the thresh-

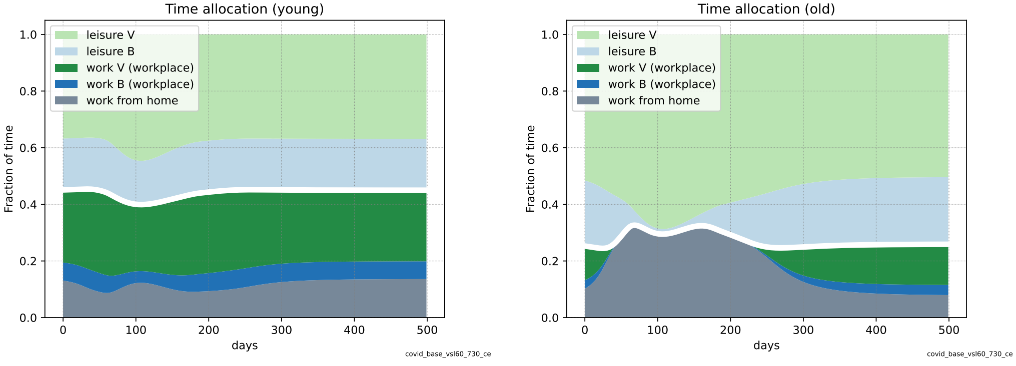

(a) Young **(b)** Old

FIGURE 3. Time allocation in rational-expectation competitive equilibrium.

old for overloading.[17] Therefore, it takes a longer time to reach herd immunity, and the epidemic is essentially over after 400 days.

As expected, the final number of recovered is the lowest in the planner scenario: a planner ensures that the herd immunity threshold is reached with the smallest amount of people getting infected in total. In the planner scenario, the final fraction of recovered is 55%, to be compared to 62% in the rational-expectations scenario and 78% in the myopic scenario. In other words, a planner avoids "over-shooting" in the number of infected and subsequently recovered.[18]

5.2 *Time allocations during the epidemic*

We now unpack the time allocations under the different scenarios. Figure 3 shows the time allocations in the competitive equilibrium with rational expectations.[19]

RESULT 1. (*Fear factor among the old*) In the rational-expectations competitive equilibrium, the old avoid all social activities.

Under the rational-expectations competitive equilibrium, the old completely stop working in the workplace from the beginning of the epidemic. They also largely stop spending time on in-puBlic leisure and instead spend their time enjoying more in-priVate leisure and working from home (recall that our definition of work includes various household activities such as cleaning and cooking). The young adjust their behavior

[17] Compare equilibrium condition 2 (in the competitive equilibrium) with optimality condition 2' (in the planner solution) in Section 2.

[18] In a standard SIR model without heterogeneity and deaths, the minimum number of recovered needed to reach herd immunity is $1-1/R_0$. With a basic reproduction number $R_0 = 2.0$ this implies 50% recovered, close to the value in the planner's solution, 55%.

[19] In myopic competitive equilibrium, the time allocations hardly change over time and are therefore not shown. Technically, there is a slight change due to deaths in the population. However, these changes are so small that they are not visible in these graphs and can be disregarded.

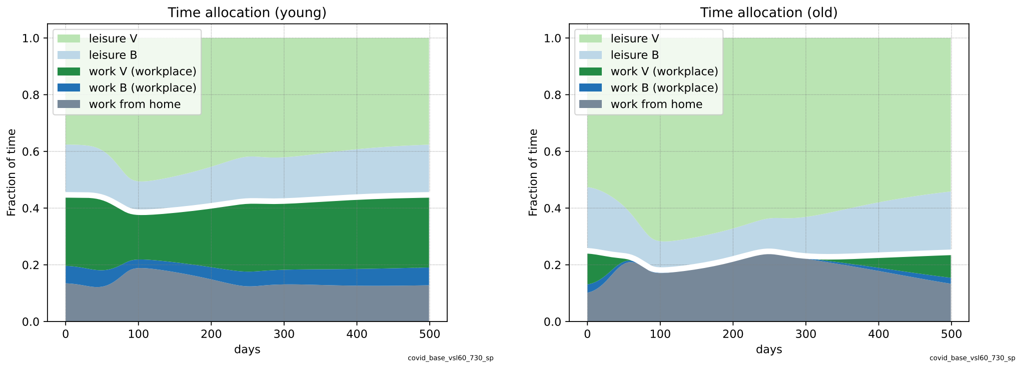

(a) Young **(b)** Old

FIGURE 4. Time allocation in the planner scenario.

as well, but to a much smaller degree. During the peak of the epidemic, when the health system is overloaded, they reduce their work in the workplace slightly but otherwise they keep their behavior relatively constant.

Figure 4 shows the time allocations in the planner scenario where, as we saw above, the time allocations are fine-tuned so as not to overload the hospital system.

RESULT 2. (*The happy old*) In the planner scenario, the old can enjoy some social leisure.

From very early on in the epidemic, the old stop working in the office. However, because infection rates are kept at a moderate level, they can still spend some time on in-puBlic leisure, and they therefore spend considerably less time working from home compared with the rational-expectations competitive equilibrium. The infection rates are kept lower because of the behavioral changes among the young. The young work less in the office and decrease their in-puBlic leisure as well, internalizing the externality of them becoming infected and subsequently infecting others.

5.3 *Aggregate variables during the epidemic*

In Figure 5, we compare the aggregate impact of the three different allocations.

RESULT 3. (*Low death rates also in well-informed markets*) The rational-expectations competitive equilibrium and the planner both radically reduce the number of deaths.

In Figure 5a, we plot the number of cumulative deaths under the three scenarios. Under the myopic market allocation, almost 1.2 percent of the population dies. The rational-expectations competitive equilibrium improves on this outcome significantly, reducing the number of deaths to 0.254%. The well-informed self interest of the old is sufficient to significantly reduce the number of deaths. In the planner scenario, the total death toll is just slightly below the rational-expectations scenario: 0.247%. However, the composition of deaths during the epidemic is very different. In the rational-expectation

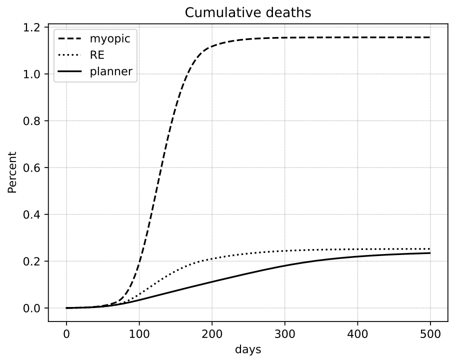

(a) Cumulative deaths

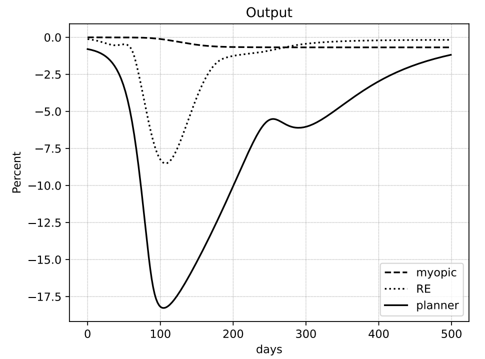

(b) Output

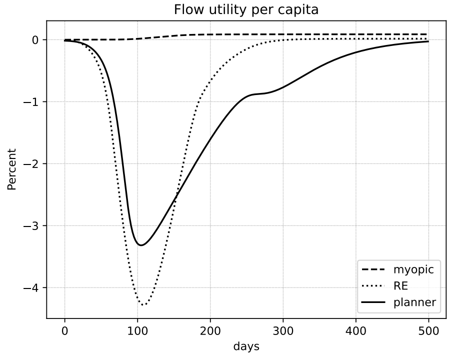

(c) Flow utility

FIGURE 5. Comparing aggregate outcomes in the three different scenarios.

scenario 44% of the deaths are in the young population. In the planner scenario, 20% of the deaths are in the young population. As discussed earlier, the final number of recovered is lower in the planner scenario than in the rational-expectations scenario. Thus, even though the total number of individuals who have once been infected is lower in the planner scenario, the total number of deaths is almost equal since there were more old individuals among the infected, and the old have a higher death rate. The difference in death rates between young and old is larger than the effect of hospital overcrowding on the death rates of the young. A planner saves young lives compared to a rational-expectations scenario, which translates into more *years* of life saved.

RESULT 4. (*The planner cuts output the most*) Output falls modestly in rational-expectations competitive equilibrium, and much more in the planner's allocation.

Figure 5b shows the corresponding responses of aggregate output. Under the myopic market allocation, output is virtually unaffected throughout the epidemic. As the population shrinks, output mechanically falls slightly. The rational-expectations competitive equilibrium has a modest fall in output during the peak of the epidemic but the annual

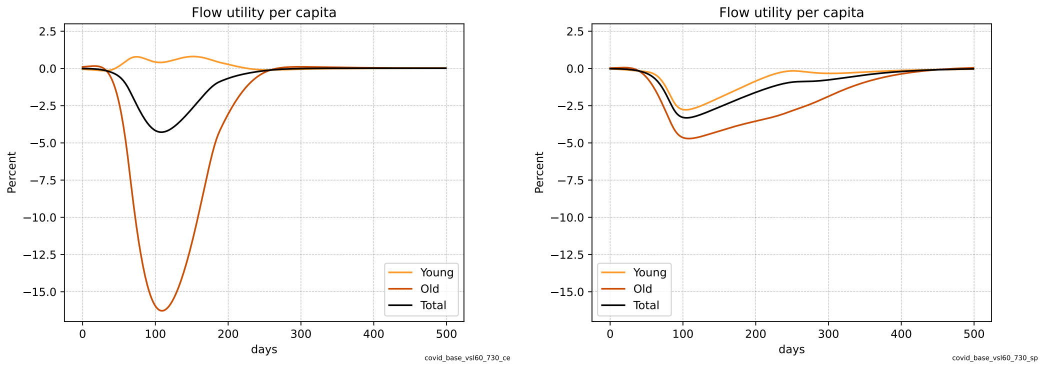

(a) Rational-expectations competitive equilibrium **(b)** Planner's allocation

FIGURE 6. Flow utility under the rational-expectations and the planner's allocations.

drop in output is a mere 2.3 percent. The planner is willing to reduce output much more than either market allocation. At the peak of the epidemic, output drops by 17 percent; during its first year, output falls by 8.7 percent.

RESULT 5. (*The importance of social leisure*) The policy trade-off is not only output vs. deaths: social leisure plays a key role.

Figure 5c shows the percentage drop in per-period per-capita utility. The myopic flow utility is essentially unaffected by the epidemic, increasing slightly due to the deaths (and the constant capital stock, leading to higher output per capita). For both the rational-expectations competitive equilibrium and the planner's allocation, flow utility per capita drops substantially during the epidemic. Note however that the planner's allocation implies a smaller fall in flow utility than the rational-expectations competitive equilibrium. In the rational-expectation competitive equilibrium, the old are essentially prohibited from any socially active activity and their utility is significantly reduced. Their utility loss is not captured by output since it is a loss of valuable leisure, not consumption, but it is a welfare loss nevertheless. It is tempting to frame a discussion of epidemic policy as a trade-off between the economy, as captured by output, and lives. This way of framing the trade-off misses that the planner is willing to sacrifice consumption utility not only in order to save lives but also to save leisure utility for the old.

Figure 6 unpacks the flow utility for the young vs. the old for the rational-expectations competitive equilibrium and the planner's allocation. It should be noted that the flow utility per capita is calculated as the flow utility per person alive in respective group; thus, the effects of deaths are not visible from these graphs, but only the instantaneous utility for the individuals who are alive.

RESULT 6. (*The planner protects the old*) Compared to the market outcome, the planner's allocation benefits the flow utility of the old at the expense of the young.

In the rational-expectations competitive equilibrium, shown in Figure 6a, it is clear that the loss in average per-capita flow utility is completely driven by the old, whose flow utility decreases by more than 10% during some critical weeks when the infection rate in the society is at its peak. During those critical weeks, the old, due to the high risk of getting infected, choose to stay at home and hardly enjoy any in-puBlic leisure at all, which drives down their utility substantially. As the old drastically reduce their B good consumption, the price of the B good falls. This benefits the young, and they even see a slight increase in utility. Further, Figure 6b shows the flow utility in the planner's solution. The planner distributes the burden of behavioral adjustment more efficiently. The young now also take a hit, with flow utility decreasing during the epidemic. However, since the old are so much better off relative to the rational-expectations scenario, the drop in average flow utility conditional on survival during the peak of the epidemic is not as severe as in the rational-expectations competitive equilibrium.

6. Alternative scenarios and parameter values

We now analyze optimal policy under different assumptions: the arrival of a cure, the value of a statistical life, and the severity of hospital overcrowding.

Result 7. (*Key assumptions*) Optimal policy can, and indeed does, change qualitatively with different assumptions about key parameters such as the belief about when a cure will arrive, the value of a statistical life, and the severity of overcrowded hospitals.

6.1 *A cure expected soon: it is optimal to "lock down"*

We first consider the situation when there is an exogenous and known end to the epidemic. We model it as the arrival of a perfect cure that is instantaneously adopted, which means that we assume that everyone infected is immediately cured from the disease once the cure arrives. Two alternative endings to the epidemic are (i) an instantaneous arrival and distribution of a vaccine or (ii) a perfect implementation of testing and tracing.[20] From a modelling perspective, both scenarios are very similar to the arrival of a cure, with the difference that with the arrival of a vaccine or testing and tracing, those infected at the point of the arrival can still die from the disease. In practice, this makes a very small difference in the model. We consider a scenario where it is known that a cure arrives after one year.

The epidemic evolutions for the rational-expectations competitive equilibrium and the planner's allocation under this scenario are shown in Figure 7. The rational-expectations competitive equilibrium is largely unaffected by the arrival of the cure, since it arrives after the epidemic is finalized and the agents' behavior is already consistent with the absence of an epidemic at this point. However, the planner's allocation qualitatively shifts towards full suppression of the epidemic if the cure arrives early enough. As Figure 7b shows, if the cure arrives within a year, the planner's strategy shifts

[20]When a planner is capable of perfectly identifying who is infected, in our model framework the planner would let the infected stay at home until they are recovered, and the epidemics would die out quickly.

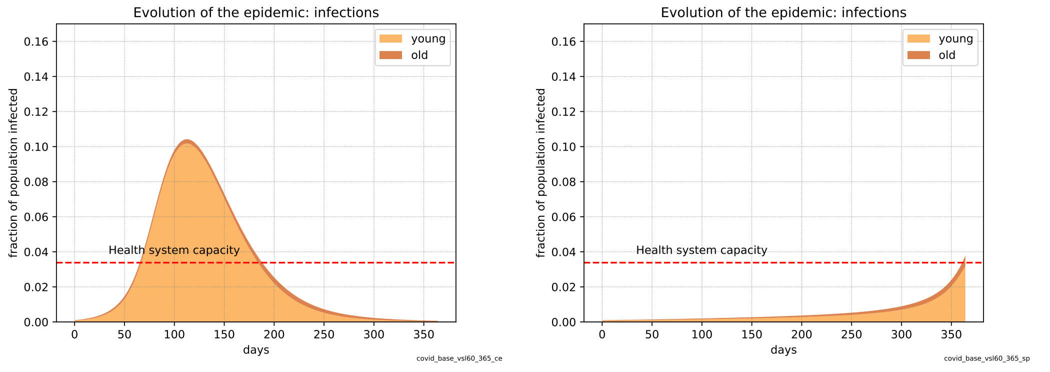

(a) Rational-expectations competitive equilibrium **(b)** Planner's allocation

FIGURE 7. The evolution of the epidemic in rational-expectations competitive equilibrium and under the planner's allocation when a cure arrives after one year.

qualitatively: from a strategy best described as "protect the healthcare system" to a "suppression" strategy. Anticipating that the cure is instantaneously and perfectly distributed after exactly a year, the planner allows for an increase in the number of infected near the day when the cure is distributed. This anticipation would disappear if, e.g., only a hazard arrival rate is known.

6.2 *A higher value of a statistical life: lock-down again*

We now highlight the importance of the choice for the value of a statistical life period for the optimal solution, as it too can qualitatively change the planner's solution. To illustrate this, we consider a scenario where the cure is known to arrive after 15 months. Using the same value of a statistical life as before, the planner's strategy is back to "protect the health care system". However, if we instead use a higher value of a statistical life period (choosing 11.4 times period consumption, following Glover et al. (2023)), the optimal strategy is to suppress the infection by harsher measures. This is illustrated in Figure 8, which shows the different approaches taken by a planner, depending on which value is assigned to a statistical life.

The "suppress" strategy, adopted in the scenario with the higher value of a statistical life, saves lives: only 0.042 percent die in this scenario, while the final death toll in the scenario where the planner has a lower value of a statistical life is 0.224 percent. However, the suppress strategy is of course costly in terms of output (and utility). Output falls by 19.9 percent over the first year, while the output fall in the protect scenario is 10.1 percent during the first year.

6.3 *Less costly hospital overcrowding: speed up infections*

We now turn to the importance of the overcrowding function. That a planner wants to suppress the epidemic to avoid overcrowded hospitals of course hinges on the assumption that overcrowding is a severe problem, and different assumptions about the severity of overcrowding of hospitals can qualitatively change the planner's solution.

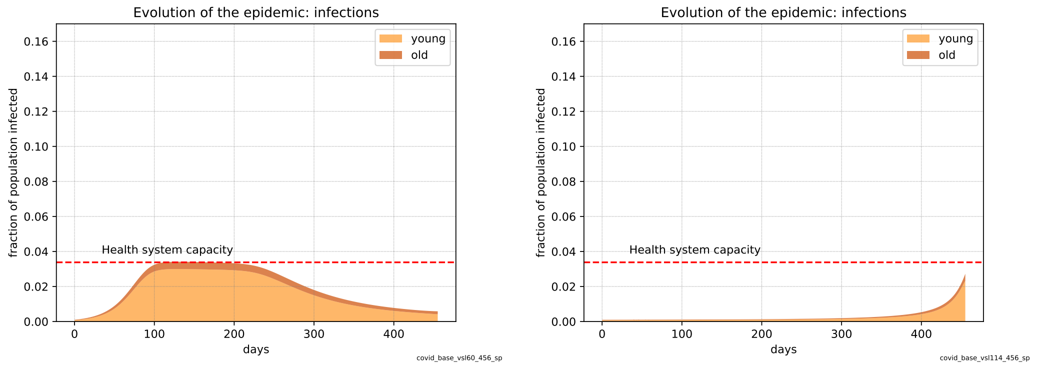

(a) With the baseline value of a statistical life **(b)** With a higher value of a statistical life

FIGURE 8. The evolution of the epidemic under the planner's allocation with two different assumptions about the value of a statistical life, assuming a cure arrives after 15 months.

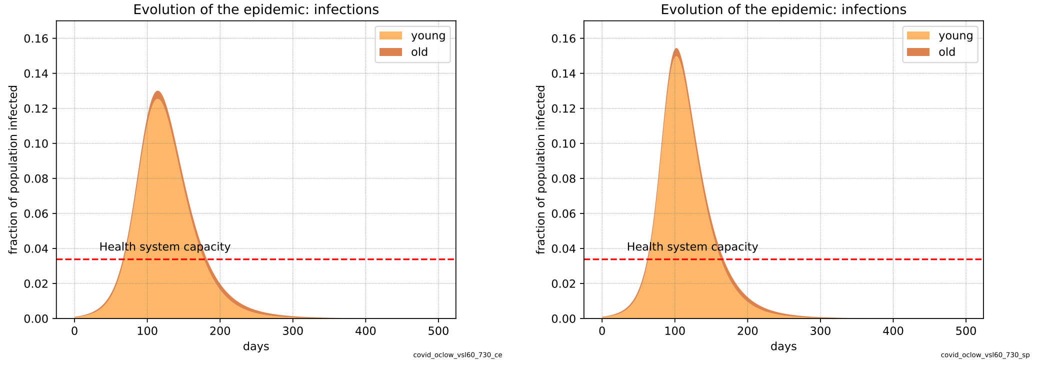

(a) Rational-expectations competitive equilibrium **(b)** Planner's allocation

FIGURE 9. The evolution of the epidemic, assuming death rates increase by only 20% once the hospital system is overcrowded.

To illustrate this, we consider a scenario where the death rate for those infected once the hospitals are overcrowded increases by only 20% (compared to 200% in the baseline). The epidemic evolution for the rational-expectations competitive equilibrium and the planner's allocation is shown in Figure 9.

In the rational-expectations scenario, people are slightly less afraid of getting infected during the peak of the epidemic, and the epidemic evolves somewhat faster than in the baseline scenario (compare to Figure 2b). However, the planner allocation shifts qualitatively: the planner implements a strategy similar to the rational-expectations outcome: let the young get infected quickly, to reach herd immunity, while the old are shielding. By speeding up the process the planner shortens the time the old have to shield at home and suffer.[21]

Note that the planner even wants the spread among the young to be faster than in the myopic model (compare Figure 9b to the lighter shaded area in Figure 2a), which implies that the planner would need to deliberately increase social activity of the young over and above the steady state “no-epidemic” levels. The implementation of such a policy would involve encouraging and facilitating the spread in the population. One way would be to arrange spread events, similar to historical “measles parties” or “pox parties”, used by parents before there were vaccines to increase the probability for their children to get infected and develop immunity at an early age (since such infections are more dangerous later in life).

7. Extended model

In this section, we enrich the model along three dimensions relevant to the covid-19 epidemic. In part, we use the richer model as a proof of concept for our methods.

First, we introduce waning immunity as in Giannitsarou et al. (2021): after some time, individuals who are recovered from infection lose their immunity and become susceptible again. In concrete terms, the law of motion for the epidemic now includes an exogenous transition rate from recovered to susceptible, which we set to $\pi_s = 1/730$, i.e., on average immunity lasts for two years.

Second, we introduce an exogenous seasonality component in the contagiousness of the disease. It is by now well understood that covid-19 spreads faster during the winter season than during summer, even though the underlying reasons for this are not entirely established.[22] In the model, we introduce this by letting the contagiousness parameters κ_B and κ_V be time varying. In the calibration, we had $\kappa_B = \kappa_V = 0.24$. Now, we let $\kappa_{B,t} = \kappa_{V,t} = 0.24 \cdot (1 + 0.5 \sin(2\pi t/365))$ so that the exogenous contagiousness varies at an annual frequency and the winter contagiousness is three times the summer contagiousness.

Third, we introduce an additional type which we label “very old”. We split the non-young (population fraction 0.27) into two types of equal size: the old and the very old. In short, the very old are even more vulnerable to the disease, have an even lower productivity, and even fewer years left to live absent any epidemic. The remaining life time of the old is $1.5 \cdot 9.2 \cdot 365$, i.e., 50% higher than the previous target, and the life time of the very old is $0.5 \cdot 9.2 \cdot 365$, i.e., 50% lower than the previous target. In the CES aggregator of labor, the production weight on the young was 0.62 and on the old 0.38. We keep these production weights but add the very old inside the CES aggregator, with productivity

[21] That the planner speeds up the spread of the disease compared to the rational-expectations scenario is a manifestation of what Garibaldi et al. (2020) call “immunity externality”: that agents shield too much in the hope that others catch the disease and reach herd immunity.

[22] The higher spread during the cold season could for instance be due to the fact that people spend more time socializing indoors, and/or that dry air dries out also the tissues lining the airways, thereby increasing the risk of getting infected. See Atkeson (2021) for a further discussion and an epidemiological framework with seasonality.

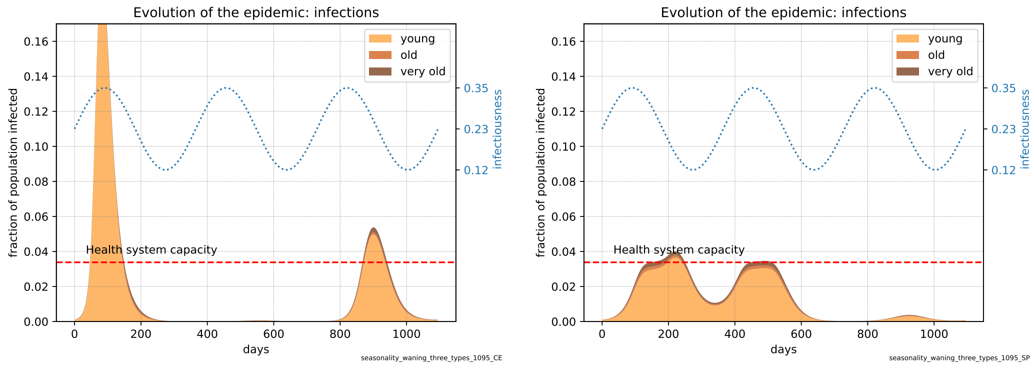

(a) Rational-expectations competitive equilibrium **(b)** Planner's allocation

FIGURE 10. The evolution of the epidemic in the extended model with waning immunity, exogenous seasonality, and three types. The dashed blue line (right-hand side axis) shows how the contagiousness parameter varies over time.

weight $0.5 \cdot 0.38 = 0.19$. Finally, the death risk of the very old is 50% higher than the death risk of the old, which is kept at the previous target.

In the following experiments, we assume that an instantaneous cure arrives after three years. We summarize the results of the experiments as follows:

RESULT 8. (*New policy trade-offs*) Waning immunity, seasonality and additional heterogeneity yield rich epidemiological dynamics and uncover new policy trade-offs.

7.1 *Results from the extended model*

Figure 10 shows the epidemic evolution in the extended model for the rational expectations competitive equilibrium and the planner's allocation.

As the figure shows, with seasonality the competitive equilibrium has a severe first wave. However, since many people get infected in this first wave, and the immunity is assumed to last on average for two years, there is no second wave one year later. However, in the third year, there is a new outbreak, since enough people have lost their immunity by then.

The planner, on the other hand, finds it optimal to smooth out the infections, and thus there is both a first wave and a second wave a year later. Towards the end of the first wave, it is optimal to temporarily exceed the health system capacity when the contagiousness of the epidemic is declining as summer is approaching. As a result of the declining contagiousness, the externality cost of having young individuals infected falls and in fact becomes negative. The reasoning behind this is straight-forward: if we need to have many individuals infected at some point, it is better to have it when the contagiousness is low, since then the risk of infecting others is diminished. Overall, the planner "protects the healthcare system" as before, but the infection curve follows seasonality. After six hundred days, herd immunity has been reached, and the third wave is just a small bump due to waning immunity.

In terms of time allocations, the main picture from the baseline model remains for the young and old group: in the rational-expectations scenario the young hardly change their time allocations, while the old all but completely stop working in the office and enjoying in-puBlic leisure. The new group, the very old, is so unproductive that they do not work at all even in pre-pandemic times, but enjoy a mix of in-priVate and in-puBlic leisure. In the rational-expectations competitive equilibrium, the very old shift over to more in-priVate leisure during the epidemic, but, unlike the old group, they do not entirely cut down on in-puBlic leisure.

In the planner's allocation, the young adjust their behavior more than in the rational-expectations scenario: they start working from home and cut down on the in-puBlic leisure to a higher extent when the planner can decide their actions. Thus both the old and the very old can enjoy more in-puBlic leisure than in the rational-expectations scenario. A full set of graphs describing the time-allocations in the extended model can be found in Supplemental Appendix B.1.

8. Other epidemics

The model can be applied to other epidemics than covid-19. Studying multiple diseases provides *cross-epidemic restrictions*: parameter values ought to be consistent with observed behavior for other diseases, not only covid-19. We choose a regular seasonal flu in order to test an epidemic that we actually see every year, and SARS to test one that is substantially more deadly than covid-19. We test both epidemics under the assumption that the epidemic is over after one year, either because the population has reached herd immunity, or because a cure/vaccine/perfect test and trace arrives and puts an end to the epidemic.

The seasonal flu To simulate a seasonal flu we set the effective reproduction number at the beginning of the flu season, R_t, to 1.3, use a death rate of 0.045%, and set the average number of days until recovery at 10. This corresponds to a regular flu season, not to a year with a particularly severe instance of the flu.[23]

Result 9. The seasonal flu simulations indicate that a relatively low value of a statistical life is more in line with observed policy actions.

When simulating a seasonal flu with a high value of a statistical life period (again choosing 11.4 times period consumption, following Glover et al. (2023)), a planner would want to lower output by 4.0 percent during the second quarter of the epidemic, and the annual drop would be 2.8 percent. As far as we can tell, this is not how policy makers have reacted historically. With the lower value of life (our baseline value), the annual drop in output is substantially lower: 1.1 percent.

[23]We verify that the chosen parameters are reasonable by comparing the death toll in our model to the actual number of deaths due to the flu in the US each year, for more details and sources for our parameters see Supplemental Appendix B.2.

SARS For SARS, we use $R_0 = 2.4$, death rates of 8% and 52% for the young and old respectively, and an average number of days until recover of 12, following Petersen et al. (2020). Both the market and planner allocations offer interesting insights.

RESULT 10. For SARS, the effective reproduction number hovers around 1 in rational-expectations competitive equilibrium. This is in contrast to covid-19, which, with age-heterogeneity, is not dangerous enough for the young to create this type of response.

SARS is perceived as dangerous enough that individuals endogenously choose to lower the amount of infectious activities. The precautionary behavior is increasing in the infection risk, which is increasing in the number of infected. However, the number of infected is decreasing in the strength of the precautionary response. The infection rate therefore stabilizes around a level which is consistent with the precautionary behavior. The same type of qualitative effect—an effective reproduction number that hovers around 1 in a market scenario—is also reported by Farboodi et al. (2021) and Bognanni et al. (2020). This result is in contrast with our simulations of the covid-19 epidemic, in which we do not find that the effective reproduction number stabilizes around 1 in the rational-expectations competitive equilibrium. Age heterogeneity is key to understanding this result: for the young, the risk of an infection does not provide a strong enough motive for a precaution that would stabilize the infection rate.

RESULT 11. Under SARS, the planner's allocation both saves lives and leads to a smaller fall in output compared to the rational-expectations competitive equilibrium.

A planner would quickly lower the amount of infectious activities to get the epidemic under control, and would thereafter not have to reduce the activities as much. In the rational-expectations competitive equilibrium people would carry on with their activities until the number of infected has increased substantially. At that point individuals would be so afraid of the epidemic that they endogenously restrict their activities greatly. The total fall in output would therefore be higher in the rational-expectations scenario. Qualitatively, this is the same type of mechanism as found in Aum et al. (2021). Again, in our calibration of the covid-19 epidemic, we do not find this effect since the covid epidemic is not perceived as dangerous enough by the young. For more information about these experiments, including details about calibration and simulation results, see Supplemental Appendix B.2 for the seasonal flu and Supplemental Appendix B.3 for SARS.

9. CONCLUSIONS

We have presented a general framework for integrated epi-econ assessment—a toolbox for policy evaluation—that we believe is an appropriate starting point when confronted with pandemic challenges. The model is built around covid-19 and we also conduct policy analysis for this case. We now emphasize some key features and outline a number of extensions. Finally, we outline how we recommend using our setting in concrete terms as a new pandemic threat surfaces.

In our model construction, we have emphasized the importance of a sociology component. We thus connect economic and social-activity decisions and show how time-use data can be employed to discipline the parameter selection. One of our key findings—not a foregone conclusion, due to our quantitative discipline—is that the value of social life can matter significantly when designing optimal policy. This is the case in our covid-19 application and even more so when we study the seasonal ordinary flu; however, when we look at SARS, the direct health consequences become more important. Throughout our covid-19 application, we list key insights comparing market and optimal-planning allocations, many of which are nontrivial and emphasize the non-linearity of the problem under study: outcomes can change qualitatively as some key parameters are changed.

In Section 7, we enrich the baseline model—which by design is minimal ("dichotomous") along key dimensions—in a number of ways, e.g., by including seasonality in the virus spread and more detailed demographics. Further elaborations are entirely feasible. For example, it is easy to introduce waves due to virus mutations, possibly with higher and higher contagiousness but lower fatality rate, thus generating a setting where a pandemic gradually turns into a (seasonal) endemic disease. As described already in our study of seasonal virus variation, new trade-offs emphasizing the timing of policy would then surface.

Relatedly, another straightforward extension is the introduction of uncertainty; uncertainty is, moreover, critical especially at the early phases of a pandemic. The hallmark of the rational-expectations literature in macroeconomics is to endow agents (and any policymaker) with a correct probabilistic assessment of all outcomes. However, to make those assumptions here does not appear reasonable. Instead, we recommend conducting separate social planning evaluations for a number of different assumptions on the virus and its consequences (and possible cures), thereby spanning a large range of outcomes. To look at market outcomes, we would recommend—as we do here—separately looking at both myopic and perfect-foresight paths, so as to bracket a range of private-sector reactions to the virus.[24] These analyses can be carried out using foreseen and unforeseen transition paths ("MIT shocks").

Our model framework is a frictionless neoclassical economic benchmark, allowing a clear comparison of the competitive equilibrium and the optimal outcome. Departures from this benchmark are possible. For example, following Eichenbaum et al. (2022a), it is straightforward to include price and/or wage frictions in our framework. This would amount to introducing price and wage setters in the competitive equilibrium, with corresponding forward-looking Phillips curves. Such extensions open up for studying optimal fiscal and monetary policy during an epidemic, which is of particular importance for economic policy makers.

A conceptually different extension but one that is also possible—and feasible—to pursue is to introduce a further element of sociology, one that goes beyond tying interaction to economic decisions: people interact in networks (family, friends, acquaintances,

[24]A straightforward extension is to consider two groups of agents within the economy: one well-informed and one myopic.

etc.) with different meeting intensities and frequencies.[25] We did not add such a feature yet simply because we are unsure of how to connect it to individual data (Azzimonti et al. (2020) offer a promising avenue). Unlike in our setting, where social interaction is a byproduct of economic decisions, such a framework would also involve pure externalities that are present even in the absence of epidemiology.

A dimension where our modeling is very stylized and harder to extend regards household-level, idiosyncratic information: in real life, individuals update their own virus status—from our perspective, the subjective probability of being in states S, I, and R, respectively—based both on the degree and timing of interaction chosen with others (these are choice variables) and on symptoms (fever, head ache, etc.); at the same time, individuals are much less informed of the corresponding probabilities of others. A full treatment of this sort (with rational Bayesian updating), involving a state variable consisting of a full distribution of endogenously evolving probabilities, is both conceptually and numerically extremely challenging. The statement and assumptions underlying an associated planning problem are even more daunting. However, simple short-cuts are feasible to pursue. For example, one can assume a publicly observable individual state (like fever) that follows a Poisson process exogenous to the individual and that entails higher infection probabilities.[26] One can also consider ad-hoc learning, such as adaptive expectations. As for medical treatments, we show that it is easy to incorporate vaccination in Boppart et al. (2022).

Turning to the practical use of the present setting, we see our framework as a skeleton setting to be filled with more detail as new data arrives. The skeleton is rich enough already to consider a number of possible cases that are believed possible given any early information about the virus: based on different scenarios, we can compare market outcomes, given different assumptions about what markets expect, with optimal outcomes. As the new disease is revealing itself and its implications for health outcomes, along with updated expectations of possible cures, these new data should continuously make us update the key model parameters. The new information should also make us revise the model by adding any heterogeneity and new features—such as the possibilities mentioned earlier in the present section—that seem to become relevant.[27] With the model framework we have proposed here, we thus believe that policymakers can quickly and robustly get a quantitative sense of the strength and weaknesses of different policy routes.

Clearly, not only government agencies and policymakers more broadly but also the general public deserve a rich apparatus with the use of which they can compare outcomes both on the individual level and on the level of designing useful epidemiological policy for our populations. The access to off-the-shelf frameworks thereby also facilitates communication and helps build trust in policymakers and their recommendations, as their decision making evolves along with the evolution of the pandemic.

[25] Spatial features could obviously be introduced here too.

[26] This can be seen as an extension of the assumption in Brotherhood et al. (2021).

[27] Relevant data, in the case of covid-19, included mobility patterns from cell-phone data (Goolsbee and Syverson, 2021), consumption-expenditure patterns from credit-card data (Andersen et al., 2022, Bounie et al., 2020, Cox et al., 2020), and quasi-experimental evidence on the economic effects of vaccines (Hansen and Mano, 2023, Barro, 2022).

References

Acemoglu, Daron, Victor Chernozhukov, Iván Werning, and Michael D. Whinston (2021), "Optimal targeted lockdowns in a multigroup sir model." *American Economic Review: Insights*, 3 (4), 487–502. [6]

Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asuman E Ozdaglar (2024), "Testing, voluntary social distancing and the spread of an infection." *Operations Research*, 72 (2), 533–548. [6]

Adda, Jérôme (2007), "Behavior towards health risks: An empirical study using the "Mad Cow" crisis as an experiment." *Journal of Risk and Uncertainty*, 35 (3), 285–305. [6]

Aguiar, Mark and Erik Hurst (2007), "Measuring trends in leisure: The allocation of time over five decades." *Quarterly Journal of Economics*, 122 (3), 969–1006. [19]

Alvarez, Fernando E, David Argente, and Francesco Lippi (2021), "A simple planning problem for covid-19 lock-down, testing, and tracing." *American Economic Review: Insights*, 3 (3), 367–82. [6]

Andersen, Asger Lau, Emil Toft Hansen, Niels Johannesen, and Adam Sheridan (2022), "Consumer responses to the COVID-19 crisis: evidence from bank account transaction data." *The Scandinavian Journal of Economics*, 124 (4), 905–929. [38]

Arias, Elizabeth and Jiaquan Xu (2019), "United States life tables, 2017." *National Vital Statistics Reports*, 68 (7). [23]

Atkeson, Andrew (2020), "What will be the economic impact of COVID-19 in the US? Rough estimates of disease scenarios." Technical report, National Bureau of Economic Research. [6, 17]

Atkeson, Andrew (2021), "A parsimonious behavioral SEIR model of the 2020 covid epidemic in the united states and the united kingdom." Technical report, National Bureau of Economic Research. [33]

Aum, Sangmin, Sang Yoon Tim Lee, and Yongseok Shin (2021), "Inequality of fear and self-quarantine: Is there a trade-off between GDP and public health?" *Journal of Public Economics*, 194, 104354. [6, 36]

Azzimonti, Marina, Alessandra Fogli, Fabrizio Perri, and Mark Ponder (2020), "Pandemic control in ECON-EPI networks." Technical report, National Bureau of Economic Research. [38]

Barro, Robert J. (2022), "Vaccination rates and COVID outcomes across U.S. states." *Economics & Human Biology*, 47, 101201, URL https://www.sciencedirect.com/science/article/pii/S1570677X22000971. [38]

Becker, Gary S (1965), "A theory of the allocation of time." *Economic Journal*, 75 (299), 493–517. [3]

Becker, Gary S (1974), "A theory of social interactions." *Journal of Political Economy*, 82 (6), 1063–1093. [3]

Bethune, Zachary A and Anton Korinek (2020), "Covid-19 infection externalities: Trading off lives vs. livelihoods." Technical report, National Bureau of Economic Research. [6]

Bodenstein, Martin, Giancarlo Corsetti, and Luca Guerrieri (2022), "Social distancing and supply disruptions in a pandemic." *Quantitative Economics*, 13 (2), 681–721. [6]

Bognanni, Mark, Doug Hanley, Daniel Kolliner, and Kurt Mitman (2020), "Economic activity and COVID-19 transmission: Evidence from an estimated economic-epidemiological model." [6, 36]

Boppart, Timo, Karl Harmenberg, John Hassler, Per Krusell, and Jonna Olsson (2024), "Supplement to 'Integrated epi-econ assessment: quantitative theory'." *Quantitative Economics Supplemental Material*. [7]

Boppart, Timo, Karl Harmenberg, Per Krusell, and Jonna Olsson (2022), "Integrated epi-econ assessment of vaccination." *Journal of Economic Dynamics and Control*, 140, 104308. [6, 38]

Boppart, Timo and Per Krusell (2020), "Labor supply in the past, present, and future: a balanced-growth perspective." *Journal of Political Economy*, 128 (1), 118–157. [21]

Bounie, David, Youssouf Camara, Etienne Fize, John Galbraith, Camille Landais, Chloe Lavest, Tatiana Pazem, Baptiste Savatier, et al. (2020), "Consumption dynamics in the covid crisis: real time insights from French transaction bank data." *Covid Economics: Vetted and Real-Time Papers*, 59, 1–39. [38]

Brotherhood, Luiz, Philipp Kircher, Cezar Santos, and Michèle Tertilt (2021), "An economic model of the covid-19 epidemic: The importance of testing and age-specific policies." [6, 38]

Chan, Tat Y, Barton H Hamilton, and Nicholas W Papageorge (2016), "Health, risky behaviour and the value of medical innovation for infectious disease." *Review of Economic Studies*, 83 (4), 1465–1510. [6]

Chang, Roberto and Andrés Velasco (2020), "Economic policy incentives to preserve lives and livelihoods." Technical report, National Bureau of Economic Research. [6]

Conley, Bryan C. (1976), "The value of human life in the demand for safety." *American Economic Review*, 66 (1), 45–55. [21]

Cox, Natalie, Peter Ganong, Pascal Noel, Joseph Vavra, Arlene Wong, Diana Farrell, Fiona Greig, and Erica Deadman (2020), "Initial impacts of the pandemic on consumer behavior: Evidence from linked income, spending, and savings data." *Brookings Papers on Economic Activity*, 2020 (2), 35–82. [38]

Eichenbaum, Martin S, Miguel Godinho de Matos, Francisco Lima, Sergio Rebelo, and Mathias Trabandt (2024), "Expectations, infections, and economic activity." *Journal of Political Economy*, 132 (8), 2571–2611. [6]

Eichenbaum, Martin S, Sergio Rebelo, and Mathias Trabandt (2021), "The macroeconomics of epidemics." *The Review of Financial Studies*, 34 (11), 5149–5187. [6, 17]

Eichenbaum, Martin S., Sergio Rebelo, and Mathias Trabandt (2022a), "Epidemics in the new Keynesian model." *Journal of Economic Dynamics and Control*, 140, 104334. [37]

Eichenbaum, Martin S, Sergio Rebelo, and Mathias Trabandt (2022b), "The macroeconomics of testing and quarantining." *Journal of Economic Dynamics and Control*, 138, 104337. [6]

Farboodi, Maryam, Gregor Jarosch, and Robert Shimer (2021), "Internal and external effects of social distancing in a pandemic." *Journal of Economic Theory*, 196, 105293. [6, 36]

Garibaldi, Pietro, Espen R. Moen, and Christopher A. Pissarides (2020), "Modelling contacts and transitions in the sir epidemics model." *Covid Economics*, 5, 1–20. [6, 33]

Geoffard, Pierre-Yves and Tomas Philipson (1996), "Rational epidemics and their public control." *International Economic Review*, 603–624. [6]

Giagheddu, Marta and Andrea Papetti (2023), "The macroeconomics of age-varying epidemics." *European Economic Review*, 151, 104346. [6]

Giannitsarou, Chryssi, Stephen Kissler, and Flavio Toxvaerd (2021), "Waning immunity and the second wave: Some projections for SARS-CoV-2." *American Economic Review: Insights*, 3 (3), 321–338. [6, 33]

Glover, Andrew, Jonathan Heathcote, and Dirk Krueger (2022), "Optimal age-based vaccination and economic mitigation policies for the second phase of the covid-19 pandemic." *Journal of Economic Dynamics and Control*, 140, 104306. [5]

Glover, Andrew, Jonathan Heathcote, Dirk Krueger, and José-Víctor Ríos-Rull (2023), "Health versus wealth: On the distributional effects of controlling a pandemic." *Journal of Monetary Economics*, 140, 34–59. [6, 18, 22, 31, 35]

Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski (2014), "Optimal taxes on fossil fuel in general equilibrium." *Econometrica*, 82 (1), 41–88. [4]

Goolsbee, Austan and Chad Syverson (2021), "Fear, lockdown, and diversion: Comparing drivers of pandemic economic decline 2020." *Journal of Public Economics*, 193, 104311. [38]

Greenwood, Jeremy, Philipp Kircher, Cezar Santos, and Michele Tertilt (2019), "An equilibrium model of the african HIV/AIDS epidemic." *Econometrica*, 87 (4), 1081–1113. [6]

Hall, Robert E, Charles I Jones, and Peter J Klenow (2020), "Trading off consumption and covid-19 deaths." Technical report, National Bureau of Economic Research. [5, 22]

Hansen, Niels-Jakob H and Rui C Mano (2023), "Covid-19 vaccines: A shot in arm for the economy." *IMF Economic Review*, 71, 148–169. [38]

Jones, Callum J, Thomas Philippon, and Venky Venkateswaran (2021), "Optimal mitigation policies in a pandemic: Social distancing and working from home." *Review of Financial Studies*, 34 (11), 5188–5223. [6]

Kapička, Marek and Peter Rupert (2022), "Labor markets during pandemics." *Journal of Economic Theory*, 204, 105520. [6]

Kaplan, Greg, Benjamin Moll, and Gianluca Violante (2020), "The great lockdown and the big stimulus: Tracing the pandemic possibility frontier for the US." Technical report, National Bureau of Economic Research. [6]

Kermack, William Ogilvy and Anderson G McKendrick (1927), "A contribution to the mathematical theory of epidemics." *Proceedings of the Royal Society of London. Series A, Containing papers of a mathematical and physical character*, 115 (772), 700–721. [6]

Kremer, Michael (1996), "Integrating behavioral choice into epidemiological models of AIDS." *The Quarterly Journal of Economics*, 111 (2), 549–573. [6]

Krueger, Dirk, Harald Uhlig, and Taojun Xie (2022), "Macroeconomic dynamics and reallocation in an epidemic: evaluating the Swedish solution'." *Economic Policy*, 37 (110), 341–398. [6]

Lucas, Robert E. (1988), "On the mechanics of economic development." *Journal of Monetary Economics*, 22 (1), 3–42. [2]

Nordhaus, William D and Joseph Boyer (2000), *Warming the world: economic models of global warming*. MIT press. [4]

Petersen, Eskild, Marion Koopmans, Unyeong Go, Davidson H Hamer, Nicola Petrosillo, Francesco Castelli, Merete Storgaard, Sulien Al Khalili, and Lone Simonsen (2020), "Comparing SARS-CoV-2 with SARS-CoV and influenza pandemics." *The Lancet Infectious Diseases*. [36]

Piguillem, Facundo and Liyan Shi (2022), "Optimal COVID-19 quarantine and testing policies." *Economic Journal*, 132 (647), 2534–2562. [6]

Shepard, Donald S. and Richard Zeckhauser (1984), "Survival versus consumption." *Management Science*, 30 (4), 423–439. [21]

U.S. Bureau of Labor Statistics (2020), "American Time Use Survey, retrieved september 15, 2020, from https://www.bls.gov/tus/data.htm." [19]

van Vlokhoven, Has (2020), "Policy during an epidemic with super-spreaders." *Available at SSRN 3600441*. [6]

Viscusi, W. Kip and Joseph E. Aldy (2003), "The value of a statistical life: A critical review of market estimates throughout the world." *Journal of Risk and Uncertainty*, 27, 5–76. [5, 22]