NBER WORKING PAPER SERIES

INTEGRATED EPI-ECON ASSESSMENT

Timo Boppart
Karl Harmenberg
John Hassler
Per Krusell
Jonna Olsson

Working Paper 28282
http://www.nber.org/papers/w28282

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2020

For feedback we thank seminar participants at the Institute for International Studies, the University of Edinburgh, and the Institute for Future Studies. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Integrated Epi-Econ Assessment
Timo Boppart, Karl Harmenberg, John Hassler, Per Krusell, and Jonna Olsson
NBER Working Paper No. 28282
December 2020
JEL No. E10,I10

ABSTRACT

We formulate an economic time use model and add to it an epidemiological SIR block. In the event of an epidemic, households shift their leisure time from activities with a high degree of social interaction to activities with less, and also choose to work more from home. Our model highlights the different actions taken by young individuals, who are less severely affected by the disease, and by old individuals, who are more vulnerable. We calibrate our model to time use data from ATUS, employment data, epidemiological data, and estimates of the value of a statistical life. There are qualitative as well as quantitative differences between the competitive equilibrium and social planner allocation and, moreover, these depend critically on when a cure arrives. Due to the role played by social activities in people's welfare, simple indicators such as deaths and GDP are insuffcient for judging outcomes in our economy.

Timo Boppart
Institute for International Economic Studies
Stockholm University
SE-106 91 Stockholm
Sweden
timo.boppart@iies.su.se

Karl Harmenberg
Department of Economics
Copenhagen Business School
Denmark
kha.eco@cbs.dk

John Hassler
Institute for International Economic Studies
Stockholm University
Sweden
John.Hassler@iies.su.se

Per Krusell
Institute for International Economic Studies
Stockholm University
106 91 STOCKHOLM
SWEDEN
and NBER
per.krusell@iies.su.se

Jonna Olsson
Department of Economics
University of Edinburgh
United Kingdom
jonna.olsson@ed.ac.uk

1 Introduction

We divide our introduction up into three pieces: a background, where we also describe our modeling approach and contribution, a results section, and a literature discussion.

1.1 Motivation and model

The current covid-19 pandemic has necessitated difficult decisions for politicians and government officials across the world. How should the public-health effects of the epidemic be mitigated and balanced against economic values such as economic output? To approach this difficult question, coherent quantitative theoretical frameworks are helpful if not necessary. Such frameworks take time to develop and test, but the endeavor is of obvious importance even if a fully satisfactory framework will not be available before the current epidemic has subsided: it will help us understand what happened during covid-19 and what alternative policy decisions could have led to, and it will prepare us for future epidemics.

This paper has the broader goal of producing an integrated assessment model of epidemics and economics. In important ways, our goal 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: what we argue for here is just that, merely in a different application.[1] And to be sure, we are far from the first paper with this aim; we discuss the literature below in more detail—including the pre-covid contributions—but in terms of the rapidly expanding very recent literature, the main paper we want to build further on is Eichenbaum et al. (2020a). Our "epi-econ" integrated assessment model of an economy during an epidemic thus develops new model features that we believe are necessary in an assessment of the trade-offs involved in the decisions that politicians and government officials have to make during an epidemic. We propose a concrete framework and evaluate it quantitatively for the covid-19 case, i.e., we solve the model for parameter values that are selected to match key economic and epidemiological facts. We then evaluate policy using this quantitative model.

A key feature of our epi-econ integrated assessment model is an element of sociology, i.e., an aspect of economic activity that describes how humans interact at work and while enjoying leisure. Thus, a more appropriate attribute than epi-econ might be epi-socio-econ, but since our proposed sociological elements are modeled with standard microeconomic tools—as opposed to through theory constructs used in sociology, such

[1] See, e.g., Nordhaus and Boyer (2000) and Golosov et al. (2014).

as social networks—we prefer the more modest label. There are two main reasons for more explicit sociological elements. One is the obvious one: human interaction is key in epidemiology. The second and no less important reason is that we perceive social interaction to be very important in any welfare evaluation in this area. An example is the changed nature of leisure that materializes during times of restricted social activities. Concrete, extreme expressions of this concern include reports of increases in the number of cases of domestic violence and sharply increased activity on suicide prevention help-lines.[2] But more generally we believe that the costs of restricting behavior—that may be necessary during an epidemic—need to go beyond merely counting lives and GDP: we need to also factor in leisure and its nature.

Our model builds a macroeconomic framework determining the allocation of production, consumption, and leisure based on market forces. An important task for us is to consider its efficiency properties, to be discussed extensively below. The framework features population heterogeneity both with respect to epidemic risk and economic productivity. Since many infections disproportionately hit specific subgroups of the population, and since subgroups of the population vary in productivity, this heterogeneity matters for the quantitative assessment of the effects of an epidemic. Our framework models social activities explicitly. From the epidemic side, an epidemic spreads when people meet and socialize, either in the workplace or in their spare time, not when they consume goods in general. Moreover, we measure the intensity of interaction by the number of hours spent in different activities. Along the same lines, leisure can, but does not necessarily, have a social-activity component and we make this concrete but modeling the number of hours spent in leisure activities with different degrees of social content.

The sociological elements we entertain in our setting obey standard, microeconomic principles. We distinguish goods by their degree of social interaction; for example, enjoying leisure at a live football game involves more social interaction than watching TV. Thus, we give goods a social-activity attribute. However, our formulation builds on the assumption that people do not care about this social activity per se: goods and leisure activities are merely enjoyed through the amounts by which they are consumed (e.g., the number of goods/services of different types and the number of hours of different kinds of leisure). In concrete terms: although there are more people around you in a football game, and you may in reality enjoy the game somewhat more in a packed stadium, we assume that consumers simply derive utility from how much time

[2]This year (as of mid-December), the number of calls to the Swedish help-line for suicidal prevention has increased by 35 percent compared to previous year. Isolation and loneliness due to corona measures are said to be the main culprits. See https://sverigesradio.se/artikel/7623261.

they spend on going to football games. Our microeconomic formulation of different degrees of social activity thus does not model what from another perspective would seem like "externalities": my going to a football game will increase the utility of other people in the stadium (and in other contexts, of course, even the direction of this externality is unclear). Our framework is simple—and, in this sense, it is a proof of concept—in that we only consider two levels of social activity (zero and positive). However, it is readily generalizable in this dimension. A convenient feature of our setup is that, abstracting from epidemiological concerns, the market equilibrium is efficient. Moreover, in the specific framework we employ, we use the construct of a "representative family" (consisting of a large number members of different age as well as health statuses). This means that welfare comparisons across different policies vis-à-vis covid-19 can be made straightforwardly.

An integrated assessment model of the economy during an epidemic needs to be quantitative in nature. Our calibration uses the American Time Use Survey to match the time spent on social leisure, non-social leisure and work (in the workplace or at home), by young and old individuals. We argue that data on time allocations at this level of disaggregation are necessary for studying the interaction between individual behavior and epidemic spread.

Furthermore, a quantitative model necessitates numerical solution methods if we want to go beyond qualitative insights. Though the model we formulate is quite non-linear, not just in its SIR dynamics, and many decisions are made and prices are set within each model period, we offer a highly computationally tractable numerical routine for computing both the social planner's allocation and the market allocation. It builds on two blocks. One is a standard computable general equilibrium (CGE) block for every time period (a day in our calibration), with two added features. First, it takes as given a population structure (overall population size and its health/age composition) that of course in our dynamic model is generated by the past. Second, it takes as given a vector of end-of-period "multipliers" that are forward-looking measures of the utility value of each of the different types of agents: how much utility is generated by the population structure that next period will inherit from today. Then, in our dynamic model, periods are connected by envelope conditions and SIR dynamics and computation can be performed in a straightforward algorithm that iterates over the path for the multiplier vector.

1.2 Key findings

Households reallocate their time substantially as a response to the epidemic–*if* they are fully informed. In the case of covid-19, the death toll in a myopic scenario with no cure is reduced by 80 percent by assuming rational expectations and endogenous adjustments.

Still, the comparison of the rational-expectations scenario and the social planner's optimal scenario shows that there is substantial scope for government policy to improve reallocations and thus welfare beyond simply informing. A social planner would want to reduce the number of lost years of life even further, and, importantly, distribute the burden of behavioral adjustment more efficiently in the population. This would come at a cost: output would fall substantially more in the social planner's optimal scenario. However, the cost in terms of per-capita flow utility is smaller: during the worst month of the epidemic, the loss in flow utility is actually larger in the rational-expectations scenario than in the social planner's optimal scenario. 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 social planner is willing to sacrifice consumption utility not only in order to save lives but also save leisure utility for the old. In the rational-expectation scenario, the old and thus vulnerable individuals are voluntarily staying home to protect their lives. However, this minimization of social interactions comes at a large utility cost. The old alone bear the burden of reducing lives lost in the rational-expectations scenario, while the burden is more equally distributed by the social planner.

In the baseline scenario when no cure is expected to arrive, 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 over-burdened, and so that "overshooting" in terms of number of infected before herd immunity is reached is minimized. However, the optimal approach can be qualitatively different depending on two crucial assumptions.

First, the optimal strategy depends qualitatively on when a cure arrives. The sooner a cure to the disease is expected to arrive, the more likely it is that a social planner chooses to adopt a strategy best described as "suppress": to keep the number of infections low by lowering social activity (and thus output) to a very large extent. The longer it takes until the cure arrives, the more likely it is that the optimal scenario tilts over to the "protect the health care system".

Second, the optimal strategy depends crucially on what we assume about the value of a statistical life. We exemplify this in a scenario that assumes that a cure arrives after 18 months. If the social planner uses a value of a statistical life on the higher end of estimates in the literature, the optimal strategy is to suppress the epidemic while waiting for the cure. However, if the value of a statistical life is on the lower end, the optimal strategy tilts over to the "protect the health care system" strategy. Thus, the value of a statistical life and the expected arrival of the cure are two crucial assumptions that impact the optimal strategy not only quantitatively but also qualitatively.

To validate the model, we simulate two other diseases: a seasonal flu and SARS. This exercise helps validating our assumptions, especially the critical assumption of the intrinsic value of a statistical life. There are numerous estimates in the literature, and we do not (and should not!) take a stance on the correct value. However, we show that a value of a statistical life from the higher end of the spectrum implies that a social planner would want to slow the economy substantially to stop the spread of a seasonal flu, by lowering output by 4.6 percent during a critical quarter, and by 3.4 percent over a year. This, as far as we can tell, is not how policy makers react. Thus, although we cannot say whether a chosen values of a statistical life is the correct one, we can say that a value from the lower range is more in line with observed policy actions. When we test our model by simulating SARS, a disease both more contagious and more lethal than covid-19, we get that a social planner would impose a very strict lock-down even with the lower value of a statistical life (assuming a year until we can end the epidemic exogenously).

Finally, an important part of our motivation of this work is that we perceive the focus on deaths and GDP alone to be too narrow. What does our model then tell us in this respect? The difference between just a crude consumption measure and a broader measure of utility, including socially active leisure, is clear from the experiments. In the main scenario, a social planner reduces output by 9.7 percent the first year of the epidemic while the utility loss incurred during the first year corresponds to a consumption fall of 9.9 percent.[3] Thus, in the social planner scenario, these two measures are reasonably close. In the rational-expectation scenario, output falls by a mere 1.5 percent the first year. However, utility falls by an amount that is equivalent to a fall in consumption of 8.1 percent, i.e., more than 5 times higher! The reason is the severe utility loss the old experience during the epidemic in the rational-expectation scenario, in which it is simply too dangerous for them to go out and enjoy social leisure.

Note that this is a utility measure that does not take into account the future utility loss due to lost lives, but only the utility in the economy during the first year of the

[3] Calculated as the change in consumption level for all types of goods that, given constant time allocations, would give the same welfare loss during the first year of the epidemic.

epidemic. The social planner saves years of life and therefore the total utility in the social planner's scenario is (of course) higher.

1.3 Related literature

The task of reviewing the exploding literature on economics and epidemiology is daunting, especially since many papers are still preliminary (and yet relevant for us). As a consequence, we will confine ourselves here to making a few comments on some ways in which we believe that our approach overlaps with that taken elsewhere.

Atkeson (2020) describes the core model developed by epidemiologists that we use here (Kermack and McKendrick, 1927). Epidemiological models had been used in economics prior to covid-19 in analyses of other viruses (e.g., Geoffard and Philipson (1996), Kremer (1996), Adda (2007), Chan, Hamilton and Papageorge (2016), and Greenwood, Kircher, Santos and Tertilt (2019)), but Eichenbaum, Rebelo, and Trabandt (2020; ERT) was 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 build 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, in our view, hallmarks of an appropriate way forward in this literature, and we thus follow ERT in many ways. Our main addition, in terms of our modeling and quantitative approach, is to move toward explicitly describing time use. 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. Here, our work has overlaps with Brotherhood et al. (2020), Glover et al. (2020), and Kaplan et al. (2020).

We also emphasize heterogeneity: people's vulnerability toward covid-19 as well as their economic detail (such as productivity) differ in our model, and there are different sectors (differing in the degree of social activity). At least one of these features are captured not only in the papers just mentioned but also in Acemoglu et al. (2020a), Acemoglu et al. (2020b), Bodenstein et al. (2020), Giagheddu and Papetti (2020), Krueger et al. (2020), van Vlokhoven (2020), Kapicka and Rupert (2020), and Aum et al. (2020).[4]

We conduct a systematic comparison between a laissez-faire market allocation and a social optimum: many other papers have the same aim.[5] Perhaps the closest to ours is Farboodi et al. (2020); many of their qualitative results also resemble ours. Our analysis of externalities in social activities is closely related to that in Garibaldi et al. (2020).

To us, surprisingly few papers include an explicit account of the value of a statistical life; even fewer discuss it in detail and provide robustness analysis with respect to it.[6] Relatedly, as far as we could tell, none of the studies relate policy prescriptions for covid-19 to those that would be implied for other known (and experienced) viruses. Finally, in our setting we take a representative-family perspective, thus not considering compensatory transfers across groups in the population. A few of the papers in the literature have this focus, e.g., Glover et al. (2020) and Kaplan et al. (2020).

Clearly, our paper misses some important elements. One is the connection to data that is now available on social and economic activity during the covid crisis. Here, Aum et al. (2020), Bognanni et al. (2020), Farboodi et al. (2020), Giagheddu and Papetti (2020), and Krueger et al. (2020) stand out. Second, we do not review new and relevant contributions to epidemiological modeling here (some of which have been made by economists). Third, we do not allow analysis of individual awareness of health status; clearly, such studies are key for allowing us to analyze test-trace-and-isolate strategies, and we aim to develop our current framework further in this direction. Existing papers with such elements include Acemoglu et al. (2020a), Acemoglu et al. (2020b), Alvarez et al. (2020), Aum et al. (2020), Bethune and Korinek (2020), Brotherhood et al. (2020), Eichenbaum et al. (2020b), Farboodi et al. (2020), Kapicka and Rupert (2020), Krueger et al. (2020), and Piguillem and Shi (2020).

[4] van Vlokhoven (2020) allows heterogeneity in the propensity to pass the virus on to others.

[5] These include Acemoglu et al. (2020a), Acemoglu et al. (2020b), Alon et al. (2020), Alvarez et al. (2020), Aum et al. (2020), Bethune and Korinek (2020), Bodenstein et al. (2020), Brotherhood et al. (2020), Chang and Velasco (2020), Eichenbaum et al. (2020b), Farboodi et al. (2020), Garibaldi et al. (2020), Giagheddu and Papetti (2020), Giannitsarou et al. (2020), Glover et al. (2020), Kapicka and Rupert (2020), Kaplan et al. (2020), Krueger et al. (2020), Jones et al. (2020), and van Vlokhoven (2020). Some papers perform maximization based formulations of social welfare that are not explicitly microeconomic rooted and although they contain interesting insights, we do not list them here.

[6] The sole focus in Hall et al. (2020) is on this issue; other nice contributions include Alvarez et al. (2020), Alon et al. (2020), Glover et al. (2020), and Krueger et al. (2020)

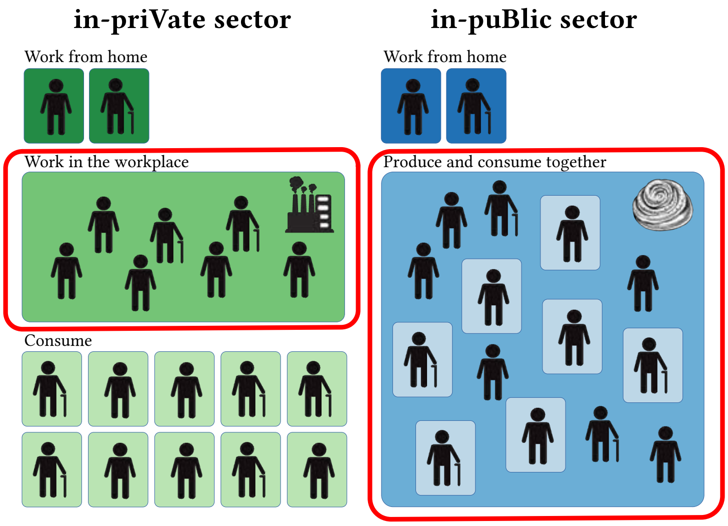

The areas marked in red are where the virus spreads.

Figure 1: Illustration of the model.

1.4 Roadmap for the rest of the paper

We formalize our epi-econ integrated assessment model in steps. In Section 2 we describe the economic model without any epidemic and its calibration to pre-pandemic time-use data and production data. Thereafter, in Section 3, we add an epidemic dimension to a static version of the model. Then, in Section 4, we present the full dynamic version of the model. For all steps, we describe the social planner's problem, and in the end comment on the differences between the social planner's solution and a decentralized market solution. Section 5 contains the results. Finally, Section 6 concludes.

2 The pre-pandemic economic model

Figure 1 schematically illustrates the model we have in mind. There are two types of individuals, young and old. They can spend their leisure in a sector we label as the in-priVate sector (illustrated by the figures in the bottom left in the figure). When people spend their time on this type of leisure they are on their own (e.g., they are at home watching Netflix). In other words: they do not interact with others and hence 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 what we label as the in-puBlic sector. In this sector, consumption and work take place jointly, and the virus can be spread between those enjoying their 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.

In the beginning of the epidemic, a small fraction of the population is infected, while the majority are still susceptible, and no-one has yet recovered. As the epidemic spreads, susceptible individuals become infected, while some infected individuals recover and a fraction of the infected dies. The infection fatality rate is higher for the old than for the young, and it increases if the hospitals become overcrowded. We assume full immunity, so that once an individual has recovered, he/she cannot be re-infected.

Individuals in the model receive utility from leisure, from consumption, and from the intrinsic value of being alive.

We will describe the model outlined above stepwise. As a first step, we in this Section describe the economy before the epidemic, and discuss how we calibrate the model to pre-pandemic data.

2.1 Formal description of the pre-pandemic economic model

The formal description of the economic model without any epidemiological considerations is straight-forward. The economy consists of a continuum of identical families. Each family consists of a continuum of young individuals, with mass ϕ^y, and a continuum of old individuals, with mass ϕ^o. Individual utility is a function of consumption (c) and leisure (h). Leisure is of two sorts: leisure involving social interactions (e.g., restaurants or movies) and leisure not involving social interactions (e.g., watching TV, reading the newspaper); we use the indexes B (in-puBlic) and V (in-priVate) for these, respectively.

The consumption good is also of two sorts: goods and services consumed while being socially active (e.g., restaurants or movies), c_B, and goods and services consumed in private (e.g., television), c_V.

The family utility function is $\sum_{i \in \{y,o\}} \phi^i u(c_B^i, h_B^i, c_V^i, h_V^i)$ and the social planner maximizes the utility of the families,

$$\max \sum_{i\in\{y,o\}} \phi^i u(c_B^i, h_B^i, c_V^i, h_V^i)$$

subject to

$$\begin{aligned} 1 &= h_B^i + h_V^i + n_{Bh}^i + n_{Bw}^i + n_{Vh}^i + n_{Vw}^i, & i \in \{y, o\}, && (1) \\ \sum_{i\in\{y,o\}} \phi^i c_j^i &= F_j(\phi^y n_{jh}^y, \phi^y n_{jw}^y, \phi^o n_{jh}^o, \phi^o n_{jw}^o), & j \in \{B, V\}, && (2) \end{aligned}$$

and non-negativity constraints for both time and consumption quantities. Constraints (1) are the time constraints. Both young and old have one unit of time at their disposal, and can spend time on leisure (h_B and h_V) or working (n).

Hours worked can take place in two different locations: either at home (indexed h) or in the workplace (indexed w) and each of these types of work can be carried out in two sectors: producing goods and services for the in-puBlic sector (e.g., restaurants or movies, indexed by B), or producing goods and services for the in-priVate sector (everything else, indexed by V).

The constraints (2) are the resource constraints. Production in both sector B and in sector V is a function of the labor worked from home and the labor in the workplace, for both young and old.

We now turn to the calibration of this model.

2.2 Parametrization of the economic model

Utility function: The flow utility for an individual is given by

$$\begin{aligned} u(c_B, h_B, c_V, h_V) &= \log \mathrm{CES}(\tilde{c}_B, \tilde{c}_V; \lambda, \varepsilon), && (3) \\ \tilde{c}_B &= \mathrm{CES}(c_B, h_B; \lambda_B, \varepsilon_B), && (4) \\ \tilde{c}_V &= \mathrm{CES}(c_V, h_V; \lambda_V, \varepsilon_V), && (5) \end{aligned}$$

where the constant-elasticity aggregator CES($\bullet$) is defined by

$$\mathrm{CES}(x_1, x_2; \lambda, \varepsilon) = \left(\lambda x_1^{\frac{\varepsilon-1}{\varepsilon}} + (1-\lambda) x_2^{\frac{\varepsilon-1}{\varepsilon}}\right)^{\frac{\varepsilon}{\varepsilon-1}}. \qquad (6)$$

The nested CES structure captures that the consumer needs to spend leisure time to derive utility from a good or a service irrespective of whether the activity is social or non-social.

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

$$
\begin{aligned}
F_j(n_{jh}^y, n_{jw}^y, n_{jh}^o, n_{jw}^o) &= k_j^{\alpha} \tilde{n}_{jh}^{\nu} \tilde{n}_{jw}^{1-\alpha-\nu}, &(7)\\
\tilde{n}_{jh} &= \text{CES}(\phi^y n_{jh}^y, \phi^o n_{jh}^o; \varphi, \theta), &(8)\\
\tilde{n}_{jw} &= \text{CES}(\phi^y n_{jw}^y, \phi^o n_{jw}^o; \varphi, \theta), &(9)
\end{aligned}
$$

where the CES function is given by (6). Given our focus on short-run analysis, the sector-specific capital stocks k_B and k_V are fixed.

2.3 Calibration of the economic model

We define young as individuals aged 15-60, while old refers to individuals above 60. We assume a unit mass of people, thus, calibrating to US demography, we set ϕ^y, the young population share, to 0.73 and $\phi^o = 1 - \phi^y$.

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 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.[7] With these definitions, we compute the

[7]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

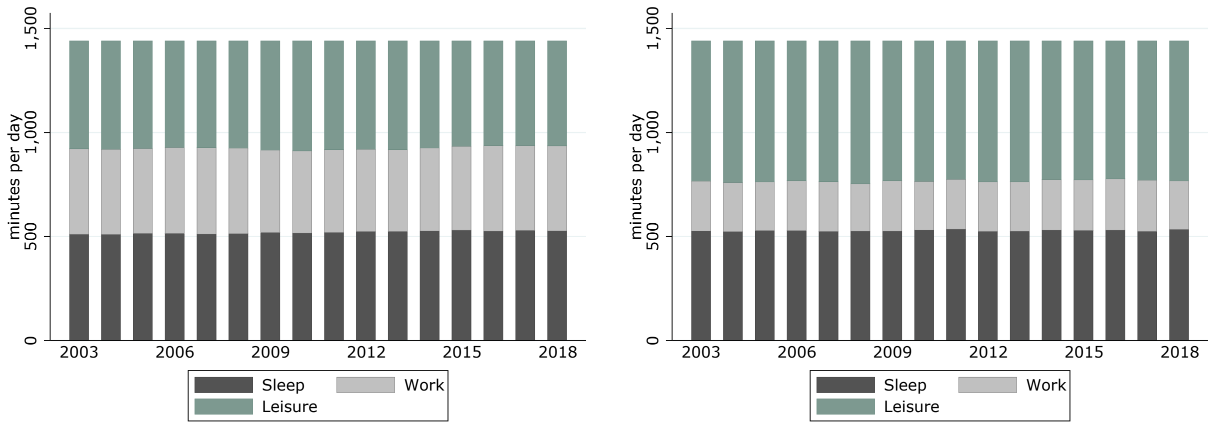

(a) Young (b) Old

A full day is $24 \cdot 60 = 1,440$ minutes.
Source: ATUS.

Figure 2: Time spent on sleep, work, and leisure over time.

share of sleep, work, and leisure for the young and the old respectively, with the result shown in Figure 2. 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 work place (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. Figure 3 shows the time spent in socially intense leisure vs. not socially intense leisure for young and for 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.[8]

Our definition of work includes market work, household work, core housework and

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.".

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

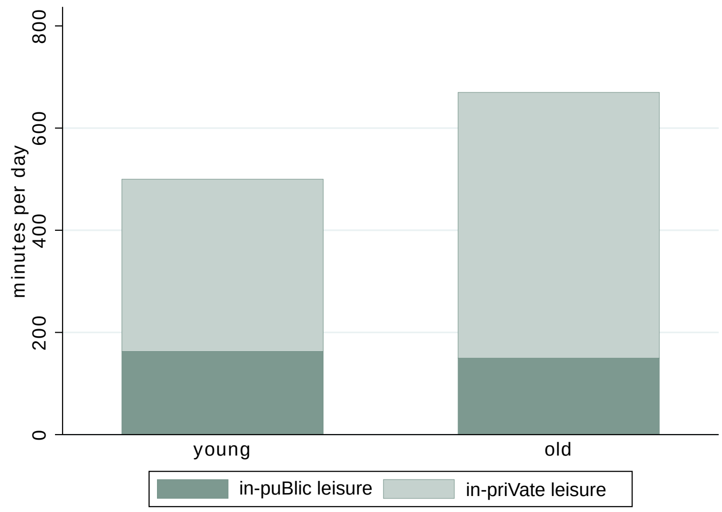

Socially intense activities include associated traveling.
Source: ATUS 2018.

Figure 3: Average minutes per day spent in socially intense leisure activities and not socially intense activities.

home production, necessity shopping, and time spent in education. In the same way as for leisure, we classify all work activities according to where they were performed: in the home or outside home. Figure 4 shows how many minutes of the average working day is spent at home (112 minutes on average) and outside home, mainly at the workplace (255 minutes on average). As Figure 4 also shows, there is no clear time trend in how large fraction of the working time that is spent at home. The slight downward trend in work done outside home can mainly be attributed to a compositional effect: the fraction of old individuals has increased slightly during this time period (21% in 2003 compared to 27% in 2018), and they work less, especially outside the home.

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 with producing for the socially active bundle. The lion's share (43%) of the workforce working in the socially intense

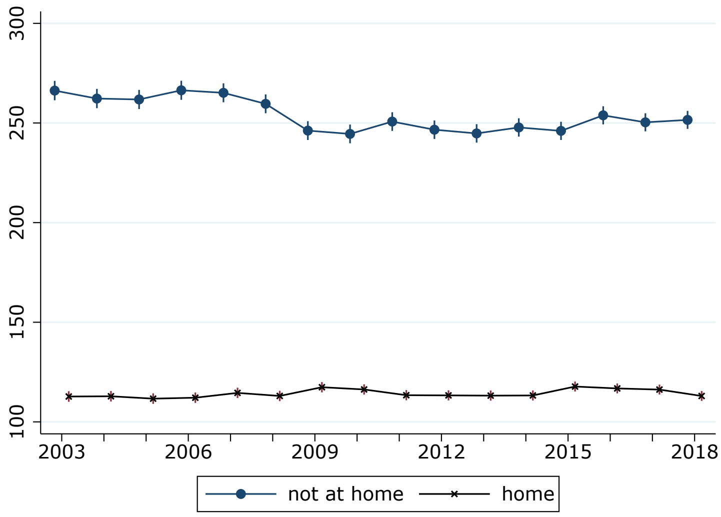

Source: ATUS.

Figure 4: Average minutes per day spent working home vs. not from home.

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 Appendix B.

Calibration of the production functions: The parameters of the production functions are externally calibrated. We assume $\alpha = 1/3$ for standard reasons and $\theta = 10$, reflecting a high but finite substitutability between young and old workers. In the data, working hours are on average distributed between work from home (n_h) vs. in the workplace (n_w) such that $n_w/n_h = 2.3$ (as seen in Figure 4). 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 $\nu \approx 0.202$. The output elasticity of work that can be done from home is 0.202 whereas it is 0.465 for work that can only be done at the workplace.

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. Then output 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 a huge amount. If n_w/n_h falls to 1/3 (2 hours worked at the workplace plus 6 hours worked from home out of a 8 hours workday) the output loss is 25 percent. We find these losses reasonable

Target	*Parameter*
Marginal product of capital equal across sectors	$\rbrace k_B, k_V$
Marginal product of capital net of depreciation = interest rate	
Output share of public sector = 0.2	λ
Leisure young = 0.56	$\rbrace \lambda_B, \lambda_V, \varphi$
Leisure old = 0.74	
in-puBlic leisure/in-priVate leisure for young = 0.45	

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

Table 1: Calibration targets for the internally calibrated economic parameters.

in magnitude, supporting a choice of $\nu = 0.202$.

Remaining internally calibrated parameters: Aside from the preference elasticities $\varepsilon, \varepsilon_B$, and ε_V, the rest of the economic parameters are jointly internally calibrated, with targets summarized by Table 1.

Based on employment figures in the two types of goods/services production, we target an aggregate production share of the B good of 0.2. This target is primarily determined by the parameter λ, the utility weight on the B good.

Choosing the capital stocks k_B and k_V is simply a normalization. Changing the units of, e.g., k_B necessitates adjusting the utility weights λ_B and λ but leaves the maximization problem of the planner intact. For ease of interpretability, we set k_B and k_V such that the implicit price of the B good and the V good are equal and the return on capital is equalized across sectors, which means that k_B is 0.2 of the total capital stock (i.e., $k_B/k_V = 1/4$).

The young spend 56% of their non-sleeping time on leisure. We therefore impose that $h_B^y + h_V^y = 0.56$. The old population spend more of their available time on leisure, 74%. We therefore impose that $h_B^o + h_V^o = 0.74$. We classify leisure time as spent either at home, which is thought of as leisure of the non-socially-intense V type, or spent outside home, which is then classified as leisure of the socially intense B type. The young spend 31% of total leisure time on the active B type of leisure, while 69% is spent on the non-active V type. This gives us an additional constraint: $h_B^y/h_V^y = 0.45$ (the corresponding ratio for the old is endogenous). These three time-use moments primarily determine the utility weights on leisure, λ_B and λ_V, as well as the relative productivity of the old, φ.

Given that the young and old are imperfect substitutes, φ does not directly correspond to the difference in marginal productivity. Our calibrated value of φ is 0.62, yielding a marginal productivity of the old which is 72 percent of the marginal productivity of the young. To compare this estimate directly to the data is difficult, since the low labor force participation, especially for the old, makes the selection bias severe: remember that our group of old contains everyone above the age of 60, and in the group of the oldest old, say above 80, we have extremely few wage observations, and it is easy to argue that the few observations we see is not a random selection of the population. A raw estimate of the average observed wages in the same age groups from the CPS (pooling data from $2003-2017$) gives an average wage for the old of 0.92 (normalizing the youngs' average wage to 1), which we argue is well in line with our estimate, given the extremely few wage observations in the ages above 70.

Elasticities in the utility function: Above we took as given the elasticities in the utility function. In this section we discuss how $\varepsilon, \varepsilon_B$ and ε_V are chosen. 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 $\varepsilon = 1.0$, yielding additive separability between the two consumption-leisure bundles.[9]

To pick ε_B and ε_V, we put additional restrictions on the utility function. First, we require that the income effect should dominate 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, 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 $\varepsilon_V = 0.80$.[10]

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

[9] We acknowledge that this is not an obvious conclusion, and therefore perform robustness checks with respect to the value of ε, found in Appendix D. 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 $\varepsilon = 1.0$ seem plausible.

[10] See Appendix C for more details.

Parameter	Description	Value
Preference parameters		
λ	Weight on $\tilde{c}_B$	0.25
λ_B	Weight on c_B	0.93
λ_V	Weight on c_V	0.69
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
Technology		
α	Capital share	1/3
ν	Home work labor share	0.202
θ	Elasticity of substitution between young and old	10
φ	Production weight on young	0.62
k_B/k_V	Relative capital stock	0.25
Demographics		
ϕ^y	Fraction young	0.73

See text for description of sources and methodology.

Table 2: Summary of economic parameters.

to not only the intensive margin elasticity but the aggregate elasticity including also the extensive margin, we think a value of 1.1 sounds reasonable.

We summarize the parameters used in the economic model in Table 2.

3 A static version of the model with an epidemic

Before we describe the full dynamic model, we unpack how the model works by considering a static version of the model with a reduced-form epidemic side. The static model provides important insight and will also be a key step in the solution of the dynamic model. We simply assume that the utility costs of additional infections are exogenously given and equal to ξ^{T^y} and ξ^{T^o} for young and old respectively. Moreover, we assume that the SIR state in the economy, i.e., the fraction of the young and the old population being susceptible, infected, and recovered, is given exogenously. The planner knows the SIR state, but cannot observe who is susceptible, infected, or recovered. Given this information, the social planner maximizes total utility net of the cost of infections

$$\sum_{i\in\{y,o\}}\left(\phi^i u(c_B^i, h_B^i, c_V^i, h_V^i) - \xi^{T^i} T^i\right)$$

subject to the economic constraints

$$1 = h_B^i + h_V^i + n_{Bh}^i + n_{Bw}^i + n_{Vh}^i + n_{Vw}^i \qquad i \in \{y, o\}, \tag{10}$$

$$\sum_{i\in\{y,o\}} \phi^i c_j^i = F_j(\phi^y n_{jh}^y, \phi^y n_{jw}^y, \phi^o n_{jh}^o, \phi^o n_{jw}^o), \qquad j \in \{B, V\}, \tag{11}$$

$$\phi^i = S^i + I^i + R^i \qquad i \in \{y, o\} \tag{12}$$

and the transmission equations,

$$T^i = \widehat{\pi}_B(h_B^i + n_{Bw}^i)S^i + \widehat{\pi}_V n_{Vw}^i S^i \qquad i \in \{y, o\}, \tag{13}$$

$$\widehat{\pi}_B = \pi_B \frac{\sum_{m\in\{y,o\}} I^m(h_B^m + n_{Bw}^m)}{\sum_{m\in\{y,o\}}\left[(S^m + I^m + R^m)(h_B^m + n_{Bw}^m)\right]}, \tag{14}$$

$$\widehat{\pi}_V = \pi_V \frac{\sum_{m\in\{y,o\}} I^m n_{Vw}^m}{\sum_{m\in\{y,o\}}\left[(S^m + I^m + R^m)n_{Vw}^m\right]} \tag{15}$$

as well as non-negativity constraints for both time and consumption quantities.

Constraints (10) and (11) are the same as in the static economic problem. Constraint (12) is an accounting identity, the population of type i is composed of susceptible, infected, and recovered individuals. Constraint (13) gives the number of transmissions T^i for type i, depending on the infection risks (per time unit) $\widehat{\pi}_B$ and $\widehat{\pi}_V$, the time spent socially in the B sector $(h_B^i + n_{Bw}^i)$ and in the V sector (n_{Vw}^i), as well as the number of susceptible individuals S^i of type i. Finally, constraints (14) and (15) describe how the infection risks depend on the parameters π_B and π_V as well as the share of infected individuals, weighted by time spent in the sector.

The two new parameters, π_B and π_V, capture how likely it is that the disease is transmitted, given the amount of time spent on the contagious activities. These parameters will be calibrated once we have a dynamic model with a full SIR block. For now, we set $\pi_B = \pi_V = 0.24$, which we later will show is a reasonable calibration for covid-19 in this model if we assume that a period in the model corresponds to one day.

3.1 Static externalities

We can characterize the effect of the epidemic on the (static) economy with four wedges. For expositional purposes, assume that type i works from home as well as in both the B work place and the V work place (i.e., the solution is interior). By taking the first-order conditions in the social planner's problem with respect to n_h^i, n_{Bw}^i, n_{Vw}^i, we get the following relationships between the marginal productivities at the work place w_{Bw}^i and w_{Vw}^i compared with the marginal productivities working from home w_h^i, as well as the marginal utility of puBlic and priVate leisure,

$$w_{Bw}^i = w_h^i + \tau_B^i, \tag{16}$$

$$w_{Vw}^i = w_h^i + \tau_V^i, \tag{17}$$

$$\frac{\partial u^i}{\partial h_B^i} = \frac{\partial u^i}{\partial h_V^i} + \tau_B^i, \tag{18}$$

where $w_{Bw}^i = \frac{\partial u}{\partial c_B}\frac{\partial F_B}{\partial n_{Bw}^i}$ is the marginal utility productivity of office work in the B sector for type i, $w_{Vw}^i = \frac{\partial u}{\partial c_V}\frac{\partial F_V}{\partial n_{Vw}^i}$ is the marginal utility productivity of office work in the V sector for type i, and $w_h^i = \frac{\partial u}{\partial c_B}\frac{\partial F_B}{\partial n_{Bh}^i} = \frac{\partial u}{\partial c_V}\frac{\partial F_V}{\partial n_{Vh}^i}$ is the marginal utility productivity of home work for type i.[11]

The wedges τ_B^i and τ_V^i (which differ by age) are the implied costs, per time unit, of being in the puBlic area (either as a worker or as a consumer) and the priVate-sector office, summarizing the entire epidemic dimension of the social planner's problem. Explicitly, the wedges are given by

$$\tau_B^i = \underbrace{\frac{S_t^i}{S_t^i + I_t^i + R_t^i}\widehat{\pi}^B \xi^{T^i}}_{\text{Infection-risk } B \text{ wedge}} \tag{19}$$

$$+ \underbrace{\pi^B \left(\frac{I_t^i}{S_t^i + I_t^i + R_t^i} \cdot \frac{1}{\sum_m (S_t^m + I_t^m + R_t^m)(h_B^m + n_{Bw}^m)} - \frac{\sum_m I_t^m (h_B^m + n_{Bw}^m)}{\left(\sum_m (S_t^m + I_t^m + R_t^m)(h_B^m + n_{Bw}^m)\right)^2} \right) \xi^{\pi^B}}_{\text{Static epidemiological externality } B \text{ wedge}},$$

[11]Since the social planner ensures that marginal utilities are equalized for young and old, e.g., $\frac{\partial u^y}{\partial c_V^y} = \frac{\partial u^o}{\partial c_V^o}$, we omit superscripts on the marginal utilities. Further, since the social planner ensures that the marginal utility productivity from home work is equalized across the two sectors, we omit the sector subscript for home work marginal utility productivity.

and

$$\tau_V^i = \underbrace{\frac{S_t^i}{S_t^i + I_t^i + R_t^i}\widehat{\pi}^V \xi^{T^i}}_{\text{Infection-risk } V \text{ wedge}} \tag{20}$$

$$+ \underbrace{\pi^V \left(\frac{I_t^i}{S_t^i + I_t^i + R_t^i} \cdot \frac{1}{\sum_m (S_t^m + I_t^m + R_t^m) n_{Vw}^m} - \frac{\sum_m I_t^m n_{Vw}^m}{\left(\sum_m (S_t^m + I_t^m + R_t^m) n_{Vw}^m\right)^2} \right) \xi^{\pi^V}}_{\text{Static epidemiological externality } V \text{ wedge}}$$

where ξ^{π^B} and ξ^{π^V} are the Lagrange multipliers on Constraints (14) and (15) respectively.

The infection-risk wedges are intuitive, and they are both present in the rational-expectations allocation and in the social planner's allocation. The infection-risk cost of spending time in, e.g., the B sector is the product of the probability of being susceptible $\frac{S_t^i}{S_t^i+I_t^i+R_t^i}$, the hourly infection rate $\widehat{\pi}^B$, and the cost of infection ξ^{T^i}.

The perhaps more straightforward epidemiological externality, that an infected individual may infect other individuals in the future, is a dynamic externality and it is therefore only implicitly captured in the static model by the cost of an infection, ξ^{T^i}. The static epidemiological externality wedges capture the effect that the behavior of a household has on the infection risk for other households. Under the maintained assumption that the social planner does not know who is infected, it is readily seen that the static epidemiological externality wedge is zero if the young and old are identical, in particular if the share of infected is the same in the two groups. Since the meeting technology of the standard SIR model is linear, it does not matter for the individual how many other people are in the public area, what matters is the *share* of infected individuals. If everyone double the time they spend in the public area, the share of infected individuals in the public area remains the same. For a discussion of this feature of the SIR model and alternative assumptions on the meeting technology, see Garibaldi et al. (2020).

Our model features age heterogeneity so the static epidemiological externality is not zero. Assuming a higher share of infected young people, the static epidemiological externality wedge is positive for the young and negative for the old. When a young person spends more time in the public area, she raises the share of infected in the public area and therefore the infection risk for the other individuals in the public area. Conversely, when an old person spends more time in the public area, she lowers the share of infected in the public area and thereby lowers the infection risk for the other individuals in the public area.

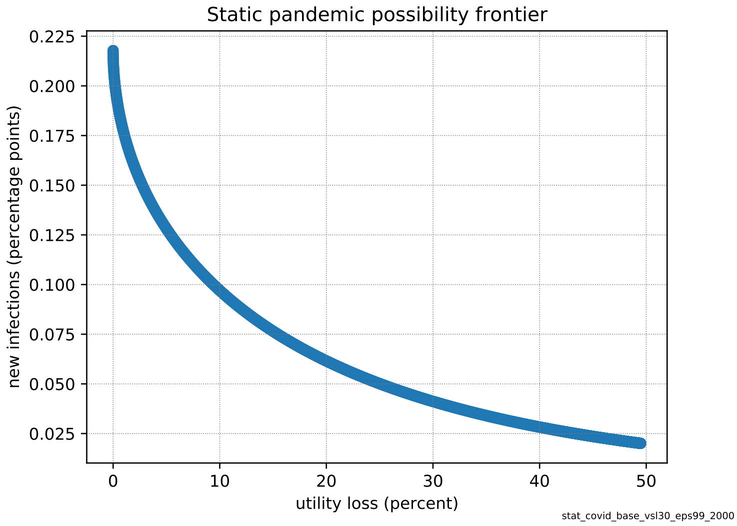

Figure 5: Static pandemic possibility frontier.

3.2 Insights from a static epidemic model

For concreteness, assume that two percent of both the old and the young are infected and that everyone else is susceptible. Furthermore, assume that the cost of an old individual becoming infected is twice as large as the cost of a young individual becoming infected.

Figure 5 displays how total utility and infections vary with the cost of an infection. The relation between infections and utility trace out a convex static pandemic-possibility frontier. A social planner that completely disregards the epidemic (i.e., sets $\xi^{T^y} = \xi^{T^o} = 0$), maximizes utility resulting in 0.22 percentage points of new infections in one day. At this point, the static pandemic possibility frontier is vertical. By the envelope theorem, the social planner can, at the margin, reduce infections at no (first-order) utility cost. However, if the social planner has a higher cost of infections, i.e., wants to significantly reduce the number of infections, then this reduction comes at a substantial utility loss.

Figure 6 displays how the optimal time allocations of the young and old vary with the cost of an infection. Absent epidemic considerations, with the cost of infections equal to zero, both the young and the old work in the workplace. They also spend time on leisure both in private and in public. As the cost of infections increases, the old reduce the amount of work they do in the workplace. Recall that we assumed that the cost of an old person becoming infected is twice the cost of a young person becoming

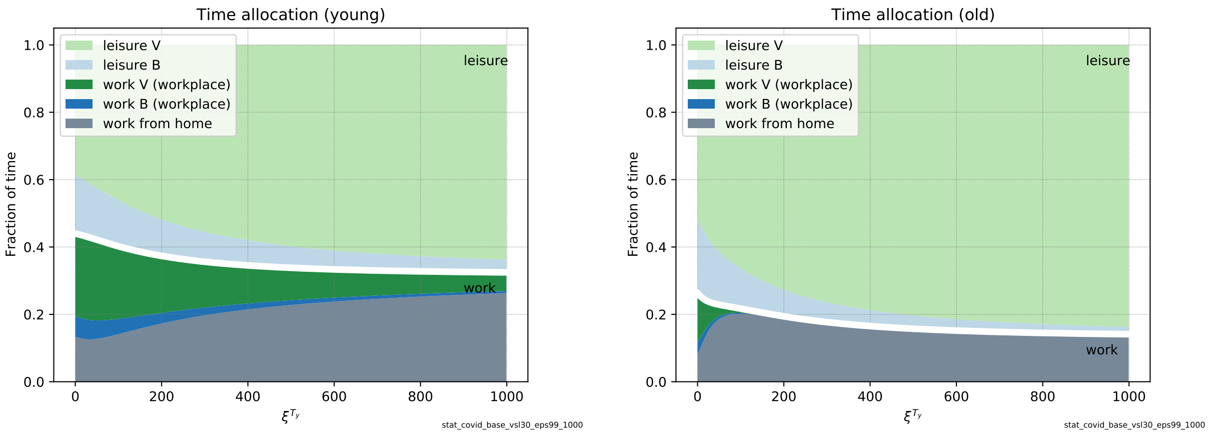

(a) Young (b) Old

Two percent of the population is assumed to be infected and the remaining population is susceptible. The cost of a new infection for the old is assumed to be twice as large, $\xi^{T_o} = 2\xi^{T_y}$. The model is specified in the next section.

Figure 6: Time allocations as a function of the cost of new infections for the young, ξ^{T_y}.

infected so the planner therefore prioritizes sending the old to work from home. As the cost of infections increases further, both the old and the young reduce the amount of leisure they spend in public and eventually also the young reduce the amount of time they spend at the workplace.

In this simple exercise we illustrated how the model works statically.[12] As the cost of new infections increases, the social planner gradually reduces epidemiologically dangerous activities. The conceptual and computational challenge is to find the true time-varying cost of an infection, taking the entire dynamic of the epidemic into account. The cost of becoming infected today depends not only on the fatality risk but also on whether it is likely that the individual would have become infected in the future anyway and whether it is likely that the individual may infect someone else in the future. In the next section, we include these dynamic considerations in the social planner's problem.

[12]In Appendix D we discuss how the choice of outer elasticity in the nested utility function, ε, affects these results.

4 The full dynamic model

We first describe the state variables in the dynamic problem, and then state the maximization problem of the planner.

State variables: S_t^i is the number of susceptible of type $i \in \{y, o\}$ (young/non-vulnerable and old/vulnerable, respectively) at time t; similarly I_t^i is the number of infected and R_t^i the number of recovered. The total population size in time t is $\sum_i \left(S_t^i + I_t^i + R_t^i\right)$, with initial conditions satisfying $I_0^y + I_0^o = 0.001, R_0^i = 0$ for $i \in \{y, o\}$ and $S_0^y + S_0^o = 0.999$. We assume that the fraction initially infected is equal in the young and the old population. In the absence of an epidemic, all young individuals live until period $\mathcal{T}^y$ and all old individuals live until period $\mathcal{T}^o < \mathcal{T}^y$.

Law of motion for the state variables: The law of motion for the state variables is then given by,

$$S_{t+1}^i = S_t^i - T_t^i \tag{21}$$
$$I_{t+1}^i = I_t^i(1 - \pi_r - \pi_{d,t}^i) + T_t^i \tag{22}$$
$$R_{t+1}^i = R_t^i + \pi_r I_t^i, \tag{23}$$

where

$$\pi_{d,t}^i = H(I_t^y + I_t^o). \tag{24}$$

The death rates $\pi_{d,t}^i$ is a function of the total number of infected individuals in the economy, since it depends on whether the hospital system is overcrowded or not.

The maximization problem: The epidemic is over at time $\mathcal{T}$ (either endogenously or by assuming that a cure instantaneously arrives at this point in time). The planner chooses all the control variables up until time $\mathcal{T} - 1$ and maximizes

$$\sum_{i \in \{y,o\}} \sum_{t=0}^{\mathcal{T}} \beta^t \left[(S_t^i + I_t^i + R_t^i) v(c_{B,t}^i, h_{B,t}^i, c_{V,t}^i, h_{V,t}^i)\right]$$
$$+ \frac{\beta^{\mathcal{T}+1} - \beta^{\mathcal{T}^o+1}}{1-\beta}(S_{\mathcal{T}}^o + I_{\mathcal{T}}^o + R_{\mathcal{T}}^o) v(c_{B,\mathcal{T}}^o, h_{B,\mathcal{T}}^o, c_{V,\mathcal{T}}^o, h_{V,\mathcal{T}}^o)$$
$$+ \frac{\beta^{\mathcal{T}+1} - \beta^{\mathcal{T}^y+1}}{1-\beta}(S_{\mathcal{T}}^y + I_{\mathcal{T}}^y + R_{\mathcal{T}}^y) v(c_{B,\mathcal{T}}^y, h_{B,\mathcal{T}}^y, c_{V,\mathcal{T}}^y, h_{V,\mathcal{T}}^y).$$

subject to the additional condition that the choice variables in period $\mathcal{T}$ maximize static period utility as in the static economic model previously described.

The epidemic is short compared to the remaining life span of the individuals: the epidemic is over at time $\mathcal{T}$, while the individuals who survive live until time $\mathcal{T}^y$ (for the young) and $\mathcal{T}^o$ (for the old). Thus, the total utility is the sum of the flow utility during the epidemic plus the remaining life span utility for the surviving individuals. We assume that the consumption and time allocations stay on their time $\mathcal{T}$ levels after the epidemic. This implies, perhaps somewhat unrealistically, that the continuation utility of all surviving individuals after $\mathcal{T}$ is unaffected by what happens during the pandemic. The above formulation of the maximization problem incorporates the entire life span of the young and the old into the analysis without explicitly modelling birth, ageing, and non-epidemic death.

The maximization problem is subject to the law of motion given by equations (21) to (24) and

$$1 = h^i_{B,t} + h^i_{V,t} + n^i_{Bh,t} + n^i_{Bw,t} + n^i_{Vh,t} + n^i_{Vw,t} \qquad i \in \{y, o\}, \quad (25)$$

$$\sum_{i \in \{y,o\}} \phi^i_t c^i_{j,t} = F_j(\phi^y_t n^y_{jh,t}, \phi^y_t n^y_{jw,t}, \phi^o_t n^o_{jh,t}, \phi^o_t n^o_{jw,t}), \qquad j \in \{B, V\}, \quad (26)$$

$$\phi^i_t = S^i_t + I^i_t + R^i_t, \qquad i \in \{y, o\}, \quad (27)$$

$$T^i_t = \widehat{\pi}_{B,t}(h^i_{B,t} + n^i_{Bw,t})S^i_t + \widehat{\pi}_{V,t} n^i_{Vw,t} S^i_t, \qquad i \in \{y, o\}, \quad (28)$$

$$\widehat{\pi}_{B,t} = \pi_B \frac{\sum_{m \in \{y,o\}} I^m_t (h^m_{B,t} + n^m_{Bw,t})}{\sum_{m \in \{y,o\}} \left[(S^m_t + I^m_t + R^m_t)(h^m_{B,t} + n^m_{Bw,t}) \right]}, \quad (29)$$

$$\widehat{\pi}_{V,t} = \pi_V \frac{\sum_{m \in \{y,o\}} I^m_t n^m_{Vw,t}}{\sum_{m \in \{y,o\}} \left[(S^m_t + I^m_t + R^m_t) n^m_{Vw,t} \right]} \quad (30)$$

for all $t \leq \mathcal{T}$ and non-negativity constraints for time and consumption quantities. Constraints (25) to (30) are the same as in the static model with an epidemic in the previous subsection, except for the time subscripts on all variables.

4.1 Calibration of the dynamic model

For most parts, the calibration of the full dynamic model is already described in the previous sections. However, there are a few new elements that we describe in turn. First of all, a period in the model corresponds to a day. We set the discount factor β such that $\beta^{365} = 0.96$.

We also need to adjust the utility function to add an intrinsic value of life, add information about the expected life length of individuals in the case of no epidemic, and calibrate the epidemic spread, recovery rate and fatality rate. We cover these three points in the following sections.

Instantaneous utility and the value of a statistical life: To properly assess the value of lost lives due to the epidemic, we need to add a "value of life" to the value function. I.e., people derive utility from merely being alive, in addition to the value they experience from leisure and consumption. Utility is additive in two terms where the first term summarizes the benefits experienced from consumption of various goods and the enjoyment of various active leisure choices. The second term is then the separate value of a statistical life, i.e., the value of being alive during that period of time. Young individuals are (naturally) assumed to live longer after the pandemic than old individuals. Thus, a planner loses more utils if a young individual dies from covid-19 than if an old individual dies, as the former loses more years of remaining life-time. We define

$$v(c_B, h_B, c_V, h_V) = u(c_B, h_B, c_V, h_V) + \underline{u}$$

with $u(\bullet)$ being defined as before in equation (3). We set $\underline{u}$ to be consistent with estimates of the value of a statistical life, using the formula

$$\frac{VSL}{c} = \frac{1-\beta^{T+1}}{1-\beta}\frac{v(c)}{cv'(c)}, \tag{31}$$

where VSL, the value of a statistical life, is expressed in period-0 units of consumption, see, e.g., Conley (1976) and Shepard and Zeckhauser (1984).

The intuition for the equation is straightforward: The utility value of an additional statistical life is $\frac{1-\beta^{T+1}}{1-\beta}v(c)$, with T being the number of expected remaining periods of life and $v(c)$ being the per-period flow utility. The utility value of an additional unit of consumption is $v'(c)$. The marginal value of an additional statistical life, in terms of consumption, is thus $\frac{1-\beta^{T+1}}{1-\beta}\frac{v(c)}{v'(c)}$. VSL/c is the value of a statistical life, expressed as a multiple of per-period consumption.

The value of a single time period, in our case a day, $VSTP$, is given by $VSTP = \frac{1-\beta}{1-\beta^{T+1}}VSL$. We therefore arrive at the formula

$$\frac{VSTP}{c} = \frac{v(c)}{cv'(c)}. \tag{32}$$

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.[13]

We use two values for the value of a statistical time period, one higher and one lower. For the higher value, we follow Glover et al. (2020) who calculate that a value of a statistical *year* of life is 11.4 times *yearly* per capita consumption, based on data from The Environmental Protection Agency (EPA) and the Department of Transportation in the US. This is a high value, relative to VSL numbers used in other contexts. Hall et al. (2020), also based on numbers from EPA, use a value of 6: a year of life is worth six times annual consumption. Others use even lower numbers. We therefore also use a number from the lower range of different estimates, and set this to 4.0: a year of life is worth roughly 4 times annual consumption, a number that is well within the range of values discussed in Viscusi and Aldy (2003).

Life span: We assume that 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.[14] 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. However, to make such an estimate requires many more assumptions. We crudely assume that the average age of a deceased in the young group is 50 years, and the corresponding age in the old group is 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), which is then the average years of life lost per death by group.

Epidemiological side of the model: We now turn to the calibration of the epidemiological parameters for the epidemic part of our model, which builds on the classical SIR model by Kermack and McKendrick (1927).

The pre-pandemic behavior implies an R_0 that we restrict our key spread parameters π_B and π_V to match. To distinguish π_B from π_V amounts to drawing distinctions between the social interactions in the leisure activity (including for those who work in that activity) and the social interactions in the workplace in the production of the good used in the in-priVate composite. We consider all interactions equally contagious and set $\pi_B = \pi_V$.

The estimates of R_0 for covid-19 are uncertain and range between at least 1.4

[13]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.

[14]The life expectancy at birth in the US is 78.6 years. However, the life expectancy, conditional on turning 60, is 83.3 (period life tables in Arias and Xu (2019)).

and 3.9; we use 2.0 in our benchmark simulations. We simulate the simplest possible SIR model with a homogeneous population given this estimate of R_0 (and a recovery rate π_r and a death rate π_d to be specified below). 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 economic model and find the $\pi_B = \pi_V$ which give the same final number of recovered in the economy, if there were no endogenous behavioral responses.[15]

In line with Atkeson (2020) and Eichenbaum et al. (2020a) 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. (2020). 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. π_d^i is a logistic function for which the midpoint of the logistic curve–where the increase in death probability is the highest–occurs at the point where the hospitals are getting over-crowded. This point is assumed to happen when the fraction of infected in the population reaches $\hat{I}$. The current death rate in any time period is thus a function of the current total number of infected:

$$\pi_{d,t}^i = H(I_t) = \pi_{d,low}^i + \frac{\pi_{d,high}^i - \pi_{d,low}^i}{1 + e^{-k(I_t - \hat{I})}} \tag{33}$$

with I_t denoting the sum of the young and the old infected. Based on US data, there were 29.4 intensive care units (ICUs) per 100,000 people at the onset of the covid crisis so we assume one ICU per 3,400 people.[16] 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.[17] Taken together, this gives us an $\hat{I} = 1/(0.03 \times 0.29 \times 3400) \approx 0.034$. 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 (2020). This means that the average infection fatality

[15] 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 slightly different time allocations and slightly different death rates.

[16] Based on information from Society of Critical Care Medicine, downloaded from `https://sccm.org/getattachment/Blog/March-2020/United-States-Resource-Availability-for-COVID-19/United-States-Resource-Availability-for-COVID-19.pdf?lang=en-US` at June 24.

[17] 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%.

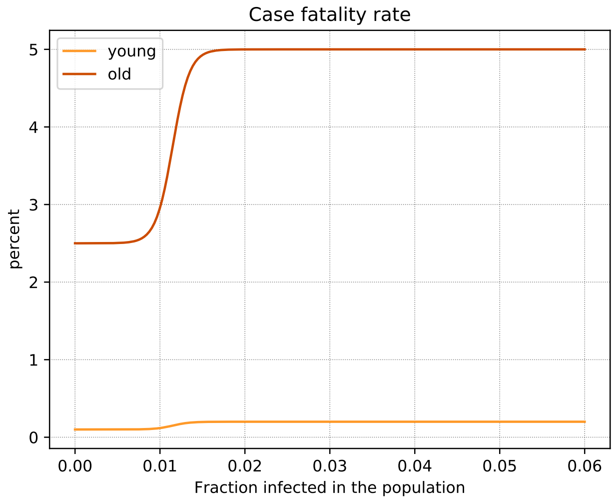

Figure 7: Infection fatality rate for young and old.

Parameter	**Description**	**Value**
Epidemic variables		
R_0	Spread factor standard SIR model	2.0
$\pi_B = \pi_V$	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]	[0.002, 0.050] · 1/18
Health care 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.

Table 3: Summary of epidemiological parameters.

rate in the population is 0.7% (if the young and the old were infected at the same rate). When the health care system is completely overburdened, the probabilities are assumed to be twice as high and thus set to $0.002 \times 1/18$ for the young and $0.05 \times 1/18$ for the old. The steepness of the curve, k, is set to 1000. The resulting mortality rates are shown in Figure 7.

A summary of the epidemiological parameters is given in Table 3.

4.2 Solution method

In Appendix E, we state the recursive reformulation of the planner's problem. The planner's problem, in a given period, can be separated into a dynamic epidemic problem and, conditional on the epidemic problem, statically computing the optimal economic allocation, as in our static model with an epidemic, described in Section 3.

The cost of a new infection ξ^{T^i} for type i depends on the difference in the marginal continuation value for the social planner between a susceptible and an infected individual,

$$\xi^{T^i} = \beta \left(\frac{\partial V}{\partial S^{i'}} - \frac{\partial V}{\partial I^{i'}} \right).$$

The epidemic dimension of the problem in a given time period is therefore reduced to the marginal continuation values $\frac{\partial V}{\partial S^{i'}}$ and $\frac{\partial V}{\partial I^{i'}}$. The planner's static economic problem is to compute the optimal economic allocation, conditional on these marginal values.

The dynamics are described by an application of the envelope theorem. The envelope theorem is here used to compute the *marginal* value of an additional susceptible, infected or recovered individual. This marginal value includes the within-period flow utility of the individual, the within-period net contribution of the individual to the aggregate budget (young are net contributors and old are net receivers of resources in the optimal allocation), the continuation value of having the additional individual in future periods, as well as the marginal effect of the individual on the epidemic. Note that we use the envelope theorem to compute the marginal value of an additional household member, which is why the expression involves the level of flow utility rather than marginal utility. We show the explicit expressions in Appendix E. There are two dynamic externalities. First, an infected individual may infect other individuals in the future. Second, an infected individual may use hospital resources, increasing the death risk for other infected individuals. The social planner takes both these externalities into account.

4.2.1 Solution algorithm

The model we have written down has a high-dimensional state space (six epidemiological states and the time dimension) and many choice variables (12 time allocations and 4 consumption allocations). Furthermore, the epidemic dynamics are distinctly non-linear. Standard approaches such as linearization or a direct application of value-function iteration are therefore not suitable for solving the model. Our method uses the insight that, within a period, the costs of infections ξ^{T^i} together with the epidemiological state summarize the intertemporal dimension.

Given the cost of an infection for both types, ξ^{T^i} for $i \in \{y, o\}$, and the epidemic state $\Omega = (S^y, I^y, R^y, S^o, I^o, R^o)$, the optimal (static) allocation is straightforwardly computed. The challenge is thus to find the correct time path for $\{\xi^{T^i}\}_{i\in\{y,o\}}$ and the correct time path for the epidemic state Ω. Our algorithm, described in detail in Appendix E, proceeds as follows. Start with a guess for the path of the costs of an infection $\{\xi_t^{T^i}\}_{t=0}^{\mathcal{T}}$ for $i \in \{y, o\}$.

Given the initial epidemic state Ω_0 and the guess for the initial-period costs of an infection $\xi_0^{T^i}$, compute the time allocations in time 0. Use the law of motion for the epidemic and the time allocations to get the epidemic state at time 1, Ω_1. Keep "rolling forward" the epidemic, using the epidemic state Ω_t and the t-period costs of an infection $\xi_t^{T^i}$ to compute the epidemic state at time $t+1$, Ω_{t+1}.

Once we have the entire time path of the epidemic, time allocations and consumption allocations, we use the envelope theorem to iterate backwards. Compute the marginal continuation values at time $\mathcal{T}$ under the assumption that the epidemic is over. Given the marginal continuation values at time $t+1$ and the allocations at time t, use the envelope theorem expressions to get the marginal continuation values at time t. Finally, use the relationship $\xi_t^{T^i} = \beta \left(\frac{\partial V}{\partial S_{t+1}} - \frac{\partial V}{\partial I_{t+1}} \right)$ to arrive at an implied path of the costs of an infection $\{\xi_t^{T^i}\}_{t=0}^{\mathcal{T}}$ for $i \in \{y, o\}$.

The new guess for the costs of an infection is a weighted average of the old guess and the implied path. Start over and iterate until convergence.

4.3 The market solution

Here we describe how individual families decide on their own time allocations. To abstract from concerns regarding redistribution of consumption between young and old, we assume that the decision unit in the market allocation is a representative family with both old and young individuals. The family does not know who is susceptible, infected, or recovered but knows how many individuals are susceptible, infected, or recovered.

We begin by describing the rational-expectations equilibrium. We then introduce a *myopic equilibrium* where households do not realize that their time allocations affect their death probabilities.

4.3.1 Rational-expectations equilibrium

The representative family (consisting of both young and old individuals) works in the two sectors of the economy as well as accrues income from owning the capital stock. It uses the income to finance consumption from the two sectors, and spends time on leisure.

The family's objective function is

$$\sum_{i\in\{y,o\}}\sum_{t=0}^{\mathcal{T}_i}\beta^t\left[(S_t^i+I_t^i+R_t^i)v(c_{B,t}^i,h_{B,t}^i,c_{V,t}^i,h_{V,t}^i)\right].$$

The family faces a sequence of budget constraints

$$\sum_{i\in\{y,o\}}(S_t^i+I_t^i+R_t^i)(p_{B,t}c_{B,t}^i+p_{V,t}c_{V,t}^i)=$$

$$\sum_{i\in\{y,o\}}(S_t^i+I_t^i+R_t^i)(w_{Bh,t}^i n_{Bh,t}^i+w_{Bw,t}^i n_{Bw,t}^i+w_{Vh,t}^i n_{Vh,t}^i+w_{Vw,t}^i n_{Vw,t}^i)$$
$$+r_{B,t}k_B+r_{V,t}k_V$$

and the time constraints

$$h_{B,t}^i+h_{V,t}^i+n_{Bh,t}^i+n_{Bw,t}^i+n_{Vh,t}^i+n_{Vw,t}^i=1$$

for $i\in\{y,o\}$.

In addition, the family takes into account how the family size and composition evolve over time:

$$T_t^i=\widehat{\pi}_{B,t}S_t^i\left(h_{B,t}^i+n_{Bw,t}^i\right)+\widehat{\pi}_{V,t}S_t^i n_{Vw,t}^i$$

for $i\in\{y,o\}$, where $\widehat{\pi}_{B,t}$ and $\widehat{\pi}_{V,t}$ are probabilities beyond the control of the family; their equilibrium determination is described below. The stocks then evolve as

$$S_{t+1}^i=S_t^i-T_t^i$$

$$I_{t+1}^i=I_t^i(1-\pi_r-\pi_{d,t}^i)+T_t^i$$

$$R^i_{t+1} = R^i_t + \pi_r I^i_t,$$

again for $i \in \{y, o\}$. $\pi^i_{d,t}$ is also a probability beyond the control of the family; its equilibrium determination is described below. This completes the description of the family planner's maximization problem.

Firms act in perfectly competitive markets and maximize profits. Hence B firms at t choose labor $(N^y_{Bh}, N^o_{Bh}, N^y_{Bw}, N^o_{Bw})$ to maximize

$$p_{B,t}F_B(N^y_{Bh}, N^y_{Bw}, N^y_{Bh}, N_{Bw^o}) - \sum_i (w^i_{Bh,t}N^i_{Bh} + w^i_{Bw,t}N^i_{Bw})$$

and, similarly for V firms, where a different combination $(N^y_{Vh}, N^o_{Vh}, N^y_{Vw}, N^o_{Vw})$ is chosen to maximize

$$p_{V,t}F_V(N^y_{Vh}, N^y_{Vw}, N^y_{Vh}, N_{Vw^o}) - \sum_i (w^i_{Vh,t}N^i_{Vh} + w^i_{Vw,t}N^i_{Vw})$$

Each of these problems result in standard conditions. The profits of the firms, which we denote by $r_{B,t}k_B$ and $r_{V,t}k_V$ respectively, are paid out as dividend to the representative household.

In equilibrium, total hours worked per age group and work type is equal to the hours per-capita times population size,

$$N^i_{Bh} = (S^i_t + I^i_t + R^i_t)n^i_{Bh}$$

and similarly for the other variables. The resource constraints in this economy are the same as those stated for the planner.

We now use bars (e.g., $\bar{x}$) to indicate aggregate variables (taken as given by the family; all families are identical so in equilibrium this distinction is dropped). The key equilibrium probabilities taken as given by the family are

$$\widehat{\pi}_{B,t} = \pi_B \frac{\sum_{m\in\{y,o\}} \bar{I}^m_t(\bar{h}^m_{B,t} + \bar{n}^m_{Bw,t})}{\sum_{m\in\{y,o\}} \left[(\bar{S}^m_t + \bar{I}^m_t + \bar{R}^m_t)(\bar{h}^m_{B,t} + \bar{n}^m_{Bw,t})\right]},$$

$$\widehat{\pi}_{V,t} = \pi_V \frac{\sum_{m\in\{y,o\}} \bar{I}^m_t \bar{n}^m_{Vw,t}}{\sum_{m\in\{y,o\}} \left[(\bar{S}^m_t + \bar{I}^m_t + \bar{R}^m_t)\bar{n}^m_{Vw,t}\right]},$$

$$\pi^i_{d,t} = H(\bar{I}^y_t + \bar{I}^o_t),$$

with $H(\bullet)$ given by equation (33).

4.3.2 Myopic equilibrium

In the myopic equilibrium, there is no feedback between the epidemic and the economy. All economic decisions are taken as if there was no connection between activities, risk of infections and deaths. Thus, the distribution of individuals over the three epidemiological states does not matter. All the decision problems, though occurring over time, are static.

5 Results

We 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. We model the exogenous end to the epidemic as an arrival of a perfect cure which is instantaneously distributed, 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 an instantaneous arrival and distribution of a vaccine or a perfect implementation of testing and tracing.[18] 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.

In the result section that follows we first focus on the case of an endogenous end of the epidemic to facilitate the understanding of the mechanisms at work. We thereafter consider a scenario with a cure arriving. Then we discuss the importance of the value of a statistical life. The section concludes with a discussion of what we can learn from testing our model on two very different epidemics: a seasonal flu and SARS.

5.1 Results with no cure of the epidemic

5.1.1 Evolution of the epidemic

In Figure 8, we show the evolution of the epidemic under the myopic market allocation, the rational-expectations market allocation, and the social planner's allocation under

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

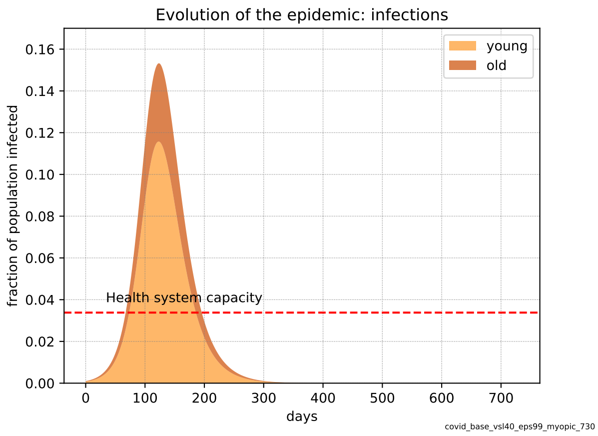

(a) The evolution of the epidemic under the myopic market allocation.

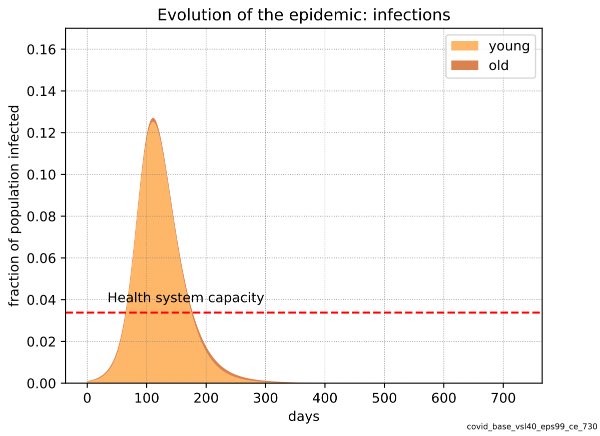

(b) The evolution of the epidemic under the rational expectations market allocation.

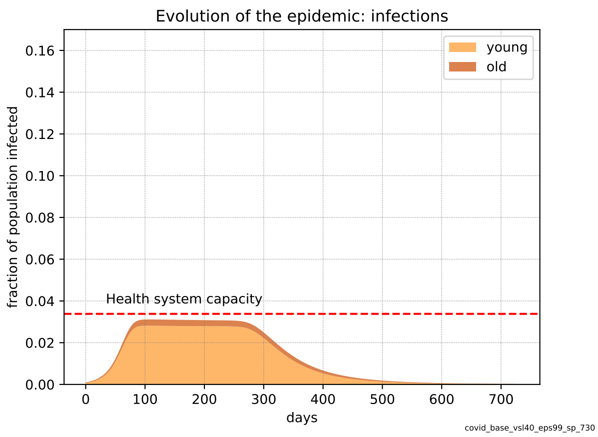

(c) The evolution of the epidemic under the social planner's allocation.

Figure 8: The evolution of the epidemic under the social planner's allocation and the two different market allocations.

the assumption that no cure arrives.[19] A first observation is that the social-planner scenario is qualitatively different compared to the myopic or the rational-expectations scenarios.

Observation 1 *The social planner "protects the health-care system" and prolongs the epidemic.*

The epidemic under the myopic market allocation 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.

The evolution of the epidemic under the rational-expectations market allocation may at first pass seem similar. The health system is overloaded and the epidemic is essentially over after 300 days. However, under rational expectations, few old become infected. The epidemic is primarily a risk for the old and under 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.

Under the social planner's allocation, the evolution of the epidemic is qualitatively different. The social planner internalizes the effects of an overloaded health system and keeps infections below the threshold for overloading. Therefore, it takes a longer time to reach herd immunity, and the epidemic is essentially over after 400 days.

Observation 2 *The social planner's allocation avoids "overshooting".*

In Figure 9, we show the evolution of the susceptible, infected, and recovered for the three scenarios. As expected, the final number of recovered is the lowest in the social planner scenario: a social planner ensures that the herd immunity threshold is reached with the smallest amount of people getting infected in total. In the social planner scenario, the final fraction of recovered is 53%, to be compared to 63% in the rational-expectations scenario and 78% in the myopic scenario. In other words, a social planner avoids "over-shooting" in the number of infected and subsequently recovered.[20]

[19] In this section, we show results from the lower value of life assumption. Results from the higher value of a statistical life assumption can be found in Appendix F. The importance of the value of a statistical life assumption will be discussed further in Subsection 5.3.

[20] Note that 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 our model, 53%.

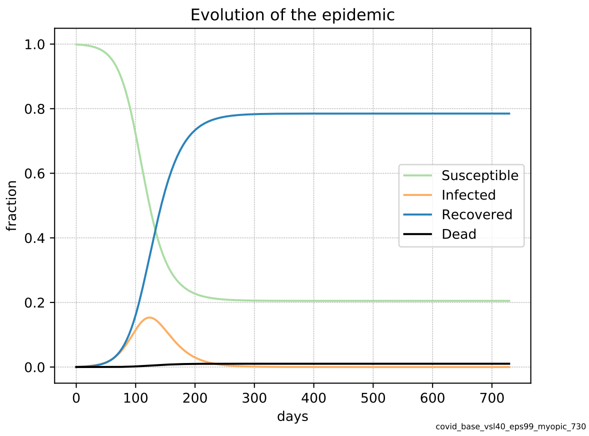

(a) SIR dynamics under the myopic market allocation.

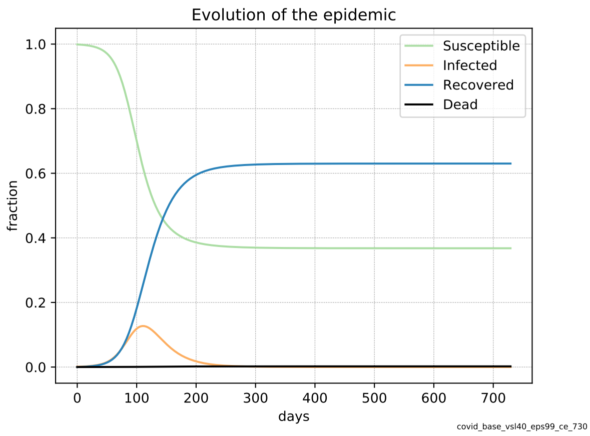

(b) SIR dynamics under the rational expectations market allocation.

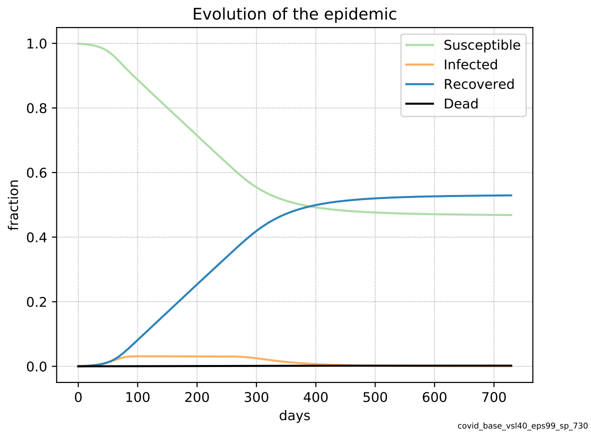

(c) SIR dynamics under the social planner's allocation.

Figure 9: SIR dynamics in the social planner's allocation and the two different market allocations.

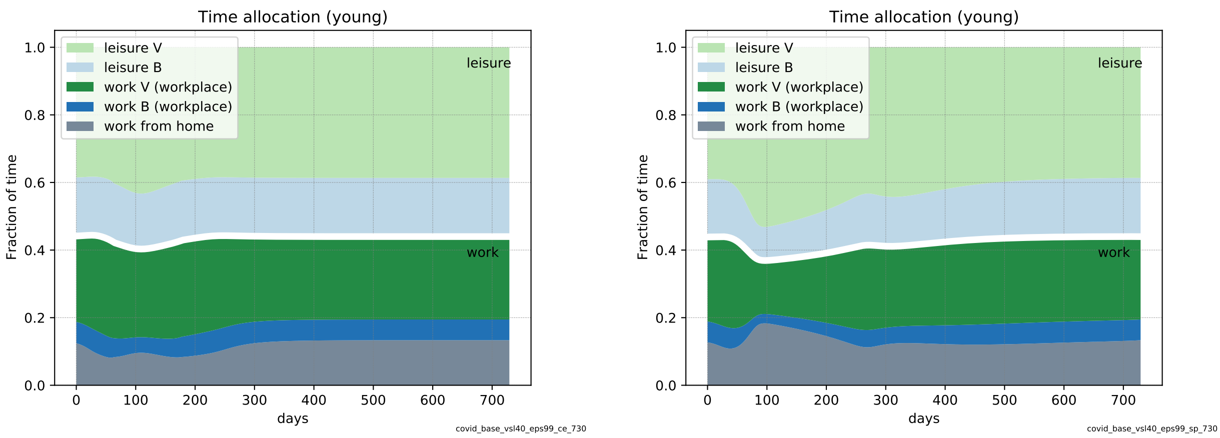

(a) Work under the rational expectations market allocation.

(b) Work under the social planner's allocation.

Figure 10: Time allocation in the social planner and the rational expectation scenario for the young.

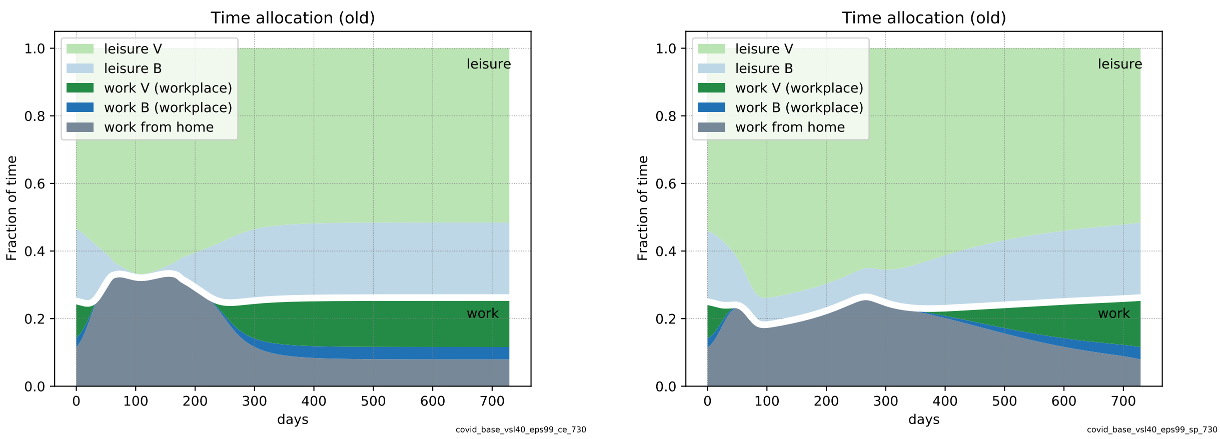

(a) Work under the rational expectations market allocation.

(b) Work under the social planner's allocation.

Figure 11: Time allocation in the social planner and the rational expectation scenario for the old.

5.1.2 Time allocations during the epidemic

In Figures 10 and 11 we unpack the time allocations under the rational-expectations and social-planner scenarios. Under the myopic market allocation, the time allocations do not change over time and are therefore not shown.[21]

Observation 3 *In the rational-expectations scenario, the old avoid all social activities.*

Under the rational-expectations market allocation, the old completely stop working in the workplace from the beginning of the epidemic. They also largely stop spending time on in-puBlic leisure, their time is instead spent working from home and enjoying a little bit more in-priVate leisure. Recall that our definition of work includes various household activities such as cleaning and cooking. The young adjust their behavior 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 but otherwise they keep their behavior relatively constant.

Observation 4 *In the social planner scenario, the old can enjoy some social leisure.*

Under the social planner's allocation, the time allocations are fine tuned to not overload the hospital system. 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 time on in-puBlic leisure, and they therefore spend considerably less time working from home compared with the rational-expectations market allocation. The young work less in the office as well, internalizing the externality of them becoming infected and subsequently infecting others.

5.1.3 Aggregate variables during the epidemic

In Figure 12, we compare the aggregate impact of the three different allocations.

Observation 5 *Both the rational-expectations allocation and the social planner's allocation substantially reduce the number of deaths.*

In Figure 12a, we plot the number of cumulative deaths under the three scenarios. Under the myopic market allocation, approximately one percent of the population dies. The rational-expectations allocation improves on this outcome significantly, reducing the number of deaths to 0.21%. The well-informed self interest of the old is sufficient

[21] Technically, time allocations do change slightly even in the myopic scenario, due to deaths in the population. However, these changes are so small that they are not visible in these graphs and can be disregarded.

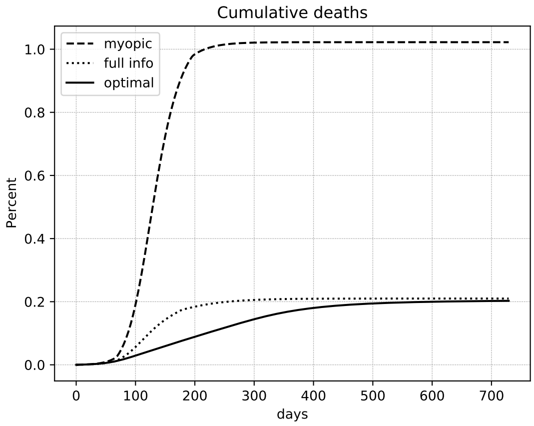

(a) Cumulative deaths under the three different scenarios.

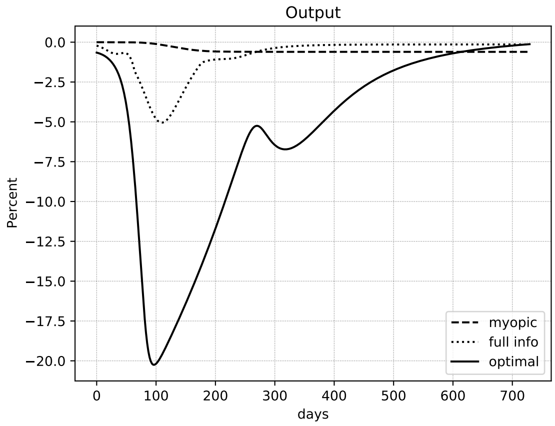

(b) Output drop under the three different scenarios.

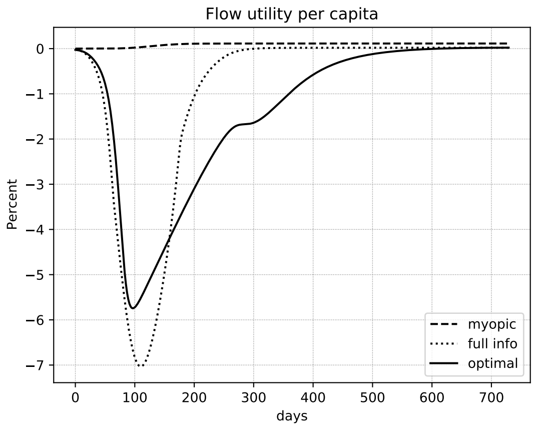

(c) Drop in flow utility under the three different scenarios.

Figure 12: Comparing the three different scenarios.

to significantly reduce the number of deaths. In the social planner scenario, the total death toll is just slightly below the rational expectation scenario: 0.20%. However, the composition of deaths during the epidemic is very different. In the rational expectation scenario 54% of the deaths are in the young population. In the social planner scenario, 24% of the deaths are in the young population. As we saw in Figure 9, the final number of recovered is lower in the social-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 social-planner scenario, the death rate 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 over-crowding on the death rates of the young. A social planner saves young lives compared to a rational-expectations scenario, which translates into more years of life saved.

Observation 6 *Output falls only modestly under the rational-expectations allocation, and substantially more in the social planner's allocation.*

Figure 12b shows the corresponding responses of aggregate output.[22] Under the myopic market allocation, output is virtually unaffected throughout the epidemic. As the population shrinks, mechanically output falls marginally. The rational-expectations market allocation implies a modest fall in output during the peak of the epidemic but the annual drop in output is a mere 1.5 percent. The social planner is willing to reduce output much more than either market allocation. During the peak of the epidemic, output drops by 20 percent and during the first year of the epidemic, output falls by 9.7 percent.

Observation 7 *The policy trade-off is not only output vs. deaths, social leisure is also an important dimension.*

Figure 12c 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 market allocation and the social planner's allocation, flow utility per capita drops substantially during the epidemic. Note however that the social planner's allocation implies a *smaller* fall in flow utility than the rational-expectation market allocation. In the rational-expectation market allocation, the old are essentially prohibited from any socially active activity and their utility is significantly reduced.

[22]Note that our measure of output is broader than only GDP, it also includes home production.

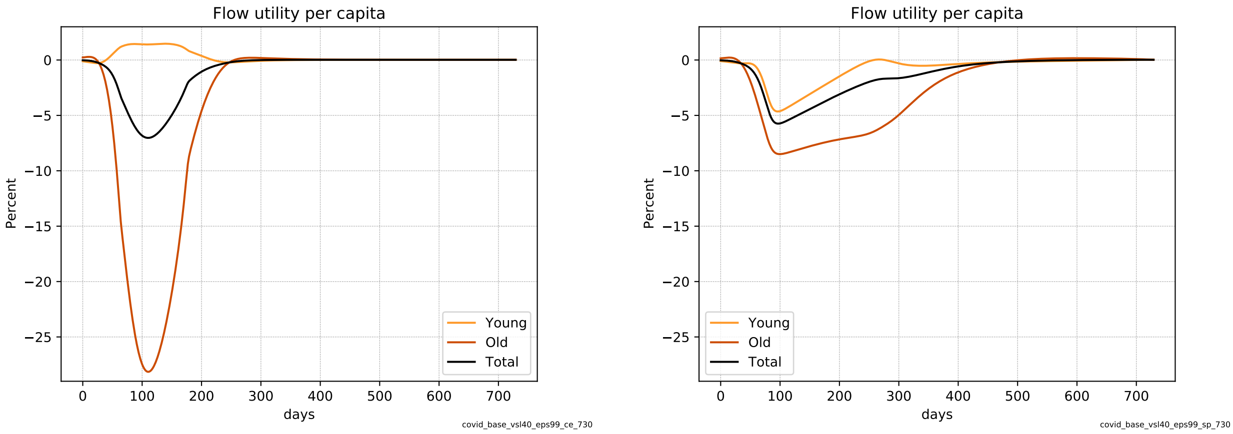

(a) Flow utility under the rational expectation market allocation.

(b) Flow utility under the social planner allocation.

Figure 13: Flow utility under the rational-expectations and the social planner's allocations.

Their utility loss is not captured by output since it is a loss of valuable leisure, not consumption, but it is an economic loss nonetheless. 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 social planner is willing to sacrifice consumption utility not only in order to save lives but also to save leisure utility for the old.

Observation 8 *The social planner's allocation benefits the flow utility of the old at the expense of the young compared to the rational-expectations allocation.*

Figure 13 unpacks the flow utility for the young vs. the old for the rational-expectations scenario and the social planner's scenario.[23] It should be noted that the flow utility per capita is calculated as the flow utility per person alive in respective group, thus the effect of deaths are not visible from these graphs, but only the instantaneous utility for the individuals that are alive. In the rational-expectations scenario, shown in Figure 13a, 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 25% during some critical weeks when the infection rate in the society is at its peak. During those critical weeks, the old have 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

[23]Under the myopic market allocation, the flow utility hardly changes and is therefore not shown. Technically, it increases slightly due to deaths in population (and the fixed capital stock), however the magnitude is so small (0.1%) that it is not visible in this graph and can be disregarded.

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 13b shows the flow utility in the social planner solution. The social planner distributes the burden of behavioral adjustment more efficiently. The young now also take a hit, with flow utility decreasing during the pandemics. 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 scenario.

5.1.4 Resulting paths for the multipliers on transmissions ξ^{T^i} and the corresponding wedges

Figure 14 shows the optimal path for the multipliers on transmissions ξ^{T^i} and the resulting wedges in the in-puBlic and the in-priVate sector for the young and the old respectively under the social planner's allocation.

Observation 9 *The social planner puts a substantial multiplier on both young and old, and the wedges distort behavior of both.*

Already in the beginning of the epidemic, the Lagrange multiplier on a newly infected old individual is high. Although the Lagrange multiplier on a newly infected young is lower than the multiplier for the old, it is also substantial.

The Lagrange multipliers are difficult to interpret directly and it is informative to translate them into "wedges". The wedges (described in equations (19) and (20)) are the hourly implied costs of being in the in-puBlic area or in the in-priVate sector workplace. As can be seen, it is generally more costly for the social planner to let old people be in those situations where they can be infected, due to their higher infection fatality rate. It is also slightly more costly to place an old individual in the workplace in the in-priVate sector than in an in-puBlic area. The reason is the different composition of people in those two areas. There is a higher concentration of infected individuals in the in-priVate workplace (in practice only young individuals work in the workplace during the height of the epidemic if the social planner were to decide), while in the in-puBlic areas there is a mix of old and young people, and the infection rate is therefore slightly lower.

In the beginning, there are extremely few infected in the economy, so even though the cost of a newly infected is high, the implied wedge is low. When the number of infected in the economy increases, the wedge also increases, even though the cost of a newly infected actually has a non-monotonic behavior for the first 100 days of the epidemic.

Towards the end of the epidemic, the Lagrange multipliers on newly infected increases again, but the implied wedges are lower than during the peak of the epidemic. The reason is of course that towards the end, the epidemic is virtually gone, and hence there are not many infected left to infect others.

Observation 10 *The rational-expectations allocation implies large multiplier on the old and large distortion of behavior of the old only.*

It is informative to compare the wedges from the social planner's allocation with the implied wedges from the rational-expectations scenario, shown in Figure 15. In this scenario, the multipliers and the resulting wedges are only functions of the direct infection risk not including the externalities. As can be seen, the differences in multipliers and wedges for the old and the young is larger. The wedges for the young are very small, while for the old, it is all but prohibitively costly to go out to enjoy any in-puBlic leisure during the height of the epidemic.

5.2 What if a cure arrives?

All previous results were under the assumption that a cure of the disease does not arrive in the foreseeable future. Now we consider what happens if a cure of the disease arrives.

Observation 11 *If a cure is expected to arrive early enough, the social planner's solution shifts qualitatively towards suppression.*

We consider a scenario where it is known that a cure of the disease arrives after one year. The epidemic evolution for the rational-expectations market allocations and the social-planner allocations under this scenario is shown in Figure 16.[24] The rational-expectations market allocations are, for all intents and purposes, unaffected by the arrival of the cure, since it arrives after the epidemic is finalized anyhow.[25] However, the social planner's allocation qualitatively shifts towards full suppression of the epidemic if the cure arrives early enough. As can be seen in Figure 16, if the cure arrives within a year, the social planner's strategy shifts qualitatively: from a strategy best described as "protect the health care system" to a "suppression" strategy.

[24]The myopic dynamics are unaffected by the arrival of a cure, so we do not show them.

[25]If the cure were expected to arrive in the midst the epidemic, also the rational-expectation allocations would adjust when the cure is close.

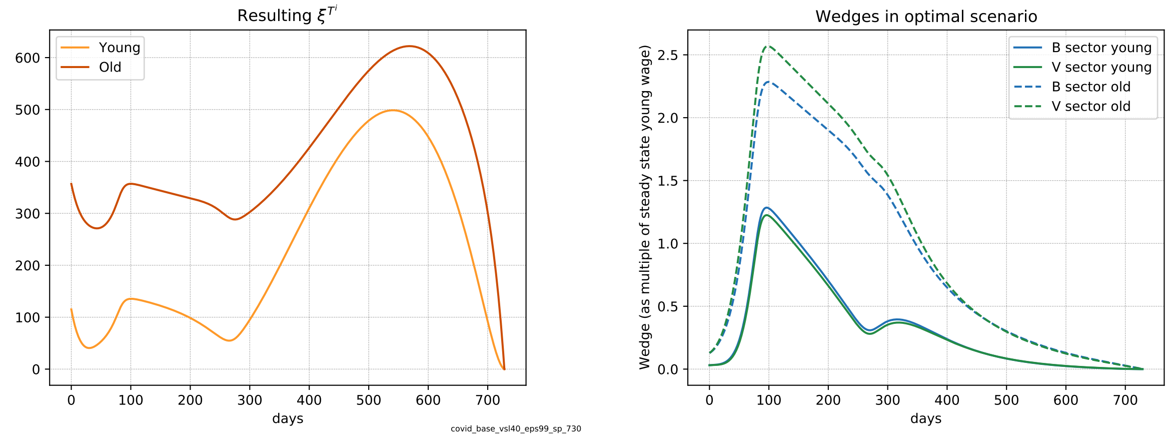

(a) Resulting ξ^{T^i} path for young and old respectively.

(b) Resulting wedges for young and old in the two sectors.

Figure 14: Paths for ξ^{T^i} and the wedges in the social planner scenario.

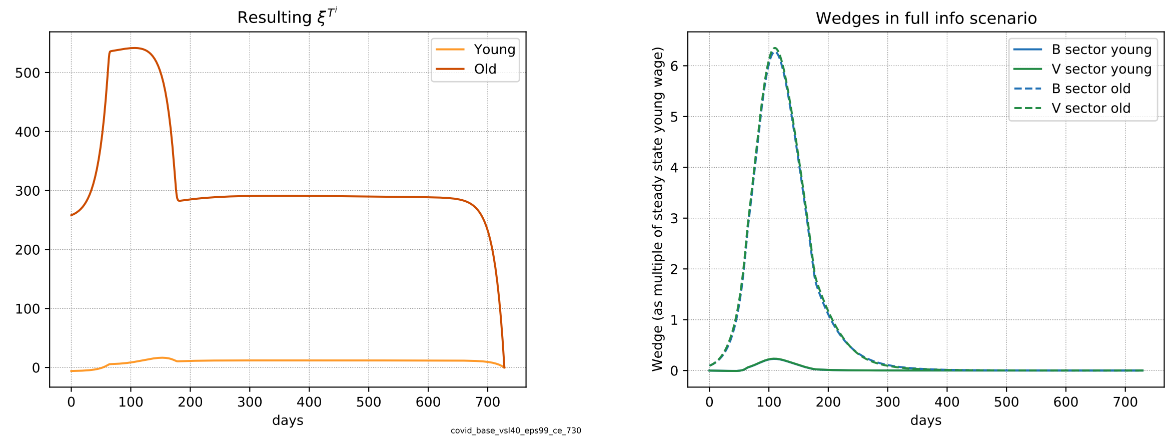

(a) Resulting ξ^{T^i} path for young and old respectively.

(b) Resulting wedges for young and old in the two sectors.

Figure 15: Paths for ξ^{T^i} and the wedges in the rational-expectations scenario.

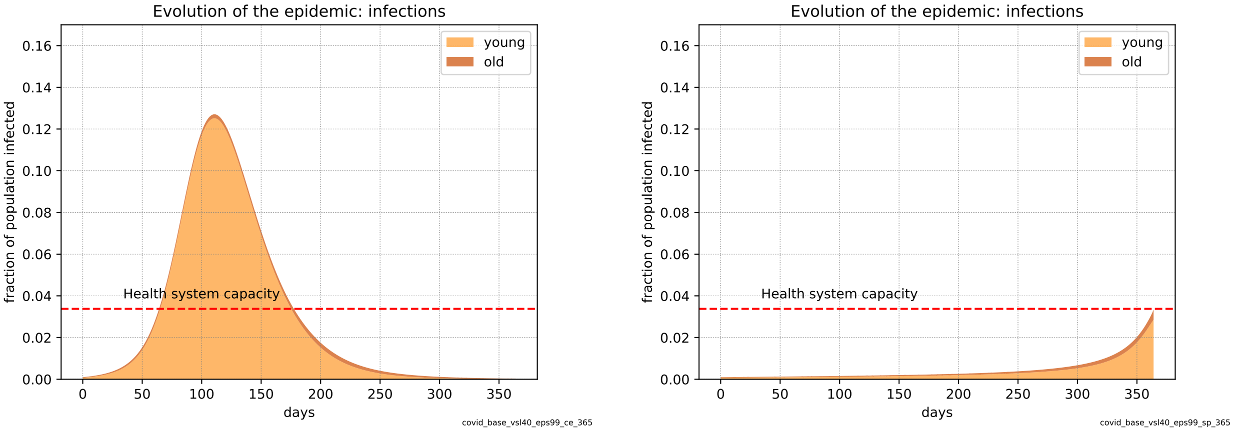

(a) The evolution of the epidemic under the rational expectations market allocation.

(b) The evolution of the epidemic under the social planner's allocation.

Figure 16: The evolution of the epidemic under the rational expectations scenario and the social planner's allocation when a cure arrives after one year.

5.3 The importance of the value of a statistical life

In the results above, we have used a value of a statistical life from the lower range of estimates used in the literature. If we instead choose a higher value of a statistical life, optimal policy can qualitatively change.

Observation 12 *The choice of value of a statistical life can qualitatively change the social planner's solution.*

We illustrate this by assuming that a cure will arrive after 18 months. When we use the lower value of a statistical life, the optimal strategy is to "protect the health care system", in other words, to flatten the curve so that it never goes above what the health care system can handle, and so that overshooting of the number of infected is avoided. However, if we instead use the higher value of a statistical life, the optimal strategy is to suppress the infection by harsher measures. This is illustrated in Figure 17, which shows the different approaches taken by a social planner assuming the lower value of a statistical life vs. the higher value of a statistical life.

The latter strategy saves lives: only 0.023 percent dies in the "suppress" scenario. When the social planner has a lower value of a statistical life the final death toll is 0.195 percent. However, the suppress strategy is of course costly in terms of output (and utility). Output falls by 22.5 percent in the suppression scenario over the first year, while the output fall in the "protect" scenario is 10.3 percent during the first

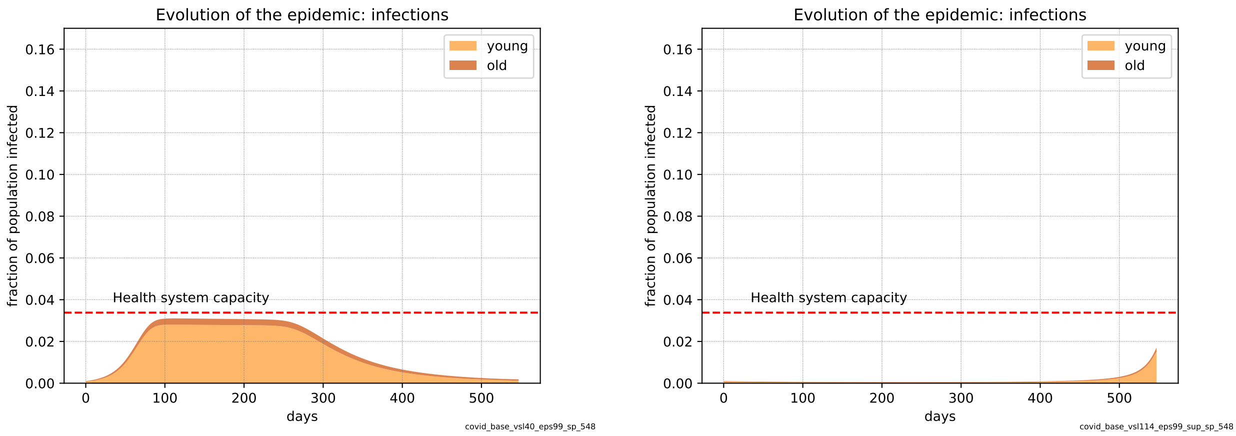

(a) Assuming the low value of a statistical life. (b) Assuming the high value of a statistical life.

Figure 17: The evolution of the epidemic under the social planner's allocation, assuming a cure arrives after 18 months, assuming different values of life.

year.

5.4 Testing different epidemics

One way to discipline the exercise and validate our model is to test it with different epidemics than 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 worse 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.

To simulate a "seasonal flu" we set the basic reproduction number, R_0, to 1.3, use a death rate of 0.045%, and an average number of days until recover of 10. This corresponds to a regular normal flu season, not to a year with a particularly severe instance of the flu.[26] 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).

A short summary of the most important insights are:

Observation 13 *The seasonal flu simulations indicate that a relatively low value of a statistical life is more in line with observed policy actions.*

[26]We also 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 Appendix G.

When simulating a seasonal flu with the higher value of a statistical life, a social planner would want to lower output by 4.6 percent during the second quarter of the epidemic, and the annual drop would be 3.4 percent. As far as we can tell, this is not how policy makers have reacted historically. With the lower value of life, the annual drop in output is substantially lower: 0.8 percent.

Observation 14 *In the case of SARS, the effective reproduction number hovers around 1 in the rational-expectations allocation. 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 so 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, that the effective reproduction number hovers around 1 in a rational expectations scenario, is also reported by Farboodi et al. (2020) and Bognanni et al. (2020).

In our calibration of the covid-19 epidemic, we do not find that the effective reproduction number stabilizes around 1 in the rational-expectations scenario. Including age heterogeneity in the model is important for our result. For the young, the risk of a covid infection does not provide a sufficiently strong motive for a precautionary response to stabilize the infection rate.

Observation 15 *In the case of SARS, the social-planner allocation both saves lives and leads to a smaller fall in output compared to the rational-expectations allocation.*

A social 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 scenario 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 so they would endogenously restrict their activities to very large extent. The total fall in output would be higher in the rational-expectations scenario. Qualitatively, this is the same type of mechanism as found in Aum et al. (2020). 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 Appendix G for the seasonal flu and Appendix H for SARS.

6 Concluding remarks

In this paper we have proposed an integrated epi-econ assessment model with visible sociological elements and we have calibrated it to pre-covid data. Aside from generating a number of insights that are specific to covid-19 policy, our quantitative model tells us that our instinct that leisure, and the extent of its social components, is quantitatively important was borne out in our experiments. We find, in particular, that this channel is very important for welfare evaluation and for designing optimal policy. That is, a mere focus on counting deaths and/or economic output leads one far astray. Let us stress that these conclusions were far from obvious at the outset: we do include this possibility in our model but the quantitative statements are entirely restricted by our calibration, which is to back out parameter values for key preference parameters using U.S. time-use data.

Of course, our framework is still restrictive in many ways. One is our modeling of the covid-19 epidemic, which is rather rudimentary. For example, cluster outbreaks, the importance of super-spreaders, and the way in which covid-19 interacts with weather and seasons seem important to consider.

Another important restriction is that we do not model the extent to which one can detect health status. In our model, if health status were publicly observable, it would be "easy" to put an end to the epidemic, so easy in fact that we do not even have to compute it: immediate isolation of all infected individuals away from social activities for a few weeks will suffice. In practice, there still seems to be a lack of consensus on how easy it is to detect covid-19; for example, what is the fraction of people who are asymptomatic but still can infect others? From the perspective of our model, thus, we would need to introduce into our framework a degree of detectability, as in some papers in the literature (see our initial literature review). This is a rather challenging task that we have postponed for future work.

One would also want to extend the model to include more heterogeneity. In our representative-family setting, this is rather straightforward. In its current version, our model distinguishes between young and old. A third subgroup that would be important to model explicitly to fully capture the impact of covid-19 is the very old (or most vulnerable). Individuals living in retirement or nursing homes cannot withdraw from social interaction with staff or nurses even if they would want to. Thus, these vulnerable individuals are exposed to the virus to the extent that the virus is present in the population working in these homes. In an attempt to stop the spread of the epidemic in nursing homes in Sweden, visitors were banned, which meant that many elderly did not see their loved ones for many months. This further underlines the broader message:

to fully evaluate the impact of more infected in the young and healthy population, we need to take into account the externalities. By how much does the risk of spreading the infection to the more vulnerable increase? But also, what is the cost of pre-emptive measures to stop the spread in terms of lost utility from social interaction?[27]

It would also seem critical from many perspectives to consider inequality, including how it interacts with the extent to which people are informed of the epidemic; our myopic equilibrium assumes that no one knows what is happening, but in reality a fraction of people is probably very well informed whereas many others are not or remain sceptical. Relatedly, we could explicitly introduce easy-to-adopt measures, such as hand-washing, keeping distances while at work, and mask wearing, that would then be used differentially in the population. These inclusions are urgent but also challenging (though not impossible).

Uncertainty also seems important to include. First, key features of the epidemic are only learned gradually, and second, when a cure/vaccine arrives is not only unknown initially but—as we show here—very important. We have chosen not to model this uncertainty explicitly at this point. However, we believe that it can be studied straightforwardly using recent computational methods, at least to a first order: see Boppart et al. (2018). I.e., it is possible to examine how information shocks (say, about an epidemiological parameter) at different points in time affect the economic and epidemiological evolution; these effects would depend on when they hit, but the methods allow for such analysis.

[27]Note also that other viruses may turn out to be particularly harmful to different groups that are also unable to protect themselves; one such group could be smaller children.

References

Acemoglu, D., Chernozhukov, V., Werning, I., Whinston, M.D., 2020a. A multi-risk SIR model with optimally targeted lockdown. Technical Report. National Bureau of Economic Research.

Acemoglu, D., Makhdoumi, A., Malekian, A., Ozdaglar, A.E., 2020b. Testing, voluntary social distancing and the spread of an infection. Technical Report. National Bureau of Economic Research.

Adda, J., 2007. Behavior towards health risks: An empirical study using the "Mad Cow" crisis as an experiment. Journal of Risk and Uncertainty 35, 285–305.

Aguiar, M., Hurst, E., 2007. Measuring trends in leisure: The allocation of time over five decades. The Quarterly Journal of Economics 122, 969–1006.

Alon, U., Baron, T., Bar-On, Y., Cornfeld, O., Milo, R., Yashiv, E., 2020. COVID-19: Looking for the Exit. Technical Report. working paper.

Alvarez, F.E., Argente, D., Lippi, F., 2020. A simple planning problem for covid-19 lockdown. Technical Report. National Bureau of Economic Research.

Arias, E., Xu, J., 2019. United states life tables, 2017. National Vital Statistics Reports 68.

Atkeson, A., 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.

Aum, S., Lee, S.Y.T., Shin, Y., 2020. Inequality of fear and self-quarantine: Is there a trade-off between GDP and public health? Technical Report. National Bureau of Economic Research.

Bethune, Z.A., Korinek, A., 2020. Covid-19 infection externalities: Trading off lives vs. livelihoods. Technical Report. National Bureau of Economic Research.

Biggerstaff, M., Cauchemez, S., Reed, C., Gambhir, M., Finelli, L., 2014. Estimates of the reproduction number for seasonal, pandemic, and zoonotic influenza: a systematic review of the literature. BMC infectious diseases 14, 480.

Bodenstein, M., Corsetti, G., Guerrieri, L., 2020. Social distancing and supply disruptions in a pandemic .

Bognanni, M., Hanley, D., Kolliner, D., Mitman, K., 2020. Economic activity and covid-19 transmission: Evidence from an estimated economic-epidemiological model.

Boppart, T., Krusell, P., 2020. Labor supply in the past, present, and future: a balanced-growth perspective. Journal of Political Economy 128, 118–157.

Boppart, T., Krusell, P., Mitman, K., 2018. Exploiting MIT shocks in heterogeneous-agent economies: the impulse response as a numerical derivative. Journal of Economic Dynamics and Control 89, 68–92.

Brotherhood, L., Kircher, P., Santos, C., Tertilt, M., 2020. An economic model of the covid-19 epidemic: The importance of testing and age-specific policies .

Chan, T.Y., Hamilton, B.H., Papageorge, N.W., 2016. Health, risky behaviour and the value of medical innovation for infectious disease. The Review of Economic Studies 83, 1465–1510.

Chang, R., Velasco, A., 2020. Economic policy incentives to preserve lives and livelihoods. Technical Report. National Bureau of Economic Research.

Conley, B.C., 1976. The value of human life in the demand for safety. American Economic Review 66, 45–55.

Dushoff, J., Plotkin, J.B., Viboud, C., Earn, D.J., Simonsen, L., 2006. Mortality due to influenza in the United States–An annualized regression approach using multiple-cause mortality data. American Journal of Epidemiology 163, 181–187.

Eichenbaum, M.S., Rebelo, S., Trabandt, M., 2020a. The macroeconomics of epidemics. Technical Report. National Bureau of Economic Research.

Eichenbaum, M.S., Rebelo, S., Trabandt, M., 2020b. The Macroeconomics of Testing and Quarantining. Technical Report.

Farboodi, M., Jarosch, G., Shimer, R., 2020. Internal and external effects of social distancing in a pandemic. Technical Report. National Bureau of Economic Research.

Faust, J.S., Del Rio, C., 2020. Assessment of deaths from covid-19 and from seasonal influenza. JAMA Internal Medicine .

Garibaldi, P., Moen, E.R., Pissarides, C.A., 2020. Modelling contacts and transitions in the sir epidemics model. Covid Economics 5, 1–20.

Geoffard, P.Y., Philipson, T., 1996. Rational epidemics and their public control. International economic review , 603–624.

Giagheddu, M., Papetti, A., 2020. The macroeconomics of age-varying epidemics.

Giannitsarou, C., Kissler, S., Toxvaerd, F., 2020. Waning immunity and the second wave: Some projections for sars-cov-2 .

Glover, A., Heathcote, J., Krueger, D., Rios-Rull, J.V., 2020. Health versus wealth: On the distributional effects of controlling a pandemic. PIER Working Paper .

Golosov, M., Hassler, J., Krusell, P., Tsyvinski, A., 2014. Optimal taxes on fossil fuel in general equilibrium. Econometrica 82, 41–88.

Greenwood, J., Kircher, P., Santos, C., Tertilt, M., 2019. An equilibrium model of the african HIV/AIDS epidemic. Econometrica 87, 1081–1113.

Hall, R.E., Jones, C.I., Klenow, P.J., 2020. Trading off consumption and covid-19 deaths. Technical Report. National Bureau of Economic Research.

Jones, C.J., Philippon, T., Venkateswaran, V., 2020. Optimal mitigation policies in a pandemic: Social distancing and working from home. Technical Report. National Bureau of Economic Research.

Kapicka, M., Rupert, P., 2020. Labor markets during pandemics. Manuscript, UC Santa Barbara .

Kaplan, G., Moll, B., Violante, G., 2020. The great lockdown and the big stimulus: Tracing the pandemic possibility frontier for the US. Technical Report. National Bureau of Economic Research.

Kermack, W.O., McKendrick, A.G., 1927. A contribution to the mathematical theory of epidemics. Proceedings of the Royal Society of London. Series A, Containing papers of a mathematical and physical character 115, 700–721.

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

Krueger, D., Uhlig, H., Xie, T., 2020. Macroeconomic dynamics and reallocation in an epidemic. Technical Report. National Bureau of Economic Research.

Nordhaus, W.D., Boyer, J., 2000. Warming the world: economic models of global warming. MIT press.

Petersen, E., Koopmans, M., Go, U., Hamer, D.H., Petrosillo, N., Castelli, F., Storgaard, M., Al Khalili, S., Simonsen, L., 2020. Comparing sars-cov-2 with sars-cov and influenza pandemics. The Lancet infectious diseases .

Piguillem, F., Shi, L., 2020. Optimal covid-19 quarantine and testing policies .

Rolfes, M.A., Foppa, I.M., Garg, S., Flannery, B., Brammer, L., Singleton, J.A., Burns, E., Jernigan, D., Olsen, S.J., Bresee, J., et al., 2018. Annual estimates of the burden of seasonal influenza in the United States: a tool for strengthening influenza surveillance and preparedness. Influenza and other respiratory viruses 12, 132–137.

Shepard, D.S., Zeckhauser, R., 1984. Survival versus consumption. Management Science 30, 423–439.

Viscusi, W.K., Aldy, J.E., 2003. The value of a statistical life: A critical review of market estimates throughout the world. Journal of Risk and Uncertainty 27, 5–76.

van Vlokhoven, H., 2020. Policy during an epidemic with super-spreaders. Available at SSRN 3600441 .

A Leisure spent at home vs. outside home

This section provides more details on leisure time spent in the two categories in-puBlic, i.e., socially intense, and in-priVate.

We classify activities as *not* socially intense if it took place in the respondent's home or yard. Moreover, we classify personal care activities (e.g., grooming and personal activities) coded with location code "Blank" in the survey as not socially intense. Lastly, 0.3% of the observations in the data are coded with "Unspecified place". For these observations, we code those where it is plausible that the activity took place in the home as not socially intense.[28]

The socially intense activities are consequently the activities that took place outside home. Examples of locations for these activities include someone else's home, store/mall, restaurant or bar, and gym/health club.

To understand what these broad categorizations mean in practice, Figure 18 shows socially intense leisure and not-socially-intense leisure broken down on a finer level. For instance, the category "Eating and drinking" shows up in both types of leisure: young spend on average 34 minutes per day eating and drinking outside their home (socially intense leisure), and 35 minutes on eating and drinking at home (not-socially-intense leisure). The largest category for leisure is "Socializing, relaxing, and leisure", both when it comes to socially intense leisure and not-socially-intense leisure. On a finer classification level, the most common subcategory within "Socializing, relaxing, and leisure" for the socially intense type is "Socializing and communicating", while it for the not socially intense type is "Relaxing and leisure", which roughly translates to watching television at home.

[28] As an example, we code the activity "Caring for and helping household children" as not socially intense, while "Participating in sports" is classified as socially intense.

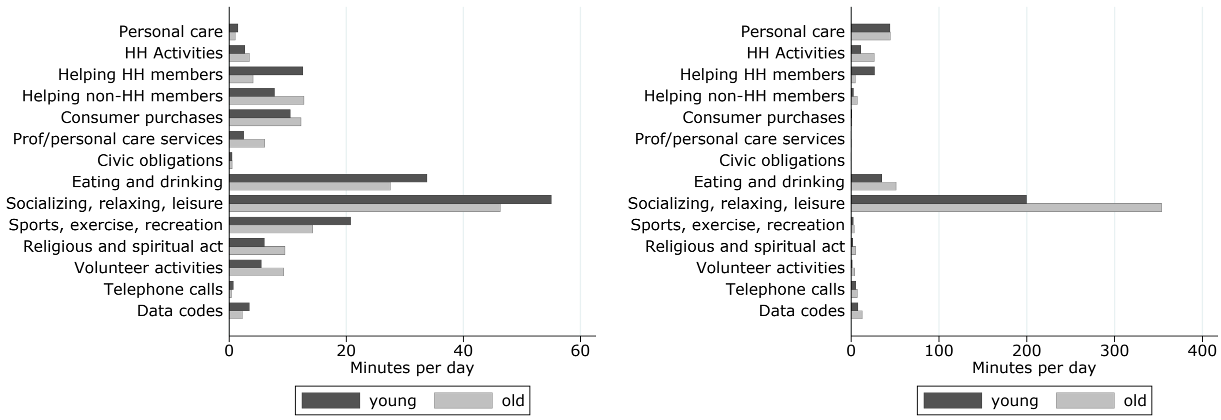

(a) Socially intense in-puBlic leisure

(b) Not socially intense in-priVate leisure

"Data codes" refer to observations where the respondent couldn't remember or refused to answer. A full day is $24 \cdot 60 = 1440$ minutes. Source: ATUS, 2018.

Figure 18: Average minutes per day spent in socially intense leisure activities and not socially intense activities, by two-digit categories (includes associated traveling).

B Sector classification

Table 4 gives an example of how different sectors are classified as being fully, to a high extent, somewhat, or not at all employing people who are in social contact with customers.

2017 NAICS	**Description**	**Employees Thousands**	**Active sector production?**
211	Oil and gas extraction	683.3	No
...			
441	Motor vehicle and parts dealers	2,021.2	No
445	Food and beverage stores	3,087.1	High extent
452	General merchandise stores	3,104.9	High extent
...			
711	Performing arts, spectator sports, and related industries	506.0	Yes
...			
7131	Amusement parks and arcades	211.0	Yes
7132	Gambling industries (except casino hotels)	120.4	High extent
...			
721	Accommodation	2,028.4	Yes
722	Food services and drinking places	11,926.3	Yes
...			
8121	Personal care services	727.8	Yes
8122	Death care services	137.1	No
...			
-	Non-agriculture self-employed	9,453.4	Somewhat
	Total number of employees	161,037.7	

Table 4: Illustrative example of sector classification.

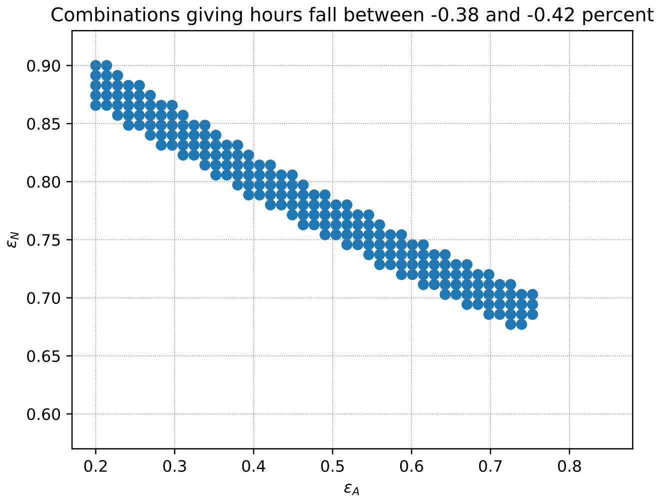

Criterion a) A realistic income/substitution effect is defined as if a TFP increase of 2% leads to hours worked falling by between 0.38 and 0.42 percent. Criterion b) Young spend a larger fraction of their leisure time in the socially intense type than the old do.

Figure 19: Combinations of ε_B and ε_V that satisfy both criterion a) and b).

C Choosing elasticities: ε_B and ε_V

To pin down our choice of ε_B and ε_B, the elasticities for the consumption-leisure bundle in the socially intense B sector and the non-social V-sector respectively, we put the two following restrictions on our utility function: a) the income effect should dominate the substitution effect in a realistic way, and b) the young should spend a larger fraction of their leisure in the socially intense B activity. These two restrictions narrow down the set of $\varepsilon_B/\varepsilon_V$ we can choose from substantially. Figure 19 shows a number of combinations of elasticities that satisfies those two restrictions. In the graph, combinations that lie "south-east" of the marked area are combinations for which the second requirement is not fulfilled.

As our base case, we pick $\varepsilon_B = 0.41$ and $\varepsilon_V = 0.8$, but note that for basically all permissible combinations, we have that $\varepsilon_B < \varepsilon_V$ and both elasticities being smaller than 1. This means that for any of the combinations we could choose as an alternative, the model behaves extremely similar and no insights of the working of the economy change.

To understand why these elasticities affect how leisure is distributed between the socially intense B type and the not-social V type, think about a marginal increase in leisure: how should it be split up between the two types of leisure? The answer is of course so that the marginal utilities of the two types still are equalized. The marginal utility with respect to the socially intense B type of leisure is given by:

$$\begin{aligned} u_{h_B} &= \frac{\partial u}{\partial \widetilde{u}_B} \cdot \frac{\partial \widetilde{u}_B}{\partial h_B} \\ &= \frac{\partial u}{\partial \widetilde{u}_B} \cdot \left(\lambda_B c_B^{\frac{\varepsilon_B - 1}{\varepsilon_B}} + (1-\lambda_B) h_B^{\frac{\varepsilon_B - 1}{\varepsilon_B}} \right)^{\frac{1}{\varepsilon_B - 1}} (1-\lambda_B) h_B^{-\frac{1}{\varepsilon_B}} \end{aligned}$$

Thus, the elasticity of the marginal utility with respect to leisure ($i \in \{B, V\}$):

$$\begin{aligned} \frac{d \log u_{h_i}}{d \log h_i} &= \frac{d}{d \log h_i} \Bigg[\log \frac{\partial u}{\partial \widetilde{u}_i} + \log(1-\lambda_i) \\ &\quad + \frac{1}{\varepsilon_i - 1} \log \left(\lambda_i c_i^{\frac{\varepsilon_i - 1}{\varepsilon_i}} + (1-\lambda_i) h_i^{\frac{\varepsilon_i - 1}{\varepsilon_i}} \right) - \frac{1}{\varepsilon_i} \log h_i \Bigg] \\ &= \frac{1}{\varepsilon_i} \left(\frac{(1-\lambda_i) h_i^{\frac{\varepsilon_i - 1}{\varepsilon_i}}}{\lambda_i c_i^{\frac{\varepsilon_i - 1}{\varepsilon_i}} + (1-\lambda_i) h_i^{\frac{\varepsilon_i - 1}{\varepsilon_i}}} \right) - \frac{1}{\varepsilon_i} \\ &= -\frac{1}{\varepsilon_i} \left(1 - \frac{(1-\lambda_i) h_i^{\frac{\varepsilon_i - 1}{\varepsilon_i}}}{\lambda_i c_i^{\frac{\varepsilon_i - 1}{\varepsilon_i}} + (1-\lambda_i) h_i^{\frac{\varepsilon_i - 1}{\varepsilon_i}}} \right) \end{aligned}$$

Hence, the elasticity of the marginal utility with respect to leisure depends on the CES elasticity (and the consumption/leisure terms within the respective bundles). The relative size of the terms in the bundles are primarily determined by the other exogenously set calibration targets and do not depend on ε_i to any larger extent. The relationship between ε_B and ε_V is therefore crucial for determining where to spend a marginal increase in leisure. With $\varepsilon_B < \varepsilon_V$, a marginal increase in leisure is spent proportionally more on the not-social V good (had the bundles been exactly the same).

D More on the choice of ε

To get more insights on how the outer elasticity in the nested CES utility function impacts the time allocations in a model with a pandemic, we simulate the same type of graphs as we did in section 3.2, using different values of ε.

Recall from section 2.3 that the flow utility for an individual is given by

$$u(c_B, h_B, c_V, h_V) = \log \text{CES}(\tilde{c}_B, \tilde{c}_V; \lambda, \varepsilon), \quad (34)$$

$$\tilde{c}_B = \text{CES}(c_B, h_B; \lambda_B, \varepsilon_B), \quad (35)$$

$$\tilde{c}_V = \text{CES}(c_V, h_V; \lambda_V, \varepsilon_V). \quad (36)$$

The nested CES structure captures that to consume a good, both social and non-social, involves spending time with the good.

For concreteness, assume that eight percent of the population is infected, and the rest is susceptible. The infected are evenly spread out in the young and the old population. Moreover, assume that the cost of an infected old is 50 times the cost of an infected young (this might sound high, but as shown in the full dynamic model this outcome is not at all extreme compared to the outcome in the rational expectations scenario).

Figure 20 shows the time allocations for the young and the old as a function of the cost of a young infection for our baseline scenario $\varepsilon = 1.0$. As discussed in section 3.2, as the cost of infections increases, individuals reallocate their time. In this example, very soon the old stop working in the workplace and all work is done from home. They also gradually reduce their in-puBlic leisure time.

In Figure 21 we show the same graphs from the same experiment, but setting $\varepsilon = 0.4$.[29] Here we see that the reallocation pattern for the old is qualitatively different. When it becomes very costly to be infected, they cut down on work in the workplace first, just as before. Then the old cut down on leisure in-puBlic, also as before. However, with an elasticity below one, the composite in-puBlic good and the composite in-priVate good are complements, and therefore the marginal utility from in-priVate leisure falls. Thus, if the old are prevented from enjoying in-puBlic leisure, their time spent watching TV will eventually fall as well. This implication we argue is implausible, and therefore advice against an $\varepsilon < 1$.

Figure 22 shows the same experiment setting $\varepsilon = 1.3$. The in-puBlic and the in-priVate composite goods are now substitutes, and thus the reallocation from in-puBlic

[29]Using another outer elasticity means that we also have to recalibrate the other remaining parameters, which we also do. See Section 2.3.

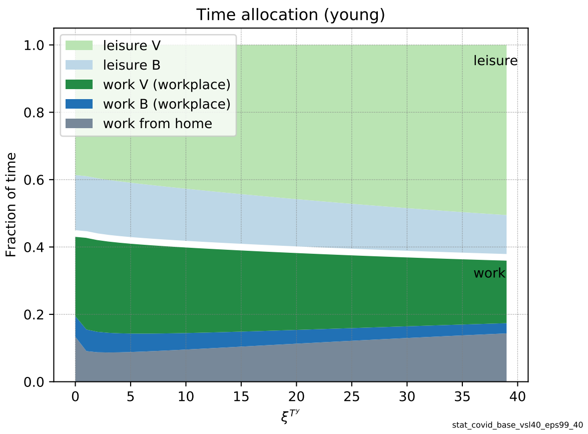

(a) Young

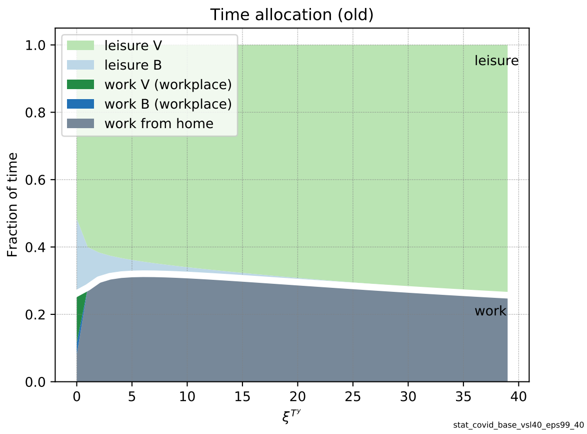

(b) Old

Assuming that eight percent of the population infected and the remaining population susceptible. The cost of a new infection for the old is assumed to be 50 times as large as the cost for the young, $\xi^{T^o} = 50\xi^{T^y}$.

Figure 20: Time allocations as a function of the cost of new infections for the young, ξ^{T^y}. $\varepsilon = 1.0$.

leisure to in-priVate leisure is somewhat stronger than in our base case. However, the difference compared to our base case is relatively small, and does not translate into any meaningful difference in results from the full model.

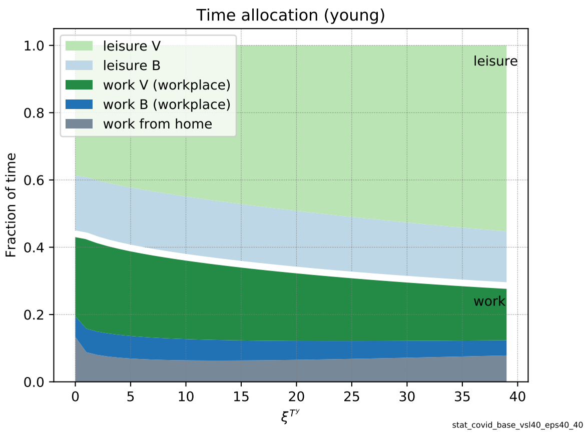

(a) Young, $\varepsilon = 0.4$.

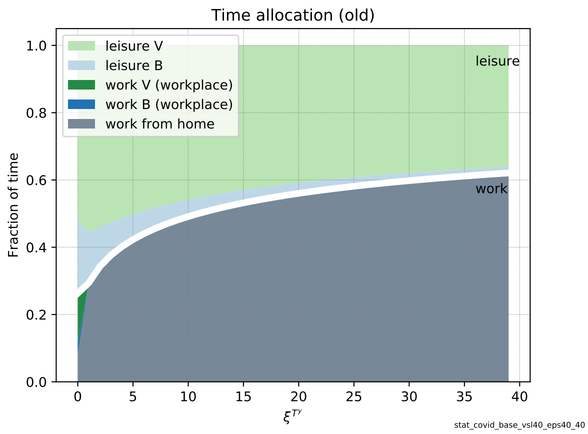

(b) Old, $\varepsilon = 0.4$.

Assuming that eight percent of the population infected and the remaining population susceptible. The cost of a new infection for the old is assumed to be 50 times as large as the cost for the young, $\xi^{T^o} = 50\xi^{T^y}$.

Figure 21: Time allocations as a function of the cost of new infections for the young, ξ^{T^y}. $\varepsilon = 0.4$.

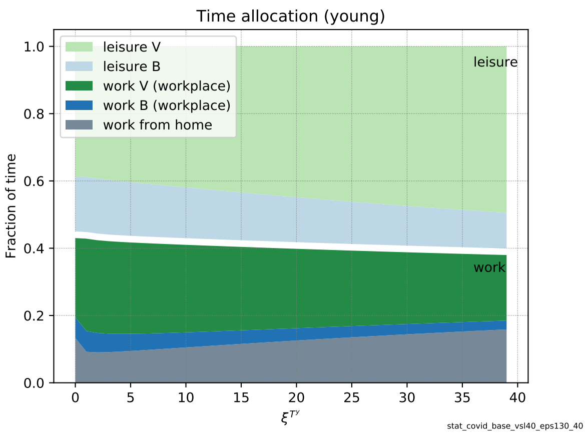

(a) Young, $\varepsilon = 1.3$.

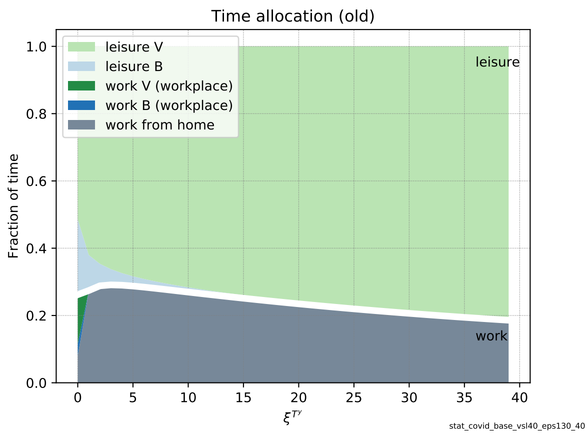

(b) Old, $\varepsilon = 1.3$.

Assuming that eight percent of the population infected and the remaining population susceptible. The cost of a new infection for the old is assumed to be 50 times as large as the cost for the young, $\xi^{T^o} = 50\xi^{T^y}$.

Figure 22: Time allocations as a function of the cost of new infections for the young, ξ^{T^y}. $\varepsilon = 1.3$.

E Solution method

E.1 Recursive formulation of the planner's problem

We denote the complete SIR state by Ω, so that $\Omega = (S^y, I^y, R^y, S^o, I^o, R^o)$. The full recursive problem of the planner is to maximize

$$V_t(\Omega) = \max \left\{ \sum_i (S^i + I^i + R^i) v(c_B^i, h_B^i, c_V^i, h_V^i) + \beta V_{t+1}(\Omega') \right\} \qquad t < \mathcal{T}$$

subject to the production technologies and resource constraints (with the bracked variables indicating the Lagrange multiplier for the different constraints, and using the short-hand notation $\phi^i = S^i + I^i + R^i$),

$$F_B(\phi^y n_{Bh}^y, \phi^y n_{Bw}^y, \phi^o n_{Bh}^o, \phi^o n_{Bw}^o) = \sum_i \phi^i c_B^i, \qquad [\lambda^{y_B}] \tag{37}$$

$$F_V(\phi^y n_{Vh}^y, \phi^y n_{Vw}^y, \phi^o n_{Vh}^o, \phi^o n_{Vw}^o) = \sum_i \phi^i c_V^i, \qquad [\lambda^{y_V}] \tag{38}$$

the time-allocation constraints (for both young and old),

$$1 = h_B^i + h_V^i + n_{Bh}^i + n_{Vh}^i + n_{Bw}^i + n_{Vw}^i, \qquad [\lambda^{n^i}] \tag{39}$$

$$h_B^i \geq 0, \qquad [\kappa^{h_B^i}] \tag{40}$$

$$h_V^i \geq 0, \qquad [\kappa^{h_V^i}] \tag{41}$$

$$n_{Bh}^i \geq 0, \qquad [\kappa^{n_{Bh}^i}] \tag{42}$$

$$n_{Vh}^i \geq 0, \qquad [\kappa^{n_{Vh}^i}] \tag{43}$$

$$n_{Bw}^i \geq 0, \qquad [\kappa^{n_{Bw}^i}] \tag{44}$$

$$n_{Vw}^i \geq 0, \qquad [\kappa^{n_{Vw}^i}] \tag{45}$$

the evolution of the epidemic (for both young and old),

$$T^i = (\widehat{\pi}^B (h_B^i + n_{Bw}^i) + \widehat{\pi}^V n_{Vw}^i) S^i, \qquad [\xi^{T^i}] \tag{46}$$

$$S^i - T^i = S^{i'}, \qquad [\xi^{S^i}] \tag{47}$$

$$(1 - \pi_d^i - \pi_r) I^i + T^i = I^{i'}, \qquad [\xi^{I^i}] \tag{48}$$

$$\pi_r I^i + R^i = R^{i'}, \qquad [\xi^{R^i}] \tag{49}$$

$$\pi_d^i = H(I^y + I^o), \qquad [\xi^{\pi_d^i}] \tag{50}$$

and the infection rates

$$\widehat{\pi}^B = \pi^B \frac{\sum_i I^i(h^i_B + n^i_{Bw})}{\sum_i (S^i + I^i + R^i)(h^i_B + n^i_{Bw})}, \qquad [\xi^{\pi^B}] \tag{51}$$

$$\widehat{\pi}^V = \pi^V \frac{\sum_i I^i n^i_{Vw}}{\sum_i (S^i + I^i + R^i) n^i_{Vw}}. \qquad [\xi^{\pi^V}] \tag{52}$$

The terminal period value is given by

$$\begin{aligned} V_{\mathcal{T}}(\Omega) = & \frac{1-\beta^{\mathcal{T}^o - \mathcal{T}+1}}{1-\beta}(S^o + I^o + R^o) v(c^o_{B,\mathcal{T}}, h^o_{B,\mathcal{T}}, c^o_{V,\mathcal{T}}, h^o_{V,\mathcal{T}}) \\ & + \frac{1-\beta^{\mathcal{T}^y - \mathcal{T}+1}}{1-\beta}(S^y + I^y + R^y) v(c^y_{B,\mathcal{T}}, h^y_{B,\mathcal{T}}, c^y_{V,\mathcal{T}}, h^y_{V,\mathcal{T}}) \end{aligned}$$

where the allocations are given by the solution to the static problem without an epidemic.

E.2 Separating the problem into a dynamic and a static dimension and the dynamic externalities

The planner's problem, in a given period, can be separated into a dynamic epidemic problem and, conditional on the epidemic problem, statically computing the optimal economic allocation. The dynamic problem generates the cost of an individual infection (which differs for young and old). Taking these costs into account, the economic problem consists of statically optimizing the economy subject to these costs, given the current number of susceptible, infected, and recovered young and old. Hence, doing this reduces the problem in each period to exactly the static problem we described in Section 3.

Formally, we back out the Lagrange multipliers on the infection rate from the Lagrange multiplier on transmissions ($\widehat{\pi}^B$ and $\widehat{\pi}^V$ denote the risk of becoming infected per time unit spent in the B and the V sectors respectively, see equations (51) and (52)),

$$\partial \widehat{\pi}^B : \qquad \xi^{\pi^B} = \sum_i (h^i_B + n^i_{Bw}) S^i \xi^{T^i}, \tag{53}$$

$$\partial \widehat{\pi}^V : \qquad \xi^{\pi^V} = \sum_i n^i_{Vw} S^i \xi^{T^i}. \tag{54}$$

The Lagrange multiplier on transmissions in turn depends on the Lagrange multipliers on the number of susceptible and infected individuals,

$$\partial T^i : \qquad \xi^{T^i} = \xi^{S^i} - \xi^{I^i}. \tag{55}$$

Finally, the Lagrange multipliers for susceptible and infected individuals depend on the marginal continuation values,

$$\partial S^{i'} : \qquad \beta V_{S^{i'}} = \xi^{S^i}, \tag{56}$$

$$\partial I^{i'} : \qquad \beta V_{I^{i'}} = \xi^{I^i}. \tag{57}$$

The epidemic dimension of the problem in a given time period is therefore reduced to the marginal continuation values $V_{S^{i'}}$ and $V_{I^{i'}}$. The planner's economic problem is to compute the optimal economic allocation, conditional on these marginal values. This economic problem, conditional on $V_{S^{i'}}$ and $V_{I^{i'}}$, is a standard static optimization problem.

The dynamics are described by an application the envelope theorem. We get the following marginal values of susceptible, infected, and recovered individuals (with the notation that N^i_{jh} and N^i_{jw} are the total amount of hours for type i in sector j, e.g., the per-capita hours worked times population size).

First, the derivative of the value function with respect to S^i:

$$\begin{aligned}\frac{\partial V}{\partial S^i} = {} & v(c^i_B, h^i_B, c^i_V, h^i_V) \\ & + \frac{\partial F_B}{\partial N^i_{Bh}} \lambda^{y_B} n^i_{Bh} + \frac{\partial F_B}{\partial N^i_{Bw}} \lambda^{y_B} n^i_{Bw} + \frac{\partial F_V}{\partial N^i_{Vh}} \lambda^{y_V} n^i_{Vh} + \frac{\partial F_V}{\partial N^i_{Vw}} \lambda^{y_V} n^i_{Vw} \\ & - (\lambda^{y_B} c^i_B + \lambda^{y_V} c^i_V) \\ & - (\widehat{\pi}^B (h^i_B + n^i_{Bw}) + \widehat{\pi}^V n^i_{Vw}) \xi^{T^i} + \xi^{S^i} \\ & + \pi^B \frac{\sum I^j (h^j_B + n^j_{Bw})}{(\sum_j (S^j + I^j + R^j)(h^j_B + n^j_{Bw}))^2} (h^i_B + n^i_{Bw}) \xi^{\pi^B} \\ & + \pi^V \frac{\sum_j I^j n^j_{Vw}}{(\sum_j (S^j + I^j + R^j) n^j_{Vw})^2} n^i_{Vw} \xi^{\pi^V}\end{aligned} \tag{58}$$

Second, the derivative of the value function with respect to I^i:

$$\begin{aligned}
\frac{\partial V}{\partial I^i} &= v(c_B^i, h_B^i, c_V^i, h_V^i) \\
&+\frac{\partial F_B}{\partial N_{Bh}^i}\lambda^{y_B} n_{Bh}^i + \frac{\partial F_B}{\partial N_{Bw}^i}\lambda^{y_B} n_{Bw}^i + \frac{\partial F_V}{\partial N_{Vh}^i}\lambda^{y_V} n_{Vh}^i + \frac{\partial F_V}{\partial N_{Vw}^i}\lambda^{y_V} n_{Vw}^i \\
&-(\lambda^{y_B} c_B^i + \lambda^{y_V} c_V^i) \\
&+(1-\pi_d^i-\pi_r)\xi^{I^i} - \sum_j (\pi_d^j)' \xi^{I^j} I_t^j + \pi_r \xi^{R^i} \\
&-\pi^B \left(\frac{(h_B^i + n_{Bw}^i)}{\sum_j (S^j + I^j + R^j)(h_B^j + n_{Bw}^j)} - \frac{\sum_j I^j (h_B^j + n_{Bw}^j)}{(\sum_j (S^j + I^j + R^j)(h_B^j + n_{Bw}^j))^2}(h_B^i + n_{Bw}^i) \right) \xi^{\pi^B} \\
&-\pi^V \left(\frac{n_{Vw}^i}{\sum_j (S^j + I^j + R^j) n_{Vw}^j} - \frac{\sum_j I^j n_{Vw}^j}{(\sum_j (S^j + I^j + R^j) n_{Vw}^j)^2} n_{Vw}^i \right) \xi^{\pi^V}
\end{aligned} \tag{59}$$

Third, the derivative of the value function with respect to R^i:

$$\begin{aligned}
\frac{\partial V}{\partial R^i} &= v(c_B^i, h_B^i, c_V^i, h_V^i) \\
&+\frac{\partial F_B}{\partial N_{Bh}^i}\lambda^{y_B} n_{Bh}^i + \frac{\partial F_B}{\partial N_{Bw}^i}\lambda^{y_B} n_{Bw}^i + \frac{\partial F_V}{\partial N_{Vh}^i}\lambda^{y_V} n_{Vh}^i + \frac{\partial F_V}{\partial N_{Vw}^i}\lambda^{y_V} n_{Vw}^i \\
&-(\lambda^{y_B} c_B^i + \lambda^{y_V} c_V^i) \\
&+\xi^{R_i} \\
&+\pi^B \frac{\sum_j I^j (h_B^j + n_{Bw}^j)}{(\sum_j (S^j + I^j + R^j)(h_B^j + n_{Bw}^j))^2}(h_B^i + n_{Bw}^i)\xi^{\pi^B} \\
&+\pi^V \frac{\sum_j I^j n_{Vw}^j}{(\sum_j (S^j + I^j + R^j) n_{Vw}^j)^2} n_{Vw}^i \xi^{\pi^V}
\end{aligned} \tag{60}$$

The first line in each equation captures the flow utility of the individual. The second and third line capture the economic contribution of the individual (marginal productivity times hours worked), and the economic cost of the individual (the individual's consumption). The subsequent lines in each equation describe the dynamic effect that the individual has on the the evolution of the epidemic. Note that the dependence on the future is summarized by the Lagrange multipliers on the continuation value function, ξ^{S_i}, ξ^{I_i} and ξ^{R_i} and the Lagrange multiplier on transmission, ξ^{T_i}.

There are two dynamic externalities. First, since ξ^{π^B} and ξ^{π^V} are in general non-zero, the behavior of an individual also has an effect on the infection risk for others. Second, since the derivative $(\pi_j^d)'$ is non-zero there is an externality from becoming infected as well. If an individual becomes infected, this increases the death risk for other infected individuals. The social planner takes both these externalities into account.

E.3 Solving the decentralized market equilibria

Rational expectations: The rational-expectations market allocation problem is equivalent to the planner's problem, with two important differences. First, the rational-expectations family planner takes the infection rates as given, i.e., lines (51) and (52) are replaced by $\xi^{\pi^B} = \xi^{\pi^V} = 0$. Second, the death rate π^i_d, is taken as given, in other words, $\partial\pi^i_d/\partial I^k = 0, i, k \in \{y, o\}$ when calculating the marginal value of an infected individual.

Myopic: The myopic market allocation is equivalent to the planner's problem described in Section E.1 with line (46) replaced by $\xi^{T^i} = 0$, capturing that the myopic household does not perceive any cost of new infections.

E.4 Solution algorithm

The decomposition of the planner's problem into a dynamic epidemic problem and a static economic problem suggests a solution algorithm as follows.

Set a final time period $\mathcal{T}$ and either conjecture that the epidemic is over at this point or assume that a cure instantaneously arrives a time $\mathcal{T}$. Given an initial epidemic state in time $t = 0$ for both young and old, $(S^i_0, I^i_0, R^i_0)_{i\in\{y,o\}}$, do the following:

1. Guess on a path of Lagrange multipliers on transmissions $\widehat{\xi^{T^i}_t}$, for $t = 0, \ldots, \mathcal{T}$.
2. Roll forward the epidemic for $t = 0, \ldots, \mathcal{T}$, starting with $t = 0$:
    - In period t, given the epidemic state $(S^i_t, I^i_t, R^i_t)_{i\in\{y,o\}}$ and the guess of the Lagrange multiplier on transmissions $\xi^{T^i}_t$, for $t = 0, \ldots, \mathcal{T}$, compute the optimal economic behavior (time allocations and consumption decisions) by solving the system of equations given by first-order conditions and constraints at time t.
    - Use the epidemic law of motion together with the economic behavior in time t to get the epidemic state at time $t + 1$.
    - Repeat until time $t = \mathcal{T}$.
3. Roll backward the marginal value functions:
    - At time $t = \mathcal{T}$, the epidemic is by definition over. Given the time allocations and consumption choices for time $\mathcal{T}$ found in the previous forward rolling

of the epidemic, calculate the marginal value of an extra individual, $V_{S^i_{\mathcal{T}}} = V_{I^i_{\mathcal{T}}} = V_{R^i_{\mathcal{T}}}$.[30] The per-period marginal value of an individual is not only the individual's utility (including the intrinsic value of life), but also the net contribution of the individual (what it brings into the family in terms of labor net of what it consumes). The total marginal value of an individual in period $\mathcal{T}$ is then the discounted sum of the marginal per-period values, from the final period of the epidemics, $\mathcal{T}$, up until the death of the individual, $\mathcal{T}^y$ and $\mathcal{T}^o$ for the young and the old respectively.

- In period $t < \mathcal{T}$, use the envelope theorem (equations (58) to (60)), given $V_{S^i_{t+1}}$, $V_{I^i_{t+1}}$, and $V_{R^i_{t+1}}$ as well as the economic behavior in time t, to compute $V_{S^i_t}$, $V_{I^i_t}$, and $V_{R^i_t}$. The marginal value of an extra individual now depends on if the individual is susceptible, infected, or recovered.
- Repeat until time $t = 0$.

4. Use $V_{S^i_t}$ and $V_{I^i_t}$ to calculate the resulting ξ_t^{Ti} (for $t = 0, \ldots, \mathcal{T}$) . Use this to update the guess for the path of Lagrange multipliers on transmission $\widehat{\xi_t^{Ti}}$.
5. Go back to step (2) and repeat until convergence.
6. Go back to step (1) and try various initial guesses to ensure that the correct equilibrium is found.[31]

We use this algorithm to solve the planner's problem, and an analogous version of the algorithm also solves the market solution.

[30] In time $\mathcal{T}$, the marginal values of an extra individual $V_{S^i_{\mathcal{T}}}$, $V_{I^i_{\mathcal{T}}}$, and $V_{R^i_{\mathcal{T}}}$ are equal since it does not matter if an individual is susceptible, infected, or recovered if there is no epidemic. We here use the notation $\partial V / \partial S^i = V_{S^i}$ to write it more compactly.

[31] The SIR dynamics and the overloading of the hospital system make the planner's problem non-convex and the problem therefore does not necessarily have a unique equilibrium. Careful evaluation of various qualitatively different initial guesses is necessary to ensure that the correct solution is found. In practice, the non-convexity appears to be modest and has only been an issue in a region of the parameter space where the optimal solution transitions from"full suppression" to "protecting the health-care system".

F Covid results with a high value of a statistical life

Figure 23 shows the outcome for the three different scenarios, assuming the higher value of life. As can be seen, the myopic scenario is (by definition) not affected by changing this assumption. However, both the rational expectations scenario and the social planner scenario display a lower death toll.

In the rational expectation scenario, people are more afraid of getting infected, and therefore they stay at home and avoid contagious activities to a higher extent than in the scenario with a higher value of a statistical life. This slows down the spread of the epidemic, and thus there is less "overshooting". In the scenario with the higher value of a statistical life, the total number of recovered is only approximately 57% (compared to the scenario with the lower value of a statistical life assumption, when it was approximately 63%). A lower number of people who have ever been infected translates into a lower number of people who die from the disease.

In the social planner scenario, the death toll is even lower. The social planner let the number of infected slowly, slowly rise, the economy is actually not in an unconstrained phase even after two years. The old are also staying at home to an even higher extent in this scenario compared to the lower-value-of-life scenario, thus they are even more shielded from the epidemic.

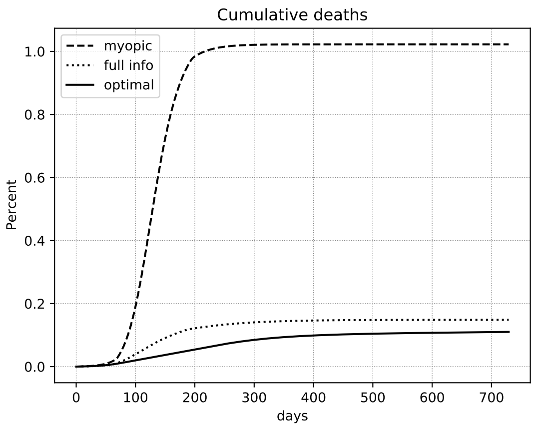

(a) Cumulative deaths under the three different scenarios.

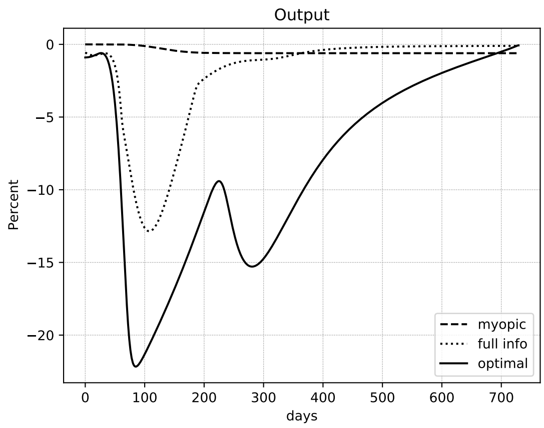

(b) Output drop under the three different scenarios.

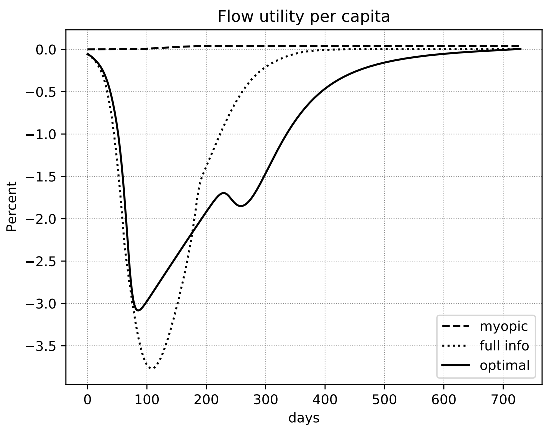

(c) Drop in flow utility under the three different scenarios.

Figure 23: Comparing the three different scenarios, assuming the higher value of a statistical life.

G Seasonal flu

To simulate a "seasonal flu" we set the basic reproduction number, R_0, to 1.3, use a death rate of 0.00045, and a recovery rate of 1/10. This corresponds to a regular normal flu season, and not to a year with a particularly severe instance of the flu, or a year with a pandemic influenza (such as the H1N1/09 virus in 2009).

A systematic review of several published estimates of the basic reproduction number for the seasonal influenza, conducted by Biggerstaff et al. (2014), found that the median estimate of R_0 for the seasonal flu was 1.3, so we use this number directly.

The infection fatality rate for a seasonal flu is more difficult to estimate, and it also varies substantially from year to year. A commonly cited case fatality rate for the seasonal flu is 0.1% (Faust and Del Rio, 2020). According to WHO, the infection fatality rate (i.e., the proportion of deaths among all infected individuals, including all asymptomatic and undiagnosed subjects) is usually well below 0.1%.[32] In order to be conservative, we set the infection fatality rate to less than half of this: 0.045%. An important difference between a regular influenza and covid-19 is also how different groups in the population are affected. According to Petersen et al. (2020), 80% of the deaths from the pandemic influenza in 2009 were below the age of 65, but for a seasonal flu the deaths are more skewed towards the elderly, but not as much as for covid-19. We set an equal death rate for the young and for the old as defined in our model.

According to Petersen et al. (2020), the proportion of infected individuals requiring intensive care is also substantially lower for a pandemic influenza than for covid-19 (less than a sixth). We do not observe any particular overcrowding problems in hospitals during a seasonal flu. To be conservative, we choose to remove the overburdening of the health care system effect, and have the same death rate regardless of the number of infected in the economy.

Further, we assume that on average it takes less time to recover from a flu than from covid-19. For the flu, we use 10 days to recover on average.

To ensure that our estimates of the reproduction number, the infection fatality rate, and estimated time for recovery are plausible, we simulate the myopic model and the rational expectation model to see what the model predicts in terms of deaths. The result is 1.9 (myopic) or 1.8 (rational expectations) deaths per 10,000 people, which rescaled to the full population is 1.5 (myopic) and 1.4 (rational expectations) deaths per 10,000 people.[33]

[32] See https://www.who.int/news-room/q-a-detail/coronavirus-disease-covid-19-similarities-and-differences-with-influenza, downloaded 2020/10/27.

[33] Our model does not include individuals below the age of 15, who constitute 21.4% of the population.

Rolfes et al. (2018) estimate the burden of seasonal influenza in the US and find that the annual deaths due to the seasonal flu in the period between the 2010/2011 season and the 2015/2016 season can have varied between 16,000 and 76,000. Using the average of this low and high number (and approximating the US population to 320 million) gives us a death toll of 1.4 per 10,000 individuals. Another estimate is given by Dushoff et al. (2006), who find an annual average number of deaths in the US from influenza of 41,400 over the period 1979 to 2001. Approximating the US population to 280 million (approximate average during this time period) gives 1.5 deaths per 10,000. A third set of estimates of the total number of deaths per year due to the seasonal flu is provided by CDC.[34] The median of their estimated number of deaths for the period 2010/2011 to 2018/2019 is 1.2 deaths per 10,000 people. Compared to these numbers, our simulated flu is slightly worse than the median flu, but not as severe as for instance the 2014/2015 flu (with 1.6 deaths per 10,000) or the 2017/2018 flu (with 1.9 deaths per 10,000). Hence, our implied estimate of $1.4 - 1.5$ deaths per 10,000 people during a normal flu year seems to be well in line with what is observed .

In sum, the flu we simulate corresponds to a reasonably normal seasonal flu, and not a year with a particularly severe flu. It is also far from a pandemic influenza such as the 2009 case. Note also that we do not take into account the burden of a seasonal flu in terms of people being sick and having to stay at home in bed for days.

Results seasonal flu Figure 24 compare the number of deaths, the output loss and the flow utility loss for a seasonal flu in the myopic market allocation, the rational-expectations market allocation, and the social planner's allocation, assuming the lower value of a statistical life. Figure 26 shows the evolution of infected individuals in the three scenarios, and as the figures show, the evolution is very similar across scenarios.

A social planner would want to decrease output by 1.4 percent during the second quarter of the epidemic, as can be seen in Figure 24b, and the annual drop in output is 0.8 percent. However, this translates into a smaller fall in flow utility, as Figure 24c shows.

Figure 25 shows the corresponding results assuming the higher estimate for the value of a statistical life. With the high value of a statistical life, the social planner would want to lower output by 4.6 percent during the second quarter, and the annual drop in output is 3.4 percent. This, as far as we can tell, is not how actual policy makers have reacted historically. Thus, although we cannot say whether a chosen value of a

We assume for simplicity that there are no deaths in this group.

[34] See https://www.cdc.gov/flu/about/burden/index.html, downloaded 2020/10/21.

statistical life is the correct one, the results from the flu simulations indicate that a value from the lower range is more in line with observed policy actions.

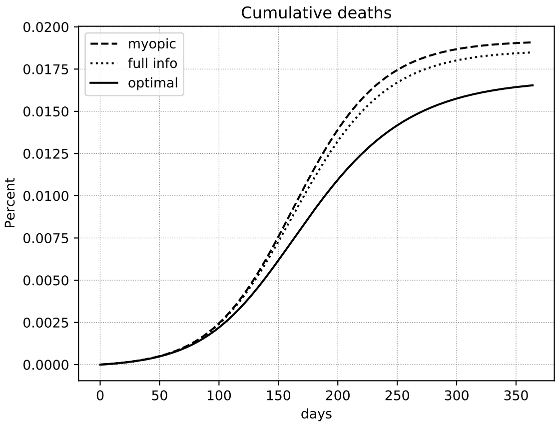

(a) Cumulative deaths under the three different flu scenarios.

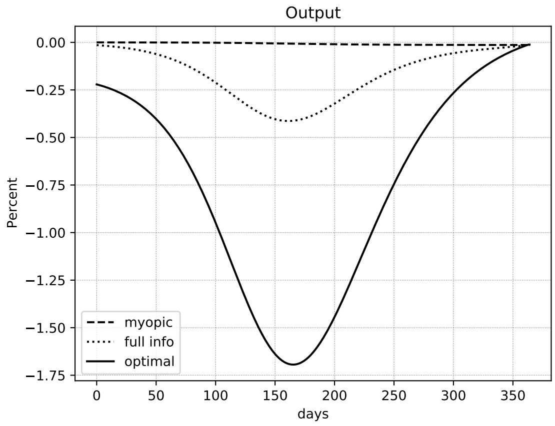

(b) Output drop under the three different flu scenarios.

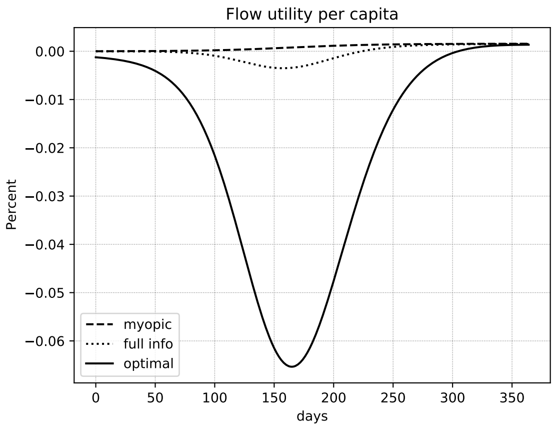

(c) Drop in flow utility under the three different flu scenarios.

Figure 24: Comparing the three different flu scenarios with a *low value of a statistical life* assumption.

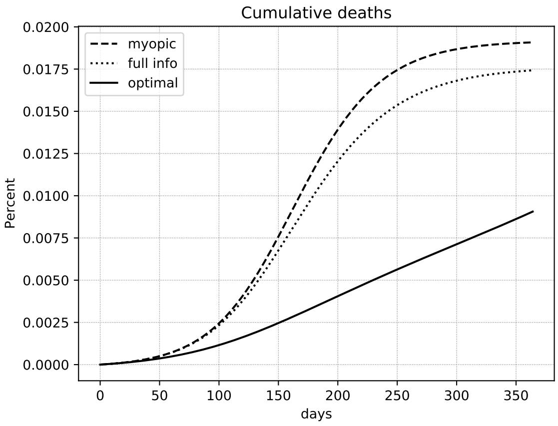

(a) Cumulative deaths under the three different flu scenarios.

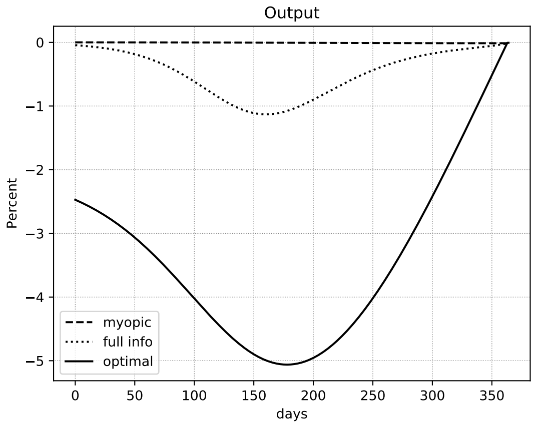

(b) Output drop under the three different flu scenarios.

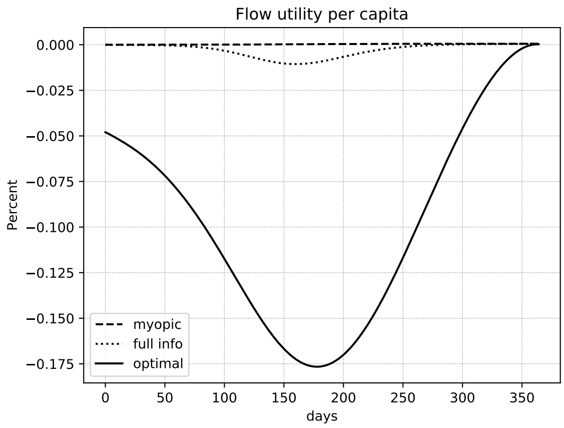

(c) Drop in flow utility under the three different flu scenarios.

Figure 25: Comparing the three different flu scenarios with a *high value of a statistical life* assumption.

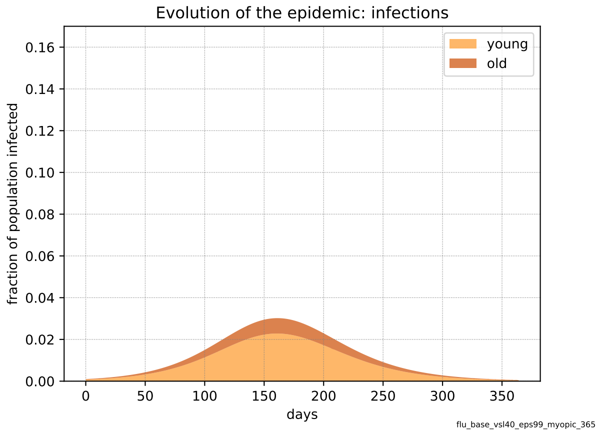

(a) The evolution of a seasonal flu under the myopic market allocation.

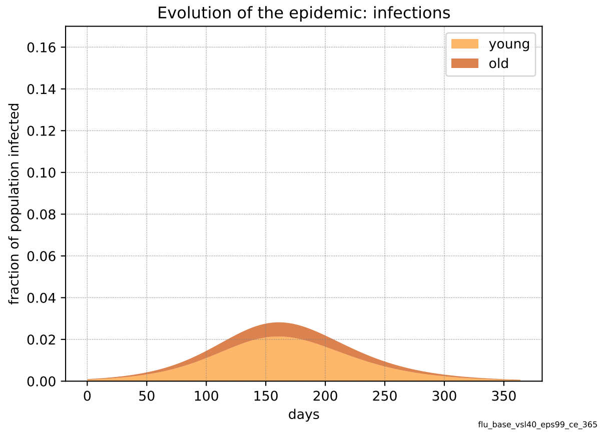

(b) The evolution of a seasonal flu under the rational expectations market allocation.

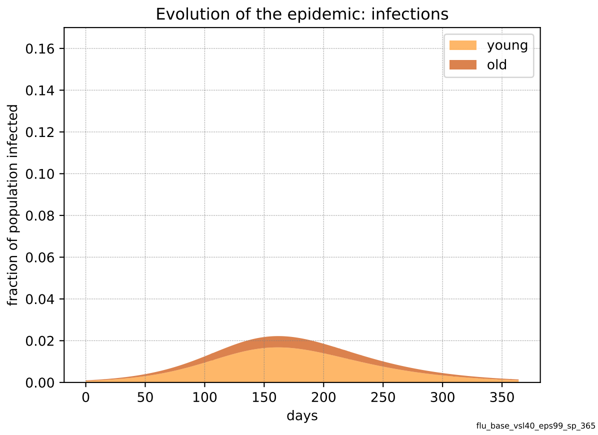

(c) The evolution of a seasonal flu under the social planner's allocation.

Figure 26: The evolution of a seasonal flu under the social planner's allocation and the two different market allocations for the *low value of a statistical life* assumption.

H SARS

To test our model with an epidemic that is substantially worse than covid-19, we use the SARS virus of 2002/2003 (severe acute respiratory syndrome coronavirus, SARS-CoV).

For the transmissability of SARS we use estimates from Petersen et al. (2020) and set R_0 to 2.4. An age-related increase in mortality was observed also for SARS-CoV (although with a far greater case fatality). In Hong Kong, the case fatality due to SARS-CoV was 0% for age group 0-24 years, 6% for those aged 25-44 years, 15% for those aged 45-64 years, and 52% for people who were 65 years and older (Petersen et al., 2020). We set the infection fatality rate to 8% for the group we define as young (15 to 60) and 45% for the group we define as old (above 60).[35]

SARS also has a slightly faster incubation period, so for average number of days until recover we use 12 days. In terms of overcrowding of hospitals, we assume that there is no overcrowding effect in the hospitals that could elevate the IFR even further.

Results SARS Figure 27 shows the evolution of the SARS epidemic in the myopic market allocation, the rational-expectation market allocation, and the social planner's allocation under the assumption that a cure arrives after one year exactly.[36] Again, the epidemic under the myopic market allocation is close to standard SIR dynamics. Many people rapidly get infected. In the case of rational expectation, the behavior is qualitatively different. SARS is dangerous enough to make people so scared of being infected that they stay away from infectious activities voluntarily to a high extent.

A social planner would lower the amount of infectious activities even more. In the case of a social planner, the epidemic is not allowed to take off at all, as Figure 27c shows (we did not forget to plot the curve in this graph!).

As Figure 28 shows, the results from the three scenarios are very different. The death toll is high in the myopic scenario, as expected. In the scenario with rational expectations, the number of deaths is reduced by more than 90%, and the social planner would reduce the number of deaths even more, as shown in Figure 28a.

Figures 28b and 28c highlight the difference in strategy taken by a social planner compared to the rational-expectations equilibrium. A social planner would quickly lower the amount of infectious activities to get the epidemic under control, and would

[35]In the case of SARS, the difference between the case fatality rate and the infection fatality rate is small.

[36]We only show the results for a low value of a statistical life, results from assuming a high value of a statistical life are very similar. Results from assuming that the cure arrives after two years are also very similar, and can be found in the appendix.

thereafter not have to reduce the activities as much. In the rational-expectations scenario people would carry on with their activities until the number of infected is too high in the economy. Then people become afraid of becoming infected, and reduce their activities. Hence, in this scenario, the effective reproduction number is around 1 all the time for the rational expectation equilibrium. The annual drop in output in the rational-expectations scenario is 26.3 percent, which should be compared to the social-planner scenario, in which it is 23.8 percent. Thus, the social planner achieves not only less deaths, but also a smaller drop in output, by setting in "lock-down" measures early, and by doing so getting the epidemic under control at an early stage.

The same type of qualitative effect, that the effective reproduction number hovers around 1 in a rational expectations scenario, is also reported by Farboodi et al. (2020) and Bognanni et al. (2020). The intuition behind the result is as follows. On one hand, the precautionary behavior is increasing in the infection risk, which is increasing in the number of infected. On the other hand, 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.

In our calibration of the covid-19 epidemic, we do not find that the effective reproduction number stabilizes around 1 in the rational-expectations scenario. Including age heterogeneity in the model is important for our result. For the young, the risk of a covid infection does not provide a sufficiently strong motive for a precautionary response to stabilize the infection rate.

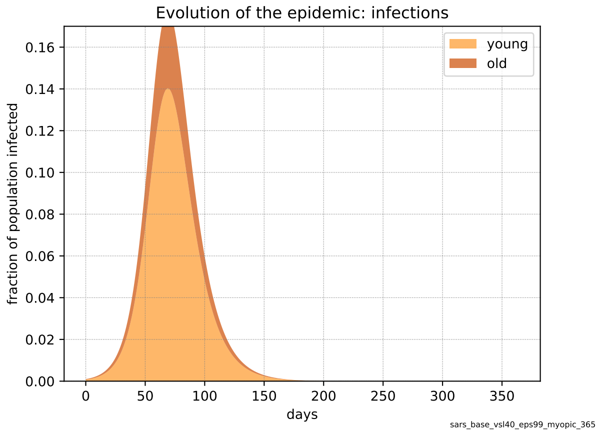

(a) The evolution of a SARS epidemic under the myopic market allocation.

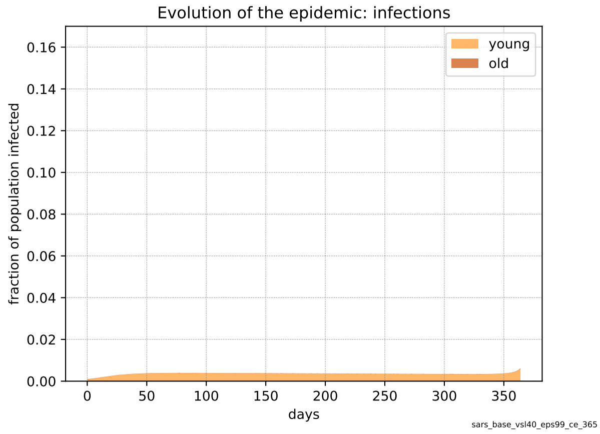

(b) The evolution of a SARS epidemic under the rational expectations market allocation.

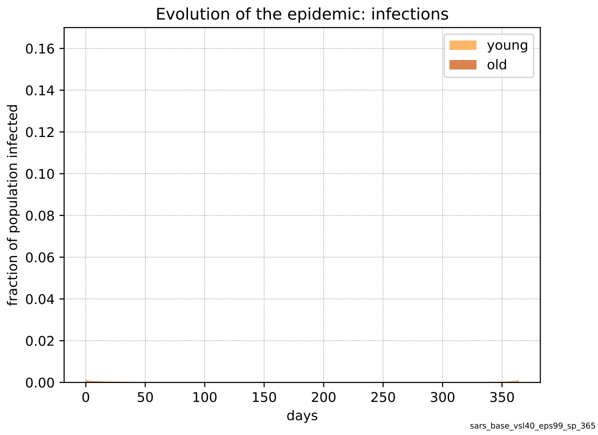

(c) The evolution of a SARS epidemic under the social planner's allocation.

Figure 27: The evolution of a SARS epidemic under the social planner's allocation and the two different market allocations.

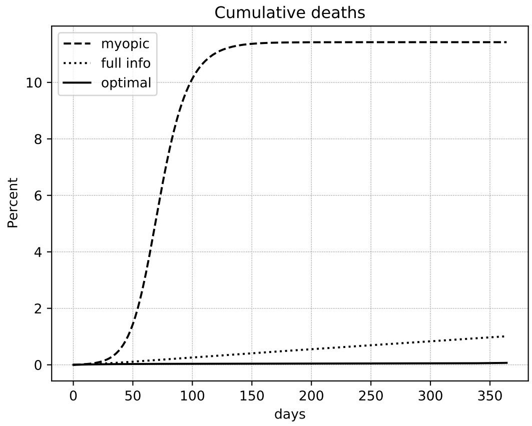

(a) Cumulative deaths under the three different SARS scenarios.

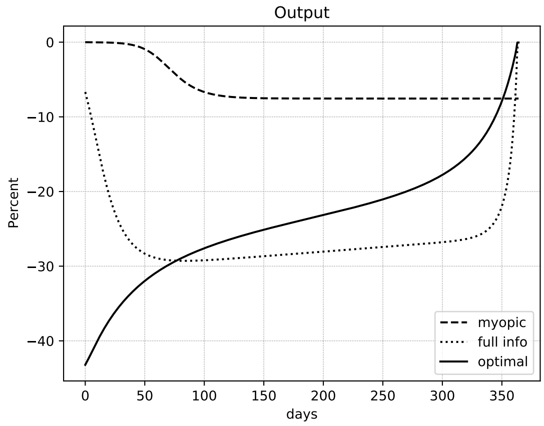

(b) Output drop under the three different SARS scenarios.

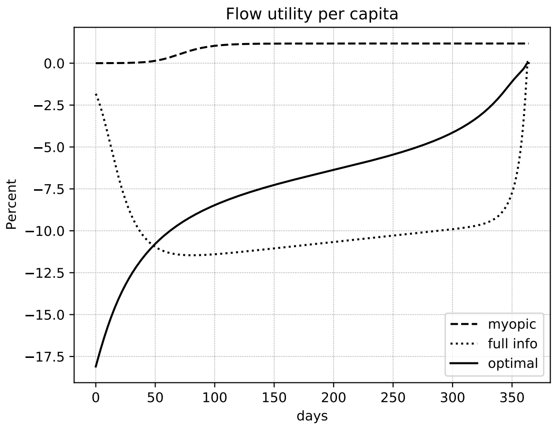

(c) Drop in flow utility under the three different SARS scenarios.

Figure 28: Comparing the three different SARS scenarios.