

An economic model of the Covid-19 pandemic with young and old agents: Behavior, testing and policies*

Luiz Brotherhood Philipp Kircher Cezar Santos
Michèle Tertilt

April 2021

Abstract

This paper investigates the importance of the age composition in the Covid-19 pandemic. We augment a standard SIR epidemiological model with individual choices on work and non-work social distancing. Infected individuals are initially uncertain unless they are tested. We find that older individuals socially distance themselves substantially in equilibrium. An optimal lockdown then confines the young more. The strictness and economic costs of the optimal lockdown depend on whether or not individuals can telework. Testing and quarantines save lives, even if conducted just on the young. When some testing is available, the optimal lockdown is much lighter and GDP rises even compared with a no-policy benchmark.

Keywords: Covid-19, testing, social distancing, age, age-specific policies

JEL codes: E17, C63, D62, I10, I18

*Replication codes are available on the authors' webpages, e.g. here: <http://tertilt.vwl.uni-mannheim.de/research.php>. A previous version was circulated in May 2020 as CEPR DP 14695 under the title "An economic model of the Covid-19 pandemic: The importance of testing and age-specific policies." Financial support from the German Research Foundation (through CRC-TR-224 (project A3) and the Gottfried Wilhelm Leibniz-Prize), the European Research Council (through ERC grant 818859), the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior—Brasil (CAPES)—Finance Code 001 and Faperj is gratefully acknowledged. We thank Suzanne Bellue and Jan Sun for excellent research assistance, Michele Belot for many insights and suggestions, Andrew Clausen, Keith Küster, and Melissa A. Marx for insightful discussions, Daron Acemoglu, Giorgio Topa and many virtual seminar audiences for feedback. The views expressed in this article are those of the authors and do not necessarily represent those of Banco de Portugal or the Eurosystem. Brotherhood: Universitat de Barcelona & BEAT (email: brotherhoodluiz@gmail.com), Kircher: Cornell University, Université catholique de Louvain, University of Edinburgh & CEPR (email: philipp.kircher@gmail.com), Santos: Banco de Portugal, FGV EPGE & CEPR (email: cezarsantos.econ@gmail.com), Tertilt: University of Mannheim & CEPR (email: tertilt@uni-mannheim.de)

1 Introduction

Many diseases affect different parts of the population differently. Some are more prevalent for men, some for women, while others affect children more than adults. Many infectious diseases, especially those caused by viruses, are much more deadly for the elderly than the young—examples include influenza, swine flu, SARS and most recently Covid-19.¹ The strong age gradient in the effects of these diseases is a first-order concern and thus raises a number of important questions: How much voluntary protective behavior do different age groups engage in? How does that interact with policy, e.g., should one target lockdown or testing to particular subgroups? Would it be optimal to confine the elderly to protect them while allowing more freedom to the young who are contributing the most to GDP? While we explore these questions in the context of the current Covid-19 pandemic, the main insights should carry over more generally to future pandemics caused by other viruses.²

To answer these questions, we use a calibrated economic model of the pandemic that features *age heterogeneity* and *individual choice*, allowing agents to choose rationally how much social distancing to undertake, taking into account future infection risk and vaccine arrival. Social distancing provides protection, but comes at the cost of forgone earnings and diminished leisure enjoyment. These can only be partially substituted by teleworking and safe leisure activities. Initial symptoms leave individuals and the government with *incomplete information* whether they have Covid, rendering testing valuable. The deadliness of the disease, the need to earn a living and the natural death probability differ between the working age population and the elderly.

We use the model to quantify the equilibrium interaction between age groups and to explore to which extent age-specific policies are promising. We do this by parametrizing our model to calibration targets taken from the Covid-19 pandemic in the US. Our calibrated model is consistent with data for the disease’s basic reproduction number (R_0), age-specific hospitalization and death rates. The model also matches data on time at work, telework, leisure outside and

¹Viral pneumonia is 10-fold more common among those aged 65+ than among those in working age (Jain et al. 2015). This leads to a corresponding difference in the number of deaths from seasonal influenza between these age groups. Reasons for this age-gradient include the presence of existing lung diseases and other medical conditions, diminished pulmonary reserves, and a waning antigen response (Meyer 2001). This 10-fold increase between age groups also characterizes the case fatality rates of respiratory virus pandemics such as the 2009 “Swine Flu” H1N1 (Wong et al. 2013), SARS (WHO 2003), and most recently Covid-19 (CDC 2020).

²The WHO (2021) reckons that the “next pandemic is most likely to be caused by influenza”, which has a similar age gradient.

time spent at home, all activities that influence infection.

The model predicts that older individuals shield themselves substantially in equilibrium, especially at the peak of the pandemic. The young also reduce work and outside leisure, but much less so due to a lower risk of dying and the need to earn a living. Though the young can telework, this is a lower-productivity activity. Relative to a purely epidemiological model where individuals do not adjust their behavior, this behavioral change decreases the overall death toll by 3/4 and cuts it by more than 80% for the elderly, at a GDP reduction of 9%.

There are two types of externalities in our framework, yielding the need for policy intervention. The first is the typical infectious disease externality: people voluntarily socially distance only to protect themselves, not to protect others. Thus, the equilibrium will feature too little distancing and too many deaths compared to what is optimal. There is another more subtle dynamic externality present in models with infectious diseases and no cure—also called *immunity externality* (Garibaldi, Moen, and Pissarides 2020b). If there is no cure or vaccine, then the epidemic will only stop once herd immunity is reached. This largely determines the death count, and distancing becomes mostly about the timing of infection.³ In the presence of young and old agents, it is better to build herd immunity by infecting the young who have a lower chance of dying (Gollier 2020b). Caution by the young can extend the duration of the pandemic, making it more costly for the old to shield themselves and leading to more deaths. But is such a dynamic externality quantitatively relevant? This is not the case in our calibrated economy with vaccine arrival after a year and a half. If a vaccine arrives very late, then the pandemic will end with herd immunity before vaccine arrival and the dynamic externality becomes relevant. Specifically, we find that if a vaccine arrives only after 15 years, then more caution by the young leads to more deaths among the old. But for plausible vaccine arrival times, we find that herd immunity-based strategies that encourage infections by the young are not optimal.

We use the model to study the effects of a variety of policies and start with lockdowns. An optimal lockdown with utilitarian welfare weights essentially goes for a no-Covid strategy: the death toll declines 99%. This optimal policy thus resembles the no-Covid strategy pursued by some countries (e.g., Australia and New Zealand). The planner confines individuals enough early on to

³As Garibaldi, Moen, and Pissarides (2020b) show, total deaths also vary somewhat as distancing affects the number of still-susceptible people once herd immunity is reached.

bring the disease to a halt. The lockdown is partially, but not fully, relaxed as the vaccine arrival approaches. The economic costs are large: GDP is 8% lower in the first year of the pandemic relative to the no-policy equilibrium (and 18% lower than the no-disease scenario). The planner treats the two age groups differently. It is optimal to confine the young *more*. The young are the ones with little private reason to protect themselves. The old, in contrast, already protect themselves a lot because consequences for their own health are severe. In a world with realistic vaccine arrival, a planner wants the young to distance more to give the old some slack relative to the no-policy benchmark.

If teleworking was not possible, the optimal lockdown would be similar. In fact, the shape of restrictions over time would be the same, but the planner would impose higher costs when individuals go outside. Yet, people would still end up engaging in more activity and the policy would be less effective. Without telework, the optimal policy decreases deaths by only 80% at a much larger GDP cost. These results show that even when teleworking is not possible and economic costs are large, a long and relatively strict confinement is still optimal, but that it no longer resembles a full no-Covid strategy.

Most countries, including the US, have not adopted a no-Covid strategy, and instead implemented lockdowns of varying degrees and lengths. Are milder lockdowns beneficial? Do they cut deaths by less but at lower GDP costs and hence are not too far from the optimum in welfare terms? To answer such questions, we explore stylized lockdowns of different lengths and strictness. The welfare benefits from milder lockdowns are much lower than those from the optimal strict and long lockdown. Some lockdowns even lead to welfare losses relative to the no-policy benchmark. For example, a short 4-week lockdown has little effect on the overall health dynamics of the pandemic. Yet, if strict enough, it can still cause a substantial economic downturn with welfare losses for the young and hardly any benefit for the old. Longer lockdowns are more successful in curtailing deaths. However, strict longer (26-week) lockdowns decrease welfare for both groups, even with fewer deaths. This lower welfare reflects large output and leisure costs relative to modest gains in survival.

A particular lockdown policy that has attracted attention in other research is confining only the elderly. In fact, virtually all existing papers with young and old agents argue for a policy that focuses confinement on the old (Acemoglu et al. 2020; Alon et al. 2020b; Bairoliya and İmrohoroglu 2020; Favero, Ichino, and Rustichini 2020; Gollier 2020b). The logic is that old individuals do not work but have a high risk of dying, suggesting a large benefit and little cost of

their confinement. While seemingly attractive, we find that optimal lockdowns confine old people *less*. Confining only the elderly reduces the welfare of this very group: the old. The reason for this result is intuitive: the elderly still care about restrictions on their leisure time, and what is optimal behavior for one elderly person remains optimal for all elderly together unless their joint behavior affects aggregates such as the overall risk of infection. But if the group is small, the latter effect is also small, and there is little room to improve upon their own choices. Why are our results so different from what the previous literature has argued? The reference point is critical. We compare our policies to an equilibrium with voluntary social distancing, while much of the literature uses the epidemiological equilibrium as reference point.⁴ This distinction is crucial. Like in the previous literature, in our model confinement of the elderly is preferred by all individuals to a purely epidemiological equilibrium where the behavior is fixed as in normal times. But mandatory confinements on top of voluntary social distancing are very costly and only beneficial if applied also to the young. A similar result also arises in simpler calibrations (Brotherhood et al. 2020b) and is confirmed in Boppart et al. (2020), a recent paper with a similar setup to ours.

We also investigate testing and find it very beneficial if carried out in large scale. Testing all individuals who are unsure of their Covid status would reduce deaths by around 50% and lead to a substantial increase in GDP, even without mandated quarantines. The reason is that individuals that are told they have Covid have some partial altruism and reduce their social interactions in our model. This frees up others to take more risk and increase their labor supply. This requires a testing capacity of 4.7% of the population in a week, though. If one additionally quarantines the infected, similar reductions in deaths can be achieved by testing half as many people, requiring a test capacity of 2.3%. If one could test everyone and apply quarantines, this would stop the outbreak on its tracks, yielding hardly any deaths and GDP improvements of 10%, at a testing capacity of 4.8%. The benefits of testing are virtually the same if one only tests the young, as the old are few and careful and hardly affect the dynamics of the disease.⁵ In fact, testing only the old slightly deteriorates the health

⁴Some recent papers in epidemiology embed reduced-form reaction functions to capture adjustments in individual behavior (Manfredi and D’Onofrio 2013). In some economics papers, like Alon et al. (2020b), individuals do make choices, e.g., over their sector of employment, but these are very crude relative to the planner, who can determine the fine details of how much individuals socially distance themselves.

⁵Obviously one might want to test the severely-ill old for Covid to identify them correctly

outcomes. This happens because the old already take a lot of precautions. If unsure whether or not they have the disease, they will continue to take such precautions. However, if told they already caught it, they become fatalistic to a degree that is not offset by partial altruism. While such fatalistic behavior has been previously pointed out by Eichenbaum, Rebelo, and Trabandt (2020b), we find that it is quantitatively relevant only for the old.

We then return to the question of optimal lockdown in a world where we test 50% of those unsure. We find that the optimal lockdown still has a similar shape, i.e., it is long and eases up only close to vaccine arrival. Yet, it is a lot less strict. Recall that the optimal lockdown without testing implies a large recession. With a strict quarantine of those that are tested positive, the planner can relax the confinement of others. Especially the young can now spend about ten additional hours outside each week and they respond by increasing their labor supply. This substantially increases GDP. The old are also allowed more time to enjoy leisure outside. The welfare of both groups sharply increase.

The next section reviews the pre- and post-Covid economic literature on infectious diseases. It makes clear that there is little pre-Covid work on how to calibrate epidemiological models with behavior. We draw on some experiences from Greenwood et al. (2019, 2017, 2013) for the HIV/AIDS epidemic, but there are clearly many differences and novel modelling choices in the current setting. Section 3 outlines the model, followed by the benchmark calibration (Section 4), baseline results (Section 5), lockdowns (Section 6), testing (Section 7), and other policies (Section 8). Section 9 concludes.

2 Literature Review

This paper contributes to the literature that combines epidemiological models (e.g., Kermack and McKendrick (1927)) with equilibrium behavioral choice. In economics, efforts to incorporate behavioral responses to disease progression through equilibrium models have mostly been theoretical. Such works have long pointed out a negative externality of too little prevention efforts by self-interested agents; see, e.g., Kremer (1996) for SI models, Quercioli and Smith (2006) and Chen (2012) for SIR models, and more recently Toxvaerd (2019). These studies consider homogeneous populations, though Kremer (1996) also considers heterogeneous preferences for risky activities. In our setting, differ-

for treatment. We do not model this medical necessity.

ences in activity are partially a consequence of different death rates.⁶

There are few quantitative economic models of disease transmission that predate Covid. Greenwood et al. (2019) develop a heterogeneous-agent choice-theoretic equilibrium model for the HIV/AIDS epidemic to analyze different mitigation policies. Within this framework, Greenwood et al. (2017) explore particular channels of selective mixing by relationship type, while Greenwood et al. (2013) allow for incomplete information in infection status. In these works, the behavioral response of agents is crucial for the results of different policies. Chan, Hamilton, and Papageorge (2016) argue in a structural model that behavioral adjustments matter for the evaluation of medical innovations. Keppo et al. (2020) use a calibrated homogeneous-agent model to argue that a substantial behavioral elasticity is necessary to match different epidemics.

In the great influx of recent economics papers studying different aspects of Covid-19, most consider homogeneous populations. Some analyze optimal containment policies that trade off economic well-being of living individuals versus lost lives (e.g., Alvarez, Argente, and Lippi (2020), Eichenbaum, Rebelo, and Trabandt (2020a), Farboodi, Jarosch, and Shimer (2020), Garibaldi, Moen, and Pissarides (2020a), McAdams (2020)). Other papers introduce uncertainty about one's infection status and the role for testing (Berger et al. (2020), von Thadden (2020), Piguillem and Shi (2020)). Only Eichenbaum, Rebelo, and Trabandt (2020b) combine testing with individual choices about social distancing as in our paper. Unlike our setting, however, infected individuals display no altruism and become more reckless if tested, which renders testing counter-productive unless combined with quarantines.

Our paper differs from these in its focus on a key heterogeneity: different ages and, hence, different risk groups. There are a few other studies that incorporate age differences: Favero, Ichino, and Rustichini (2020) and Gollier (2020a) argue that re-opening should focus on the young while shielding the old, Gollier (2020b) argues that herd immunity has less deaths when built on the young, Glover et al. (2020) analyze how a blanket lockdown affects young and old agents differently and how this leads to disagreement on optimal policy, Alon et al. (2020b) argue that shielding the old while the young work is even more important in developing countries, Acemoglu et al. (2020) characterize the optimal frontier between GDP and lives lost in a model with three age groups and

⁶Other work that considers heterogeneity includes Galeotti and Rogers (2012) who study two identical populations but with non-random mixing, and work on transmission in networks where individuals occupy different positions (Acemoglu, Malekian, and Ozdaglar 2016).

argue that the tension between lives saved vs GDP lost can be best addressed with targeted group-specific policies.⁷ Except for the latter, these papers do not consider uncertainty about the infection status nor testing. They assume that individual behavior can be finely adjusted directly through policy, but is otherwise fixed (or has a coarse dimension such as which sector to work in), and usually trade off lost production vs lives saved. This implies large benefits to confinement of the elderly who are at risk but do not produce. Our work focuses on voluntary behavioral change which has been empirically found to be a large driver behind social distancing (e.g., Maloney and Taskin (2020)) and how this interacts with age-specific policies. In our model the old confine themselves voluntarily and further mandatory confinement lowers their welfare.

3 Model

The general setup of our model is in discrete time. The economy is populated by a continuum of ex-ante identical agents of two types: young y and old o , so that age $a \in \{y, o\}$.⁸ Individuals work, enjoy leisure outside the home and domestic hours. In the presence of the coronavirus, denote the agent's status by j . A susceptible agent is denoted by $j = s$. By spending time outside the house, the agent may catch a disease, which may be Covid-19 or a common cold. Both lead to mild "fever" symptoms.⁹ These agents can be tested for coronavirus with probability $\xi_p(a)$, where the subscript denotes that this is a "policy" choice of the government. With complementary probability, they are not tested and are therefore unsure about the source of their symptoms. Call this fever state $j = f$. A tested individual knows for sure whether they are infected with Covid-19. For simplicity, assume that non-tested infected individuals discover their

⁷Brotherhood et al. (2020a) extend our framework to study heterogeneity on income and housing arrangements in developing countries (including slums). Spatial considerations for the Covid pandemic are also explored in Bognanni et al. (2020). Fernández-Villaverde and Jones (2020) study many countries and cities to analyze possible reopening scenarios.

⁸The model can easily accommodate any number of age groups, but we focus on just two. Acemoglu et al. (2020) highlight large benefits from separately targeting the elderly and the working-age population but little further improvements in sub-dividing those who work. Moreover, we intentionally limit heterogeneity to age to provide a transparent picture, though computationally other dimensions of heterogeneity could be handled.

⁹Ours is one of the few models in the literature that captures partial information—rather than full knowledge or no knowledge about infection status. The specific modelling assumption is for tractability, and evidently does not capture truly asymptomatic cases. Note, however, that many cases that are classified as "asymptomatic" in the data do have light symptoms such as headaches or a running nose. By including all these symptoms in our "fever" state, we can capture the majority of "asymptomatic" cases.

status after one period of uncertainty. If an agent knows they have Covid-19, denote this by $j = i$. Conditional on being infected, they can develop more serious symptoms that require treatment in a hospital, a status denoted by $j = h$. This happens with probability $\alpha(a)$. An agent being treated in a hospital can die with probability $\delta_t(a)$, on top of the natural death probability $\bar{\delta}(a)$ in the absence of the pandemic.¹⁰ The Covid death rate is time-dependent because the death rate for individuals who obtain a bed in an intensive care unit ($\tilde{\delta}_1(a)$) is lower than the death rate for those who do not ($\tilde{\delta}_2(a)$), and intensive care bed shortages depend on the state of the pandemic. The agent can recover from the disease with probability $\phi(j = h, a)$. If the agent recovers, they become immune (or resistant) to future infections, a status denoted by $j = r$.¹¹ Thus, $j \in \{s, f, i, h, r\}$. Agents discount the future at a common factor $\tilde{\beta}$, but since the natural survival probability $\Delta(a) = 1 - \bar{\delta}(a)$ is age-specific, their effective discount factor is $\beta(a) = \tilde{\beta}\Delta(a)$.

For production and leisure, each agent is endowed with one unit of time per period. This can be divided into outside work hours n , teleworking hours v , leisure outside the house ℓ and hours at home d ("domestic" leisure). The agent time constraint is thus:

$$n + v + \ell + d = 1. \quad (1)$$

Agents enjoy utility from consumption c , a composite leisure good when they leave home g , and domestic hours d . The good g is produced using hours ℓ and buying "intermediate" goods x according to $g = g(x, \ell)$. We normalize the utility after death to zero and capture the bliss from being alive through the parameter b . The utility function is given by:

$$u(c, g, d, v; j, a, p) = \ln c + \gamma \ln g + b + \lambda_d \ln d + [\lambda(j) + \lambda_p(j, a)] \ln(d + v).$$

The term $\lambda(j)$ expresses an additional preference for staying at home when being infected, and is a simple way of capturing altruism. We assume two levels: $\lambda(h) = \lambda(i) = \lambda_a$ and $\lambda(r) = \lambda(s) = 0$, so that individuals who transmit the virus are altruistic and the others have no need for that. Individuals in the fever state are unsure whether they are infected or not, and $\lambda(f)$ is a weighted average of these two levels, with weights equal to their belief of being infected with Covid.

¹⁰We do not add a birth process to the model as this is unlikely to be important over the time-frame of the pandemic. The death rate is included to account for differences in effective discount factors between the young and the old.

¹¹While there is uncertainty regarding permanent immunity, immunity in the short run is likely. As long as immunity lasts until a vaccine is available, all our results go through.

We suppress the dependence on the belief to ease notation. $\lambda_p(j, a)$ has a similar role, but from the point of view of the government.¹² This captures simple policies that confine everyone to staying at home ($\lambda_p(j, a) = \bar{\lambda}_p$), but can also capture age-dependent confinements ($\lambda_p(j, a) = \bar{\lambda}_p(a)$), quarantines of those who have tested positive ($\lambda_p(i, a) = \bar{\lambda}_p$, $\lambda_p(j, a) = 0$ for $j \neq i$) and more.

The wage per unit of time worked outside is denoted by w . The wage rate for teleworking hours depends on the amount of telework the agent performs, according to $w\tau(v)$, where $\tau(v) = \tau_0 - \tau_1 v$. The first few activities moved to telework do not carry significant wage penalties. However, as more work is moved to the home, it becomes more costly. This creates a clear trade-off with teleworking: the more one works from home, the less likely it is to catch a disease but the lower is the wage. Old agents receive a fixed retirement income \bar{w} . Total earnings can then be written as follows:

$$w(a, n, v) = \begin{cases} w[n + \tau(v)v] & \text{if } a = y \\ \bar{w} & \text{if } a = o. \end{cases}$$

The budget constraint of the agent is thus given by:

$$c + x = w(a, n, v). \quad (2)$$

Infections happen to susceptible people ($j = s$) when they leave their house. The longer they spend outside, the riskier it gets. Per unit of time outside the house, the transmission risk $\Pi_t(a)$ is time-varying and depends on the number of infected people and how much time these people spend outside, as discussed later. So the probability of getting infected this period is

$$\pi(n + \ell, \Pi_t(a)) = (n + \ell)\Pi_t(a).$$

An agent might also catch a common cold, which happens with probability

$$\pi^*(n + \ell) = (n + \ell)\Pi^*.$$

The probability that the agent catches either disease is

$$\pi_f(n + \ell, \Pi_t(a)) = \pi(n + \ell, \Pi_t(a)) + \pi^*(n + \ell),$$

¹²Variables with an underscore “ p ” are policy instruments that can be used by the government. With a slight abuse of notation, let p_t denote the set of policies the government undertakes at period t .

which implicitly assumes that these are mutually exclusive events.¹³ If this happens and the agent is not tested (probability $1 - \xi_p(a)$), the agent is in the fever state $j = f$ for one period in which they cannot distinguish between the common cold and Covid-19. They assign probability $\Pi_t(a)/(\Pi_t(a) + \Pi^*)$ to having Covid-19. If they are tested (probability $\xi_p(a)$), they will know immediately whether they are infected ($j = i$) or not ($j = s$). Otherwise they will learn at the end of the period whether the fever was due to coronavirus or not.

Finally, assume that a vaccine comes along in period T^* . With the vaccine, nobody gets infected with Covid-19 anymore. That is, $\Pi_t(a) = 0$, for $t \geq T^*$.¹⁴

The value function for susceptible agents is given by:

$$\begin{aligned} V_t(s, a) = & \max_{c, x, n, v, \ell, d} u(c, g(x, \ell), d, v; s, a, p_t) + \\ & \beta(a)[1 - \pi_f(n + \ell, \Pi_t(a)) + \pi^*(n + \ell, \Pi_t(a))\xi_{p_t}(a)]V_{t+1}(s, a) + \\ & \beta(a)\xi_{p_t}(a)\pi(n + \ell, \Pi_t(a))V_{t+1}(i, a) + \\ & \beta(a)(1 - \xi_{p_t}(a))\pi_f(n + \ell, \Pi_t(a))V_{t+1}(f, a) \\ & \text{s.t. (1) and (2).} \end{aligned} \quad (3)$$

The first line captures the utility from consumption and leisure. If the agent has no fever or has fever but tested negative for Covid, they continue as a susceptible person, captured in the second line. The third line captures the continuation for a feverish person who gets tested and had been infected, and the fourth line corresponds to the fever state (fever symptoms and no test).

The value function for an agent who knows that they are infected with coronavirus but do not need hospitalization is given by:

$$\begin{aligned} V_t(i, a) = & \max_{c, x, n, v, \ell, d} u(c, g(x, \ell), d, v; i, a, p_t) + \beta(a)\phi(0, a)V_{t+1}(r, a) + \\ & \beta(a)(1 - \phi(0, a))\alpha(a)V_{t+1}(h, a) + \\ & \beta(a)(1 - \phi(0, a))(1 - \alpha(a))V_{t+1}(i, a) \\ & \text{s.t. (1) and (2).} \end{aligned} \quad (4)$$

The last term in the first line captures the case in which the agent recovers from the disease and becomes resistant to the virus. The second line gives the value

¹³This is a good approximation when the probability of either event is sufficiently small, in which the chance of getting both becomes negligible.

¹⁴An alternative modeling assumption is to suppose a stochastic arrival of a vaccine as in Farboodi, Jarosch, and Shimer (2020).

for the case in which the agent does not recover and requires hospitalization. The third line is the case in which the agent does not recover and does not require hospitalization and, thus, remains infected.

To define the value for an agent in the fever state, it is convenient to denote by $\tilde{V}_t(c, x, n, \ell, d; s, a)$ the terms in lines two to four on the right hand side of (3), and by $\tilde{V}_t(c, x, n, \ell, d; i, a)$ the corresponding terms in (4). They represent the continuation values conditional on choices made this period, both for the susceptible and the infected. Those in the fever state simply get their current utility and the weighted average of these continuation values, weighted by their belief of being infected with Covid:

$$V_t(f, a) = \max_{c, x, n, v, \ell, d} u(c, g(x, \ell), d, v; f, a, p_t) + \frac{\Pi^* \tilde{V}_t(c, x, n, \ell, d; s, a)}{\Pi_{t-1}(a) + \Pi^*} + \frac{\Pi_{t-1}(a) \tilde{V}_t(c, x, n, \ell, d; i, a)}{\Pi_{t-1}(a) + \Pi^*} \quad (5)$$

s.t. (1) and (2).

For individuals in hospital care, we set their flow utility equal to that of death (i.e., to zero) to account for the harsh nature of the disease at this stage. They can enjoy the utility of normal life again if they recover. These agents provide no labor ($n = v = 0$) but we assign an exogenous level of "outside" time ($\ell = \bar{\ell}_h$) to account for the infection burden that they impose onto their carers. The value function for a hospitalized person is:

$$V_t(h, a) = \beta(a) [\phi(1, a) V_{t+1}(r, a) + (1 - \phi(1, a))(1 - \delta_t(a)) V_{t+1}(h, a)] \quad (6)$$

s.t. (1) and (2).

This captures the case of recovery as well as the chance of remaining in the hospital, and the continuation value after dying is set permanently to zero.

Finally, an agent who has already recovered and is resistant to the virus enjoys utility:

$$V_t(r, a) = \max_{c, x, n, v, \ell, h} u(c, g(x, \ell), d, v; r, a, p_t) + \beta(a) V_{t+1}(r, a) \quad (7)$$

s.t. (1) and (2).

To define the laws of motion, denote the measure of agents of each type j of age a in period t by $M_t(j, a)$. Let \mathcal{M}_t be the set of these for all j and a . Further, let $n_t(j, a)$ and $\ell_t(j, a)$ denote their times spent outside the house in equilibrium.

Let \mathcal{N}_t be the set of these equilibrium time allocations in period t for all j and a . The law of motion is a mapping from the state vector and equilibrium actions and the infection rates in period t into the number of agents of each type \mathcal{M}_{t+1} in the next period. Call this map T , so that

$$\mathcal{M}_{t+1} = T(\mathcal{M}_t, \mathcal{N}_t, \Pi_t(o), \Pi_t(y)). \quad (8)$$

It simplifies the accounting to introduce two separate sub-states of the fever state: $j = f_s$ for those with fever who are susceptible (called fever-susceptible) and $j = f_i$ for those who are infected with Covid-19 (called fever-infected). Agents do not know their sub-state, obviously, and therefore act identically in both states. We continue to denote by state $j = f$ all agents who have a fever, which encompasses those in f_i and f_s .

As an example of the laws of motion for this economy, consider the number of susceptible agents next period, which is given by

$$\begin{aligned} M_{t+1}(s, a) &= M_t(s, a)\Delta(a) [1 - \pi_f(n_t(s, a) + \ell_t(s, a), \Pi_t(a)) + \pi^*(n_t(s, a) + \ell_t(s, a), \Pi_t(a))\xi_{p_t}(a)] \\ &+ M_t(f_s, a)\Delta(a) [1 - \pi_f(n_t(f, a) + \ell_t(f, a), \Pi_t(a)) + \pi^*(n_t(f, a) + \ell_t(f, a), \Pi_t(a))\xi_{p_t}(a)] \end{aligned} \quad (9)$$

where the second line captures all situations in which susceptible individuals from last period remain alive and susceptible, as explained in connection to value function (3). The third line resembles the second except that it uses the time allocations of those in the fever state. It accounts for those who entered the period fever-susceptible and continue to remain susceptible during this period. The right hand side of (9) gives the map T_s for the susceptible agents. Appendix A provides the analogous laws of motions T_j for agents in the other states $j = f_s, f_i, f, i, h, r$, and for completeness also for Covid deaths and new infections. The aggregate mapping T is then the vector of the T_j for all states j and ages a .

Aggregation: output, infectiousness, and deaths. Aggregate output in the economy in a given period is given by the time young individuals spend at outside work or telework multiplied by the corresponding wage rate:

$$Q_t = \sum_j w[n_t(j, y) + \tau(v_t(j, y))v_t(j, y)]M_t(j, y). \quad (10)$$

For many of the exercises we aggregate these weekly output measures to get overall GDP measures for longer time periods (e.g., year).

To calculate the aggregate probability of getting infected per fraction of the period spent outside, observe that the number of infected people times their average time spent outside the house times an exogenous susceptible-infected transmission rate Π_0 (which we assume to be age-independent in light of no evidence to the contrary) yields a rate of infection of:

$$\hat{\Pi}_t = \Pi_0 \sum_{\tilde{a}, j \in \{f_i, i, h\}} (n_t(j, \tilde{a}) + \ell_t(j, \tilde{a})) M_t(j, \tilde{a}). \quad (11)$$

This can be rationalized by assuming a unit amount of common space in which agents are distributed uniformly, so that within each subunit of space an individual encounters the number of infected people represented by the sum in (11), and each transmits the virus at the exogenous rate. Expression (11) then corresponds to the probability of getting infected when this number is close to zero, which happens if either the time-weighted number of infected people or the exogenous transmission rate is low. A low number of infected people happens in particular early in a pandemic. To match the basic reproductive number (R_0), calibrations tend to keep Π_0 high, which can push (11) to exceed unity once infection levels peak. This arises because (11) does not take into account the probability of getting infected multiple times within a period, which is irrelevant if this rate is low, but becomes important at the peak. Shortening the period length is counter-productive in our setting as this would also reduce the length of uncertainty of fever individuals in our model.

Therefore we explicitly account for multiple infections within a period in a way that keeps belief updating simple and keeps our infection probabilities in line with other epidemiological models. Assume that the time outside the house represents the probability of entering a common space where one can get infected, and with the complementary probability the person is outside but in a safe space, but the individual does not know which space they are in. If in the common space, interpret (11) as a rate of encountering infections as in continuous-time epidemiological models. Conditional on being in common space, one can integrate to a (weekly) unit of time the probability of having at least one encounter that leads to infection:

$$\Pi_t(a) = 1 - e^{-\hat{\Pi}_t(a)}. \quad (12)$$

When $\hat{\Pi}_t(a)$ is small, this reduces approximately to $\Pi_t(a) \approx \hat{\Pi}_t(a)$. This is constructed under the standard random mixing assumption where everyone meets

everyone else with equal probability.¹⁵

The probability that an agent in need of hospital care ($j = h$) dies of Covid-19, $\delta_t(a)$, can vary over time. This may happen as the recovery probability depends on the supply of hospital resources (e.g., ICU beds) versus the demand for treatment (number of patients). Let Z denote the total number of hospital beds and $M_t(h)$ the total number of agents in need of hospital care in period t . Recall that $\tilde{\delta}_1(a)$ denotes the probability of dying if one is allocated a hospital bed and $\tilde{\delta}_2(a)$ if not. Assume hospital beds are allocated randomly to patients. Hence, the probability of dying $\delta_t(a)$ is given by:

$$\delta_t(a) = \tilde{\delta}_1(a) \min \left\{ \frac{Z}{M_t(h)}, 1 \right\} + \tilde{\delta}_2(a) \max \left\{ \frac{M_t(h) - Z}{M_t(h)}, 0 \right\}. \quad (13)$$

A *rational-expectations equilibrium* in this economy with initial number of agents $M_0(j, a)$ consists of a sequence of infection and death rates $\{\Pi_t(a), \delta_t(a)\}_{t=0}^{\infty}$ and equilibrium time allocations $\{n_t(j, a), v_t(j, a), \ell_t(j, a)\}_{t=0}^{\infty}$ such that these time allocations are part of the solutions to the individual optimization problems (3) to (7), and the resulting law of motion (8) and their aggregation in (12) and (13) indeed give rise to the sequence $\{\Pi_t(a), \delta_t(a)\}_{t=0}^{\infty}$.

4 Calibration

This section describes how we discipline the parameters of the model. Let the time period be one week. Regarding demographics, suppose the old people (who do not work in the model) are those above 65 years old. According to the US Census Bureau, this fraction is 0.214 of the adult population.

We start with the calibration of health-related parameters. These parameters are summarized in Table 1, while Table 3 shows how the model fits the data targets. Start with the common cold. According to Heikkinen and Järvinen (2003), the average American has between two and four colds every year. Suppose an agent in our model has an average of three colds per year. This implies a weekly infection rate of 0.058.¹⁶ In the model, this implies $\Pi^* = 0.113$.

The parameter Π_0 controls how infectious Covid-19 is. We pick Π_0 in order to match the basic reproduction number (R_0) of Covid-19. R_0 represents the

¹⁵While some (e.g., Mossong et al. (2008)) report that contact patterns are assortative with age, others report more uniform contact patterns with considerable interactions across age groups (Belot et al. 2020). In Section 8, we discuss interventions that separate parts of the outside activities by age group, leading to selective mixing and age-specific infection risk.

¹⁶For details, see Appendix B.1

average number of new infections that a random person who gets infected at the start of the pandemic is expected to generate over the course of his disease. In our model, this is closely related to Π_0 .¹⁷ We thus pick this parameter to generate an R_0 of 2.5. This falls within the range of values that Atkeson (2020) uses and close to the range Remuzzi and Remuzzi (2020) report. The corresponding parameter value is then $\Pi_0 = 13.425$.

Once the agent is infected with the coronavirus ($j = i$), there is a probability of recovering from the disease ($\phi(0, a)$), and if not recovered, a probability for developing symptoms ($\alpha(a)$). Set $\alpha(a) = 1$ such that an infected agent spends one week with mild symptoms and recovers (probability $\phi(0, a)$) or develops more serious symptoms (probability $1 - \phi(0, a)$). These numbers are close to what the WHO reports.¹⁸ The parameter $\phi(0, a)$ then also controls the fraction of agents that move to an ICU. CDC (2020) reports age-specific hospitalization (including ICU) rates for Covid-19 patients. A fraction of 3.33% of patients aged 20-64 required being moved to an ICU, whereas this number was 9.1% for those above 65 years. We thus set $\phi(0, y) = 0.967$ and $\phi(0, o) = 0.909$.

We treat hospitalized agents as those that are in an ICU. As discussed above, we assume they cannot work and do not make any decisions. We assume a flow utility level of 0, i.e. the same as death. These individuals still interact with others (e.g., doctors and nurses) for a fraction $\bar{\ell}_h$ of their time. Butler et al. (2018) estimate that patients in ICUs spend about 7.6 hours a day interacting with other people. As these patients are under carefully controlled environments, we assume their infectiousness is half as much as others. Thus, set $\bar{\ell}_h = (7.6/24)/2 = 0.158$.

Agents in need of hospitalization ($j = h$) may die (probability $\delta(a)$) or recover ($\phi(1, a)$). Verity et al. (2020) report that patients with severe symptoms were discharged after an average of 24.7 days, or 3.52 weeks. This yields $\phi(1, a) = 1/3.52 = 0.284$. CDC (2020) also reports age-specific death rates conditional on being hospitalized in the ICU: 14.2% for those aged 20-64 and 65% for those above 65 years old. Given our time period of one week and the recovery rate, this yields weekly death rates of $\delta(y) = 0.065$ and $\delta(o) = 0.738$.

For older agents, Arias and Xu (2019) report an annual survival rate for individuals above 65 years old of 0.95. Thus, set the weekly survival rate for old agents to $\Delta(o) = 0.95^{1/52} = 0.999$. For younger agents, set $\Delta(y) = 1$.

We start the pandemic with a fraction of 0.01% of the population infected

¹⁷For details, see Appendix B.2.

¹⁸See the Report of the WHO-China Joint Mission on Coronavirus Disease 2019 (COVID-19).

Table 1: Calibration – Disease Parameters

Parameter	Value	Interpretation
	0.214	Fraction of old in Population
Π^*	0.113	Weekly infectiousness of common cold / flu
Π_0	13.425	Infectiousness of Covid-19
α	1	Prob(serious symptoms no recovery from mild)
$\phi(0, y)$	0.983	Prob of recovering from mild Covid-19, young
$\phi(0, o)$	0.954	Prob of recovering from mild Covid-19, old
$\phi(1, y)$	0.284	Prob of recovering from serious Covid-19, young
$\phi(1, o)$	0.284	Prob of recovering from serious Covid-19, old
$\bar{\ell}_h$	0.158	Infections through the health care system
$\delta(y)$	0.065	Weekly death rate (among critically ill), young
$\delta(o)$	0.738	Weekly death rate (among critically ill), old
$\Delta(y)$	1	Weekly survival (natural causes), young
$\Delta(o)$	0.999	Weekly survival (natural causes), old
T^*	78	One and a half year (78 weeks) to vaccine arrival

Table 2: Calibration – Economic & Preference Parameters

Parameter	Value	Interpretation
ρ	-1.72	Elasticity of subst. bw leisure time and goods
θ	0.033	Production of leisure goods
γ	0.635	Rel. utility weight - leisure goods
λ_d	1.562	Rel. utility weight - leisure at home
λ_a	1.068	Rel. utility weight - leisure at home (infected)
b	11.0	Flow value of being alive
$\tilde{\beta}$	$0.96^{1/52}$	Discount factor
w	1	Wage per unit of time
τ_0	1.055	Parameter related to telework productivity
τ_1	0.960	Parameter related to telework productivity
\bar{w}	0.214	Retirement income

in $t = 1$. Recall that a vaccine arrives after T^* periods. After this, nobody gets infected with Covid-19 anymore. Set $T^* = 78$, such that a vaccine arrives after one year and a half (78 weeks).¹⁹

Now turn to the preference and economic parameters, see Table 2 for a summary. Let the leisure goods g be produced according to $g(x, \ell) = [\theta x^\rho + (1 - \theta)\ell^\rho]^{1/\rho}$. The parameter ρ controls the elasticity of substitution between leisure time outside the home and leisure goods. Following Kopecky (2011),

¹⁹Several vaccines have been developed by now and large scale vaccination programs are underway. Yet, getting everyone vaccinated will likely take until the summer of 2021. Thus our 1.5 year assumption is quite plausible. We will also explore alternative scenarios later.

set $\rho = -1.72$, which implies an elasticity of 0.368 so that goods x and leisure time ℓ are complements. The parameters γ (utility weight of leisure goods g), λ_d (utility weight of leisure time at home), τ_0 and τ_1 (which control the productivity of telework), and θ are jointly chosen to match five data targets. First, we match a 40-hour work week ($n = 40/112 = 0.357$) in a world without the pandemic. A fraction of 8% of these 40 hours (or 3.2 hours) are spent on telework.²⁰ Moreover, individuals spend 17.3 hours on non-working outside activities ($\ell = 17.3/112 = .154$).²¹ To calibrate the teleworking productivity parameters, we use the following additional piece of information: when the pandemic induces 36% of workers to telework, incomes decline by 10%.²² This yields $\tau_0 = 1.055$ and $\tau_1 = 0.96$. We also match a fraction of income spent on x equal to 12.5% ($x/[w(n + \tau(v)v)] = .125$).²³ Set the discount factor to $\beta = 0.96^{1/52}$.

The parameter $\lambda(i) = \lambda_a$ denotes the increase in the marginal utility of staying at home for agents that know that they are infected with Covid. Government advice is clearly to stay at home, and this parameter captures to which extent time at home actually increases absent further legal enforcement. Epidemiologists have been interested in this during other epidemics such as swine flu (H1N1). In their surveys, the overwhelming majority is willing to comply voluntarily. Rizzo et al. (2013) is the only study that also reports actual adherence, which averaged roughly 50% across Italy, Finland and Romania. We therefore choose λ_a to match an increase in time spent at home by 50%.²⁴

In the model, the parameter b represents the value of being alive over and above the value of consumption. This influences how “afraid” agents are of dying. To discipline this parameter, we target a value of statistical life (VSL) of 9.3 million dollars. This is the value used by the Environmental Protection

²⁰Bick, Blandin, and Mertens (2020) report that around 8% of individuals were working from home in February, 2020.

²¹This comprises the average hours per week spent on purchasing goods and services; caring for and helping nonhousehold members; organizational, civic, and religious activities; socializing and communicating; arts and entertainment (other than sports); sports, exercise and recreation; and travel related to leisure and sports. The data comes from the American Time Use Survey (ATUS). Note that, in our model, the old do not work. In our calibration, the old spend 23 hours in leisure outside; i.e. more than the young. This is consistent with the data; the old spend 19.5 hours in leisure outside. If we include the small reported time at work, this number rises to 25.2 hours per week in the data.

²²According to Bick, Blandin, and Mertens (2020) in May 2020, 35.2% of workers were working from home; similarly, Aum, Lee, and Shin (2020) report a 36% decline in in-loco work in the same period. Further, US GDP declined by 10% in the 2020Q2.

²³This comprises expenditures on food away from home, public transportation, medical services and entertainment. The data comes from the Consumer Expenditure Survey (CEX).

²⁴We analyze the case with no altruism (i.e. $\lambda_a = 0$) in Appendix C.1.

Table 3: Moments – Model vs. Data

Moment	Model	Data (ranges)
Common colds per year	3	2-4
R_0 , Covid-19	2.5	1.6-4
% of infected in critical care, young	3.33	3.33
% of infected in critical care, old	9.10	9.10
% in critical care that dies, young	14.2	5-24
% in critical care that dies, old	65.0	40-73
Weeks in critical care, young	3.5	3-6
Weeks in critical care, old	3.5	3-6
Hours/day interacting while in ICU	3.8	7.6 (controlled)
Life expectancy (natural), young	∞	79
Life expectancy (natural), old	20	20
% \uparrow in time @ home - mild symptoms	50	50 (H1N1)
Value of statistical life (in million US\$)	9.3	9.3
Hours of work per week	40	40
% of weekly hours in telework (normal times)	8	8
% \downarrow in output w/ 36% of workers in telework	10	10
Hours of outside activities per week	17.3	17.3
% of income on goods outside	12.5	11.1-16.1
Replacement rate - social security, %	60	46-64

Agency (see Eichenbaum, Rebelo, and Trabandt (2020a)).²⁵

Normalize the hourly wage to $w = 1$. According to Biggs and Springstead (2008), the replacement rate for social security benefits for a median-income household ranges between 46% and 64%. Set the replacement rate in the model to 60%, a value towards the upper bound of the range since households may have savings outside the official social security income. This implies $\bar{w} = .6w[n + \tau(v)v] = .6 \times .357 = .214$.

For now, assume that $Z = 1$, such that there are enough ICU beds to treat everyone. We will discuss the case with hospital bed constraints in Section 8.

5 Baseline Results

Our baseline is an economy with Covid-19 but no policy interventions. We begin with a word of caution. Data on many relevant dimensions of the model is still scarce and, even if they exist, wide ranges are reported. We know espe-

²⁵To fit the VSL, we consider a young person in a no-Covid world, increase the quarterly probability of death by 1/10,000 and compute b such that we need to give this person 930 dollars per quarter to be indifferent.

Table 4: Non-Targeted Moments – Model vs. Data

Moment	Model	Data (ranges)
Infection fatality rate (IFR), %	0.89	0.5-8.5
Daily growth of infections, outset of Covid-19, %	14.4	15-50
Deaths, old/all, %	49	≈80

cially little about the infection fatality rate because it is unclear what fraction of the population is already infected. We thus did not use some important dimensions as calibration targets. Accordingly, all quantitative results should be interpreted with caution. To give the reader a sense of what our model implies along these important dimensions, we start by reporting moments from our baseline model that were not targeted and compare to empirical estimates, where available, see Table 4.

Most empirical estimates report the case fatality rate (CFR), i.e. the number of deaths per confirmed Covid-19 case, which clearly differs from (and is lower than) the infection fatality rate (IFR) if testing is imperfect. Estimates of CFR vary considerably, ranging from 0.5 in Iceland, 3.5 in Italy and 8.5 in Mexico as of January, 2021.²⁶ We are somewhat on the low side here. However, as more antibody surveys become available, it may turn out that many more people were already infected than is currently believed, which would bring the case fatality rate down.

Second, the daily growth rate at the outset of the pandemic in the model is 14.4%. In the data, these numbers vary greatly. Countries like Australia report an initial daily growth rate of 15%; whereas the counterpart for Spain is around 30%. In countries like Italy and Norway, it reached 50%.²⁷ Note that, though our model generates a number in the lower range, in reality, the initial rise in the data was probably partly due to a ramp-up in testing as well. Third, in the benchmark 49% of all deaths are accounted for by those over 65. Some empirical estimates find a higher number, as high as 80%.²⁸

Table 5 reports our baseline results. Again, this concerns the economy with Covid-19 but no policy intervention (Column Benchmark). It takes about 15 weeks for an unchecked pandemic to reach its peak in terms of seriously ill

²⁶See the University of Oxford's ourworldindata.org. A study from Ischgl finds an IFR as low as 0.24 based on testing 79% of the population for antibodies. See *The Telegraph*, "High coronavirus immunity found in 'super-spreader' Austrian ski resort," June 25, 2020.

²⁷Again, see the University of Oxford's ourworldindata.org website.

²⁸See <https://www.cdc.gov/nchs/nvss/vsrr/covid19/index.htm>

Table 5: Benchmark Results

	Benchmark	Epidem.	Age ext. partial	Age ext. general	No disease
Wks to peak hospitalizations (yng)	15.00	12.00	15.00	17.00	
Wks to peak hospitalizations (old)	11.00	12.00	11.00	12.00	
Hospitalizations p/ 1,000 @ peak (yng)	0.76	12.81	0.17	0.12	
Hospitalizations p/ 1,000 @ peak (old)	0.20	11.11	0.20	0.07	
Dead p/ 1,000 1year (yng)	1.22	4.04	0.32	0.24	
Dead p/ 1,000 1year (old)	4.03	31.40	4.03	1.62	
Dead p/ 1,000 1year (all)	1.82	9.89	1.11	0.54	
Dead p/ 1,000 LR (yng)	1.66	4.04	0.46	0.38	
Dead p/ 1,000 LR (old)	5.79	31.40	5.79	2.47	
Dead p/ 1,000 LR (all)	2.55	9.89	1.60	0.83	
Immune in LR (yng), %	35.12	85.29	9.74	8.10	
Immune in LR (old), %	8.67	45.81	8.67	3.71	
Immune in LR (all), %	29.46	76.84	9.51	7.16	
GDP at peak - rel to BM	1.00	1.13	0.73	0.97	1.14
GDP 1year - rel to BM	1.00	1.09	0.73	0.95	1.10
Cost p/ life saved, million \$	—	—	18.82	2.31	
Hrs @ home (yng) - peak	76.29	57.97	104.50	79.27	57.97
Hrs @ home (old) - peak	104.44	88.99	104.44	94.75	88.99
Hrs @ home (yng) - 6m	74.66	57.97	103.63	79.07	57.97
Hrs @ home (old) - 6m	103.31	88.99	103.31	94.65	88.99

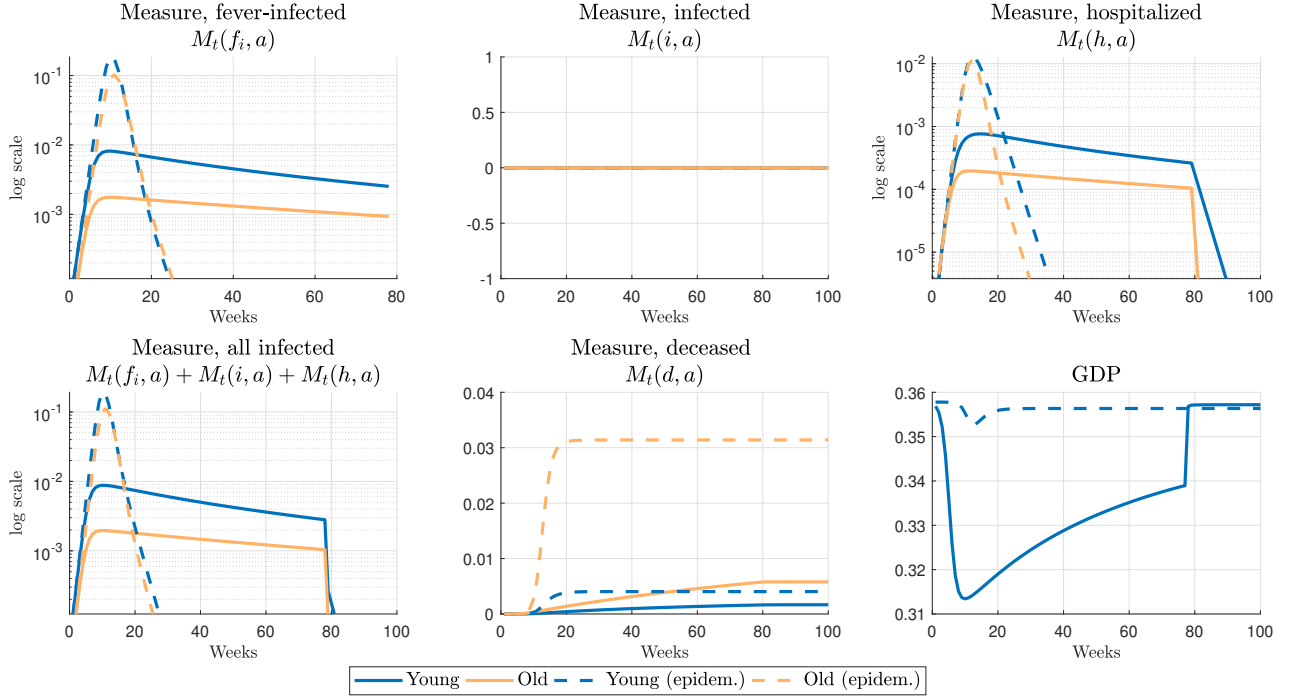
patients. Recall that a vaccine arrives after 78 weeks. As can be seen in Figure 1, at this time, the pandemic ends. The death count is substantial: 2.55 deaths per 1,000 people. This number masks considerable age heterogeneity: the death rate is 5.79 per 1,000 old individuals and 1.66 for the young.²⁹ Most of these deaths happen within the first year of the pandemic, related to the quickness with which it reaches its peak. By the time the vaccine arrives, about 30% of the population had been infected at some point and recovered.

We compare the benchmark results to an epidemiological version of the model (Column Epidem. in Table 5), which assumes there is no change in behavior compared to a world without the disease (last column in Table 5).³⁰ In the benchmark people adjust their behavior substantially, i.e., the individual risk of dying leads people to increase their time at home substantially (see left panel in Figure 2). As the old do not work, they cut leisure outside substantially. The young cut both time at work and leisure outside (again, see Figure 2). This reduces the overall number of infected people in the long run by a little over 45 percentage points. The total death rate declines from 9.89 to 2.55 per 1,000 people, but markedly more so among the old (see the middle lower panel in Figure

²⁹Note that our (realistic) assumption that the old may also die from natural causes does play a role here. Without this assumption, the old would be even more careful leading to a death rate of only 3.5 per 1,000 instead.

³⁰Recall that infected agents are partially altruistic. We explore the behavioral responses in a world without altruism in Appendix C.1.

Figure 1: Aggregate variables (Benchmark equilibrium)




1). The timing of the disease does not change much, but the peak is much lower (see Figure 1). The economic costs of this self-preservation is sizeable: GDP for the first year of the pandemic is reduced by 10% relative to a no-disease world (see Table 5 and Figure 1 for the time path of GDP). The no-pandemic scenario has a GDP 14% higher than the baseline at the peak of the disease. In other words, voluntary reductions in activity reduce GDP substantially.³¹

The young have more incentives to leave their house because of work. They can then contribute more to the spread of the disease; and the burden will fall more heavily on the old. This is a multi-group equivalent of the immunity externality discussed in Garibaldi, Moen, and Pissarides (2020b). To get a sense of the quantitative relevance of this dynamic externality, we run two counterfactuals. Suppose the preferences of the young feature the same death and symptoms probabilities as those of the old (keeping the actual transition rates at their true levels). That is, the young, who still need to work for their income, believe they are subject to the same risks as the old. One counterfactual runs such scenario in a partial equilibrium sense: we observe the difference in behavior of the young assuming they cannot affect the aggregate infection rates. The other

³¹Using a similar mortality rate and a simple utilitarian welfare function, Hall, Jones, and Klenow (2020) estimate people's willingness to pay to avoid Covid-19 deaths to be 18% of annual consumption.

Figure 2: Choices of susceptible agents (Benchmark equilibrium)



counterfactual performs the same thought experiment in general equilibrium. In partial equilibrium (third column in Table 5), the young become substantially more careful and considerably increase their hours at home. This lowers the infections among this group and they consequently die in much lower numbers. In general equilibrium, this extra cautiousness of the young affect the old in two ways. First, by being more careful, less infections take place. On the other hand, if the young are more reckless and face more infections, they end up contributing more to herd immunity and less of a burden falls on the old (Gollier 2020b). In the experiment with our baseline calibration (second to last column in Table 5), when the vaccine arrives for everyone after only 78 weeks, the economy remains far from herd immunity. So, more careful young individuals have a positive effect on the old and the death rate among the latter group declines. However, if the vaccine takes a longer time to arrive, a smaller contribution from the young increases the burden on the old. Suppose the vaccine only arrives after 15 years (Table C2 in Appendix C.2). Then, as the young become more cautious, deaths among the old rise. This effect only materializes at long horizons, and is not present even with vaccine arrival after 10 years (results omitted), and so remains a theoretical possibility rather than a plausible scenario.

6 Shelter-at-Home Policies

Throughout the pandemic, several governments around the world implemented shelter-at-home policies; i.e. lockdowns. This section investigates the impact of such policies on the dynamics of the disease and the economy. We start with an analysis of the optimal policy and then move to analyze stylized policies of

Table 6: Optimal Lockdown

	Benchmark	Optimal policy
Wks to peak hospitalizations (yng)	15.00	79.00
Wks to peak hospitalizations (old)	11.00	79.00
Hospitalizations p/ 1,000 @ peak (yng)	0.76	0.04
Hospitalizations p/ 1,000 @ peak (old)	0.20	0.03
Dead p/ 1,000 1year (yng)	1.22	0.00
Dead p/ 1,000 1year (old)	4.03	0.02
Dead p/ 1,000 1year (all)	1.82	0.01
Dead p/ 1,000 LR (yng)	1.66	0.01
Dead p/ 1,000 LR (old)	5.79	0.08
Dead p/ 1,000 LR (all)	2.55	0.03
Immune in LR (yng), %	35.12	0.30
Immune in LR (old), %	8.67	0.13
Immune in LR (all), %	29.46	0.26
GDP at peak - rel to BM	1.00	1.10
GDP 1year - rel to BM	1.00	0.92
Cost p/ life saved, million \$	–	1.99
Hrs @ home (yng) - peak	76.29	63.76
Hrs @ home (old) - peak	104.44	92.59
Hrs @ home (yng) - 6m	74.66	80.39
Hrs @ home (old) - 6m	103.31	92.87
CEV rel. to BM (yng), %	–	0.91
CEV rel. to BM (old), %	–	6.47

different strictness and lengths.

6.1 Optimal Lockdowns

We start our discussion of lockdowns with optimal shelter-at-home policies. In particular, we solve the problem of a utilitarian planner that maximizes the weighted utility of the old and the young according to their population shares. To do this, the planner chooses how strict the confinement must be. To operationalize such confinements, we assume the planner can choose $\lambda_p(a)$ to alter individual choices. Yet, when evaluating welfare, individual utility with $\lambda_p = 0$ enters the planner's problem. Utility thus captures the impact of distorted choices, without counting the additional effect of the policy instrument per se.

Table 6 reports the results for the optimal policy and compares it with the benchmark. The optimal lockdown essentially goes for a no-Covid strategy. That is, the confinement is strict enough that it essentially eradicates the disease for most of the pandemic until the vaccine arrives. The death toll is 99% lower compared with the no-policy benchmark. As Figure 3 shows, the measure of infected individuals is essentially zero throughout most of the pandemic. Close to the 18-month mark, it starts to increase. But, with the arrival of the vaccine,

Figure 3: Aggregate variables in the optimal policy


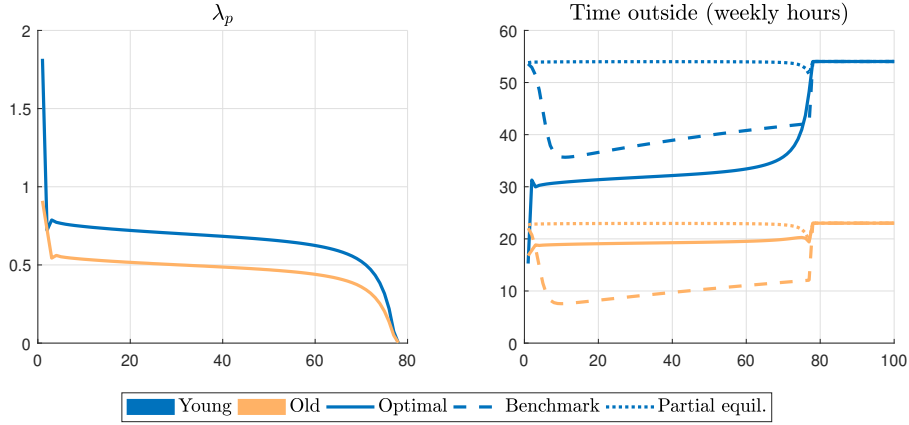


Figure 4: Optimal Lockdown



the disease is soon over. Welfare (measured as consumption equivalent variation, CEV) increases substantially with this optimal policy for both groups, and especially so for the old.³²

How does the planner achieve such an outcome? Figure 4 provides answers. Focus on the young first. In the optimal policy, the young spend less time outside compared with the no-policy benchmark. Part of the extra time at home is spent in productive telework while part of it is devoted to leisure. With less time outside, the young cut both leisure outside and the more productive time at work outside.³³ This reallocation of the time of the young causes GDP to fall below the benchmark (Figure 3). In the first year, output is 8% lower. The lockdown is stricter at the very beginning of the pandemic. With this strictness, the planner can bring the disease under control right away. As the arrival of the vaccine nears, the lockdown is relaxed and the young is allowed more time

³²Appendix B.3 details how we compute welfare .

³³Figure C1 in Appendix C.3 reports time allocation for all activities under optimal lockdown.

outside.

The policy prescription for the old, however, is markedly different. The planner allows the old to spend *more* time outside compared with the benchmark that features a more widespread disease (see Figure 4, right panel). Even though the old do not work, they value leisure, which is taken into account by the planner. This is in contrast with frameworks that only take into account the output produced by the different age groups (as in Acemoglu et al. (2020) for example). The planner places restrictions on the old, though, as otherwise they would free-ride too much on the low infection risk. If the old were allowed to choose without constraints, they would spend even more time outside (see the partial equilibrium choices, represented by dots in the right panel of Figure 4). They do face less restrictions than the young, as shown in the left panel of Figure 4. The old spend less time outside anyhow and are intrinsically more responsive to the disease, which requires less planner intervention.

6.1.1 Optimal Lockdown without Teleworking

The results from the optimal lockdown discussed in the previous section implied an aggressive strategy that virtually eradicated the disease. This drastic outcome is achieved at a substantial economic cost. This cost, however, could have been even larger if teleworking was not possible. With telework, individuals can still produce at home under lockdown (even if at a lower productivity). What would happen then if a planner had to choose the optimal lockdown in a world without telework?

The overall paths of the optimal lockdown policies with and without telework are similar. The young are required to spend much more time at home compared with the benchmark than the old. Yet, the disease is no longer eradicated. The death toll declines 80% but still adds up to 0.74 deaths per 1,000 individuals. The economic costs are also much higher: output in the first year of the pandemic is 33% lower than in the benchmark. Finally, as individuals have a stronger desire to go outside since they cannot produce at home, the planner must raise the utility costs of the confinement (higher λ_p) in order to incentivize people to stay at home. See Appendix C.3 for the details.

6.2 Stylized Lockdowns

The optimal lockdown discussed in the previous section essentially amounts to a no-Covid strategy in which the planner aims at eradicating the disease. Some

countries seemed to have followed a similar approach (e.g., Australia and New Zealand). However, several others did not. To study different strategies, we now turn to lockdowns of various lengths and strictness.

Table 7 has the results for several constellations of the policy. To evaluate the various policies, we always report the welfare gain relative to the benchmark. The last two rows of the table report the welfare gain relative to the epidemiological model. Since the epidemiological model ignores voluntary cautious behavior, the welfare gains, especially for the old, of all policies are extremely large. As we will see, once we evaluate policy relative to the equilibrium, policy gains are much more nuanced.

Start with the mild and short (4-week) lockdown, in which λ_p is chosen such that young agents would spend an extra 25% of time at home in the no-Covid scenario (yielding $\lambda_p = 0.3836$). The effects on the statistics related to the disease are small. The disease slows a bit with the peak being achieved two weeks later. The death rate declines slightly. Given the mild requirement to stay at home and the possibility to telework, GDP hardly changes. The lower death rate and the small effects on output compound to marginally increase the welfare of both groups.

The 4-week strict lockdown increases the shelter-at-home requirement to 90%, yielding $\lambda_p = 32.3598$. The strictness of the policy brings the pandemic down to a point of delaying the peak of the disease by 9 weeks, even though the policy only lasts for 4 weeks. The total death rate declines by 11.4%, with a slightly bigger impact among the old (a fall of 12.1% within this group). Although the policy is in place for a short period of time, its strictness adds to a fall of 2% in yearly GDP, notwithstanding the possibility of telework. This fall in income affects the young. As they are now poorer and face a tough shelter-at-home requirement, their welfare declines. The old, on the other hand, experience a small rise in their welfare since their retirement income is guaranteed. This type of disagreement may explain the controversies surrounding the implementation of lockdowns.

Now consider longer policies that last for 26 weeks. The mild 26-week lockdown decreases the total death rate by about 16%, with a higher decline among the young. Though the policy is mild, its long duration leads to a fall in GDP of 3% compared with the benchmark. As both groups experience better health outcomes and the fall in income among the young is not too deep, the welfare for both the young and the old rise. A strict 26-week lockdown has more profound consequences. This policy can flatten and delay the infection curve

Table 7: Policy Experiments: Shelter at Home

	Benchmark	4-week lockdown		26-week lockdown		26-week, only old	
		Mild	Strict	Mild	Strict	Mild	Strict
Wks to peak hospitalizations (yng)	15.00	17.00	24.00	37.00	56.00	15.00	15.00
Wks to peak hospitalizations (old)	11.00	14.00	21.00	34.00	52.00	12.00	31.00
Hospitalizations p/ 1,000 @ peak (yng)	0.76	0.75	0.73	0.63	0.61	0.76	0.75
Hospitalizations p/ 1,000 @ peak (old)	0.20	0.19	0.19	0.17	0.17	0.19	0.17
Dead p/ 1,000 1year (yng)	1.22	1.16	0.99	0.80	0.09	1.22	1.21
Dead p/ 1,000 1year (old)	4.03	3.83	3.22	3.01	0.42	3.92	2.30
Dead p/ 1,000 1year (all)	1.82	1.73	1.46	1.27	0.16	1.79	1.44
Dead p/ 1,000 LR (yng)	1.66	1.62	1.49	1.36	0.83	1.66	1.66
Dead p/ 1,000 LR (old)	5.79	5.62	5.09	4.95	2.65	5.68	4.12
Dead p/ 1,000 LR (all)	2.55	2.48	2.26	2.13	1.22	2.52	2.18
Immune in LR (yng), %	35.12	34.22	31.53	28.68	17.46	35.10	34.96
Immune in LR (old), %	8.67	8.43	7.66	7.46	4.06	8.51	6.23
Immune in LR (all), %	29.46	28.70	26.42	24.14	14.59	29.41	28.81
GDP at peak - rel to BM	1.00	1.00	1.00	1.01	1.00	1.00	1.00
GDP 1year - rel to BM	1.00	1.00	0.98	0.97	0.88	1.00	1.00
Cost p/ life saved, million \$	–	3.51	3.19	3.40	4.50	-0.15	-0.09
Hrs @ home (yng) - peak	76.29	76.22	76.30	75.79	76.80	76.26	75.95
Hrs @ home (old) - peak	104.44	104.38	104.26	103.72	103.64	104.78	110.37
Hrs @ home (yng) - 6m	74.66	75.01	76.00	77.63	110.13	74.63	74.33
Hrs @ home (old) - 6m	103.31	103.47	103.91	99.52	110.10	103.74	110.32
CEV rel. to BM (yng), %	–	0.02	-0.23	0.16	-1.66	0.00	0.01
CEV rel. to BM (old), %	–	0.19	0.09	1.04	-0.83	-0.00	-2.44
CEV rel. to epidem. (yng), %	1.67	1.70	1.44	1.84	-0.02	1.68	1.68
CEV rel. to epidem. (old), %	20.11	20.34	20.22	21.39	19.09	20.10	17.10

Note: A mild lockdown is given by the λ_p that makes young individuals spend an extra 25% of time at home in the no-Covid scenario, while a strict lockdown raises this requirement to 90%.

(see Figure 5). This delay and the arrival of a vaccine after 78 weeks combine to decrease the death toll by more than 50%. This result comes at a substantial economic cost as GDP declines by 12% in the first year of the pandemic. The time path of output actually exhibits two downturns. The first downturn is caused by the lockdown. The second downturn is due to the endogenous response of individuals. As the disease starts to spread, people cut their time outside and this cut includes a decrease in the young's time at the more productive outside work (see Figure 6). In the end, though the young experience better health outcomes, this decline in their income reduces their welfare. The retirement income of the old does not change with the lockdown and they also face lower death rates. However, the strict lockdown reduces their leisure outside so much that the policy leads to a decline in their welfare as well.

Some of the policies discussed so far are able to decrease fatalities in the pandemic. However, this outcome comes with large drops in GDP. To avoid these large economic costs, one possibility is to implement group-specific policies. The old can be confined since they are the more at-risk group while the produc-

Figure 5: Aggregate variables (Shelter in place, strict, 26 weeks)

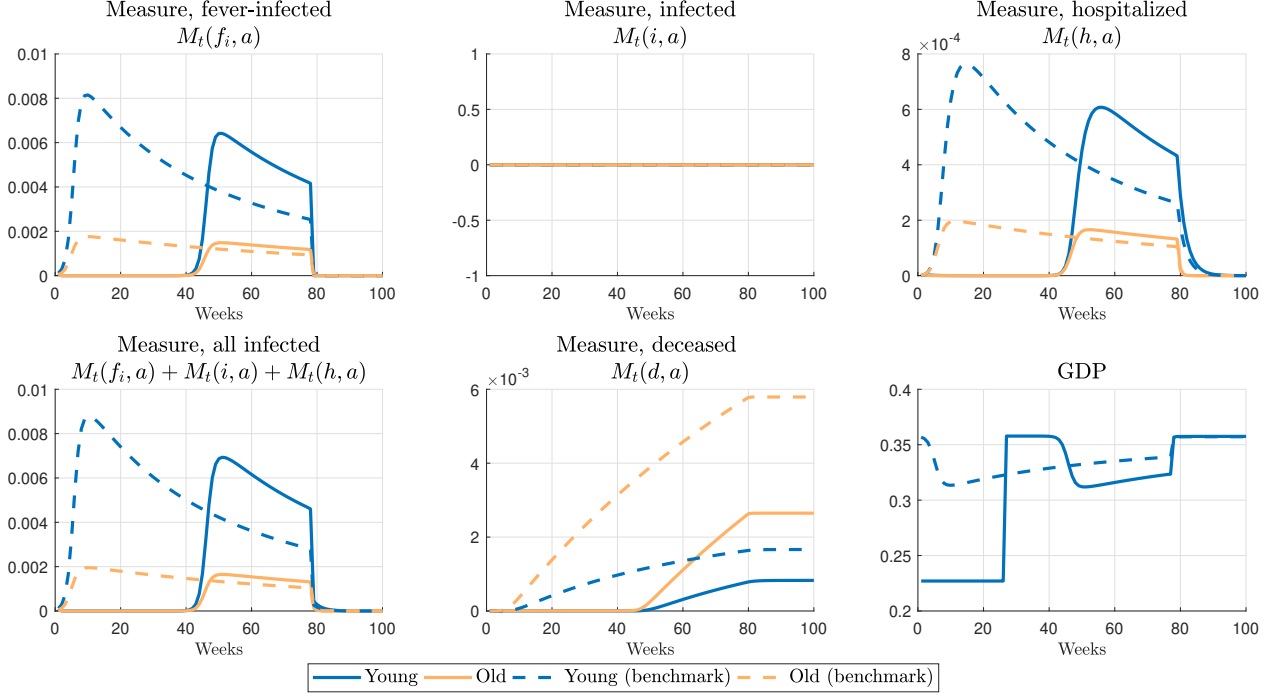



Figure 6: Choices of healthy agents (Shelter in place, strict, 26 weeks)



tive young are allowed to work in order to contain the economic damages. This type of age-specific policy has been suggested both in public discourse and in academic works; e.g., Acemoglu et al. (2020), Alon et al. (2020b), Bairolia and Imrohoroglu (2020), Favero, Ichino, and Rustichini (2020), and Gollier (2020a).

The last two columns in Table 7 report the results for two 26-week lockdowns (one mild and one strict) that only confine the old. A first observation is that the health of the young is hardly affected by these lockdowns: the death toll and the fraction of immune within this group are virtually the same as in the benchmark. Since the old comprise a relatively small group and they already take significant precautions in the benchmark, the extra time they spend

at home during the lockdown does not alter the situation much for the young.³⁴ Second, as the young are not confined, they are free to work outside. As a consequence, GDP is essentially the same as in the no-policy baseline. The effects on the old are more pronounced. Take the strict policy (last column in Table 7). The death rate in this scenario is almost 30% lower among the old compared with the benchmark. The decline is larger than 40% in the first year. Moreover, as the old are retired, their income is not affected by the lockdown. However, the welfare of the old falls substantially—by 2.4% in consumption equivalent units. This happens since the old are required to stay home for longer than they would have chosen to.³⁵ The reference point for welfare comparisons is crucial. Most papers in the literature compare the welfare of the old under the policy with the epidemiological world with no behavioral adjustments (what we do in the last two lines in Table 7). This comparison suggests that the old would prefer to be confined. However, with endogenous behavioral change in equilibrium, the old socially distance a lot. Since the old are a small group, especially when weighted by their low activity levels, they do not induce large equilibrium effects, and so their individually optimal behavior is close to the collectively optimal behavior of the old as a group. Therefore, substantial further mandatory lockdowns of just this group decrease the welfare of precisely the group it is intended to protect.

All shelter-at-home policies discussed so far were implemented at the very start of the outbreak. What happens if such a policy is implemented only after the disease is developing? Consider a strict lockdown that only starts 8 weeks after the outbreak and lasts for 26 weeks. The results are in Figure C3 in Appendix C.4. As the disease develops in the first 8 weeks, a wave of infections develops. With this policy, infections decline substantially. When the policy is lifted after 26 weeks, a second wave appears. GDP declines at the outbreak due to individuals' endogenous social distancing. It falls further due to the lockdown and it recovers when the policy is lifted. When the second wave develops, the endogenous behavior of agents again leads to a decline in output.

In sum, shelter-at-home policies can decrease the death toll of the pandemic if they last long enough until close to the arrival of a vaccine. Stricter lockdown

³⁴The elderly are a small group precisely because they are so careful: they are 21.4% of population, but only 10.39% of steady-state interaction because they do not need to work, and only 5.45% of social interactions at the peak of the pandemic because of their additional precautions.

³⁵Belot et al. (2020) find that individuals 65 or older experience worse negative non-financial effects from lockdowns. Andersson et al. (2020) report that the old in Sweden are willing to pay more to avoid confinement.

policies are particularly powerful to curb deaths, but they cause substantial economic costs. A utilitarian planner, according to our calibration, would attempt to follow a no-Covid strategy and aim for the virtual eradication of the disease if individuals are able to telework. In this optimal policy, the planner would do so by confining the young more than the old since the old have fewer contacts anyhow and are intrinsically more responsive to the pandemic.

7 Testing

Health authorities around the world emphasize the importance of testing. Accordingly, many countries invested heavily to ramp up their testing capacities since the start of the pandemic. But what exactly is the benefit of testing? Does it work by itself or only in combination with quarantines? In the limit, if everyone was tested every day and, if found Covid-19 positive, immediately quarantined, the pandemic would stop almost immediately. Yet, such a policy is clearly not very realistic. So what is the benefit of partial testing? What about imperfect quarantines? This section considers various combinations of testing and quarantining policies to assess their effect on deaths and GDP.

We start with policies that only test individuals without further quarantines. In our baseline calibration, upon infection with the coronavirus, individuals remain one week in the “fever” state, unsure whether they have Covid-19 or a common cold. By testing them, this uncertainty disappears earlier. Since the agents are partially altruistic, they are more cautious about leaving their homes.

Table 8 provides the results of testing all individuals (column Testing all), only the young (Testing young) or only the old (Testing old). Start with the universal policy. Figure 7 shows that, with universal testing, the mass of infected agents that are unsure whether they have the disease (M_{fi}) goes to zero. Instead, agents now know they are infected (M_i) and thus act according to their partial altruism. This leads the disease to develop at a slower pace; i.e. flattening the curve. The peak now takes about three months longer to arrive and is less pronounced (see Figure 7). This translates into fewer deaths for both age groups. Fewer people catch the disease overall as can be seen from the lower immunity rates in the long run. Note that this universal testing policy is a massive undertaking: in the week with the maximum number of tests performed, about 4.7% of the population is tested. In the US, this implies about 15 million tests in a single week. This is accompanied by a 5% rise in GDP compared with

Table 8: Policy Experiments: Testing

	Benchmark	Testing all	Testing young	Testing old	Testing and quarantine		
					50% testing	100% testing	100% testing only young
Wks to peak hospitalizations (yng)	15.00	28.00	27.00	15.00	25.00	3.00	3.00
Wks to peak hospitalizations (old)	11.00	25.00	25.00	11.00	22.00	3.00	3.00
Hospitalizations p/ 1,000 @ peak (yng)	0.76	0.36	0.35	0.77	0.26	0.01	0.01
Hospitalizations p/ 1,000 @ peak (old)	0.20	0.12	0.12	0.20	0.09	0.01	0.01
Dead p/ 1,000 1year (yng)	1.22	0.58	0.57	1.23	0.46	0.00	0.00
Dead p/ 1,000 1year (old)	4.03	2.33	2.32	4.06	1.97	0.01	0.01
Dead p/ 1,000 1year (all)	1.82	0.95	0.94	1.83	0.78	0.00	0.01
Dead p/ 1,000 LR (yng)	1.66	0.84	0.84	1.68	0.69	0.00	0.00
Dead p/ 1,000 LR (old)	5.79	3.40	3.39	5.82	2.94	0.01	0.01
Dead p/ 1,000 LR (all)	2.55	1.39	1.38	2.57	1.17	0.00	0.01
Immune in LR (yng), %	35.12	17.77	17.64	35.42	14.64	0.05	0.06
Immune in LR (old), %	8.67	5.11	5.09	8.72	4.42	0.02	0.02
Immune in LR (all), %	29.46	15.06	14.95	29.71	12.45	0.04	0.05
Max. n. of tests in a week, %	0.00	4.72	4.24	0.48	2.36	4.76	4.27
GDP at peak - rel to BM	1.00	1.07	1.08	1.00	1.09	1.14	1.14
GDP 1year - rel to BM	1.00	1.05	1.05	1.00	1.06	1.10	1.10
Cost p/ life saved, million \$	–	-2.47	-2.47	–	-2.40	-2.14	-2.14
GDP gain per test, 1 year, \$	–	1431.03	1575.79	–	3286.91	2282.61	2540.06
Hrs @ home (yng) - peak	76.29	67.58	67.32	76.53	64.82	58.72	58.72
Hrs @ home (old) - peak	104.44	99.18	99.00	104.53	96.99	90.09	90.08
Hrs @ home (yng) - 6m	74.66	67.61	67.31	74.86	64.76	57.97	57.97
Hrs @ home (old) - 6m	103.31	99.17	98.96	103.40	96.90	88.99	88.99
CEV rel. to BM (yng), %	–	0.76	0.77	-0.01	0.88	1.43	1.43
CEV rel. to BM (old), %	–	3.33	3.35	-0.06	3.89	6.66	6.66

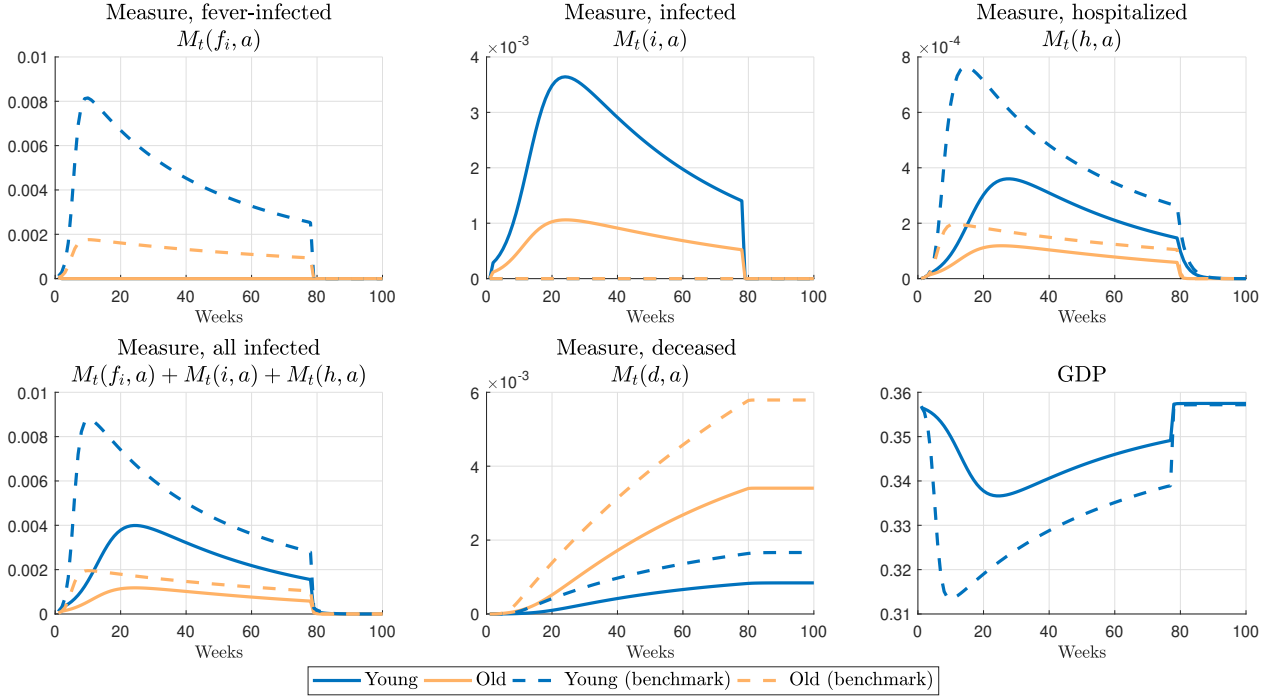
Note: the quarantine intensity is given by the $\lambda_p(i, a)$ that makes young individuals spend an extra 90% of time at home in the no-Covid scenario ($\lambda_p(i, a) = 32.3598$).

the benchmark without testing due to a milder pandemic. In the end of the first year, this translates to about 1,431 dollars in extra GDP per test performed.

Table 8 also reports age-specific testing policies. First, testing exclusively the old causes only minor changes to the development of the disease. Since this group comprises a minority of the population and they protect themselves more, their behavioral change is not strong enough to generate considerable changes in the economy. In fact, a slightly negative effect materializes. As the old get tested and are found positive, they decrease the substantial prevention effort they engage in during the no-policy benchmark. That is, they become fatalistic and undertake more activities, a result reminiscent of Eichenbaum, Rebelo, and Trabandt (2020b). This has a small detrimental effect on others, both via a higher death toll and lower welfare. Testing the young has the opposite effect. As this group consists of 80% of the population and they protect themselves less, the results are quite similar to those obtained with a universal policy.

While testing works quite well, and testing with quarantines even better (as

Figure 7: Aggregate variables (Testing all)



we will see below), implementing this policy faces some hurdles.³⁶ First, in the model we test a fraction of those with fever. In reality, one needs to know who to test. One possibility is to focus on contact tracing. A second issue is that our model assumes instantaneous test results, while in reality this is not the case. However, rapid tests that deliver results within an hour exist by now. Finally, unlike in our model, in reality testing involves costs. Scaling up the current test methods to the quantities needed is not easy. Even if scaling up testing comes at a substantial cost, these costs are likely worth paying, especially since our analysis has shown that there are substantial gains in GDP per test performed, which are more than enough to cover the costs.

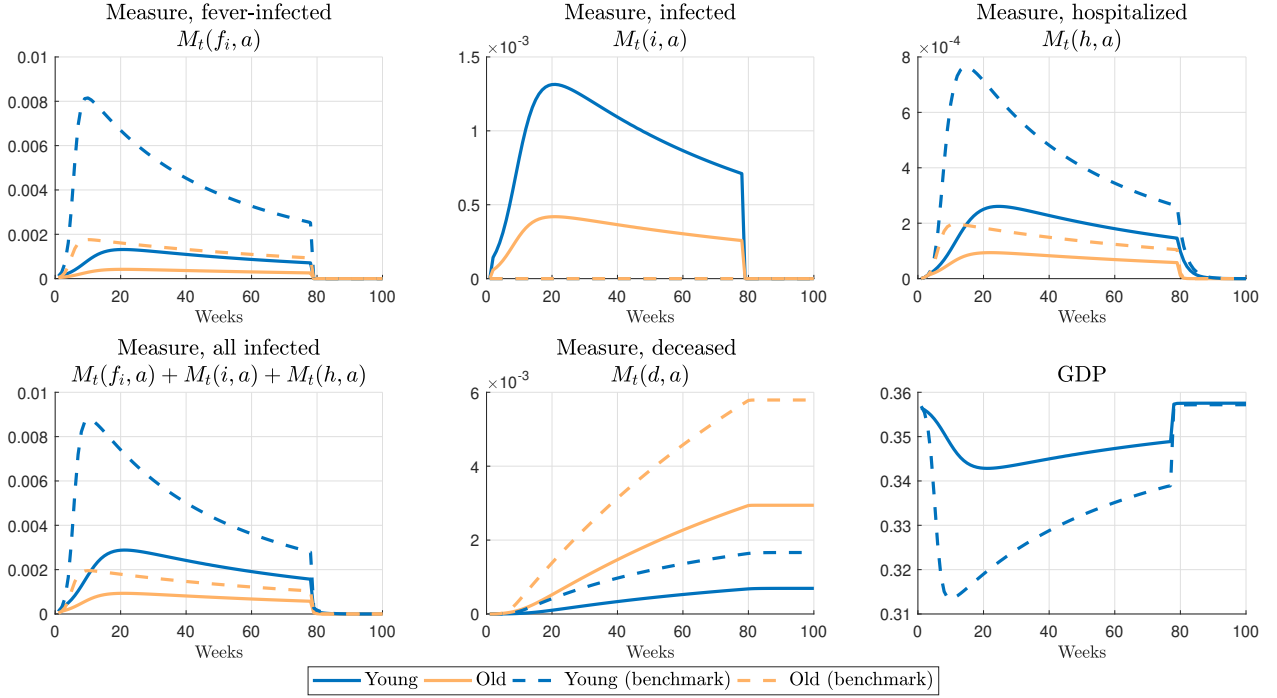
7.1 Test and Quarantine

The previous section discussed the impact of testing policies that rely on individual's altruism to curb the spread of the disease. On top of testing people, one could quarantine those that tested positive (which most countries try to do). The last three columns of Table 8 report results of three such experiments.

Consider a policy where 100% of those with a fever are tested, and those that test positive are forced to quarantine through an increase in $\lambda_p(i, a)$ that makes

³⁶See The Economist, April 25, 2020 "Test of Reason" for difficulties of scaling up testing.

Figure 8: Aggregate variables (Test and Quarantine, 50%)



young individuals spend an extra 90% of time at home in the no-Covid scenario ($\lambda_p(i, a) = 32.3598$). This policy leads the young to spend 92% of their time at home (the old even more). So while it is a relatively strict quarantine, it is not 100%. This captures the fact that enforcement of quarantines is never perfect in reality. It can be thought of as people still escaping from their quarantine 8% of the time (about 8h/week), or that 8% of agents are not complying at all. This policy is extremely effective. In the long run, less than a tenth of one percent of the population catches the disease and are thus immune. The death count is almost zero. With the lower burden of the disease, GDP goes up by 10% in the first year of the pandemic. Given that the number of immune individuals is low in the long run, this may be a drawback of this policy: in the event of a new outbreak, the economy would be far from the herd immunity level.

Consider a test-and-quarantine policy that focuses on the young only (last column in Table 8). Since the young comprise the lion's share of the population and spend more time outside, testing this group only is essentially as effective as testing the entire population. In fact, if one factors in that tests cost some money, then this is the better policy since fewer people are tested, which increases the GDP gain per test to more than \$2,500. Alternatively, a testing policy that applied to the old only (not shown) would not be as useful.

In our model, we assume immediate test results. This is not true in reality, as

the results usually take a few days. Focus then on a test-and-quarantine policy that targets 50% of the population (see Table 8). Alternatively, this could be thought of as testing with results delayed to the middle of the week of infection. With only half of those with a fever being tested, some agents will move from the unsure "fever" state (M_{fi}) to the infected state (M_i), but not all. See Figure 8. Compared with the benchmark, the peak occurs at a later date and is lower; i.e. the curve is flattened (see Figure 8). There will thus be a much lower fraction of hospitalized individuals (M_h) at the peak. Relative to the baseline, the number of deaths declines by more than 50% to 1.17 per 1,000 people. This decline is steeper among younger individuals (a drop of about 58% for this group). The total number of Covid-19 cases in the long run falls from almost 30% of the population in the benchmark to around 12%. GDP goes up in the first year by about 6% relative to the baseline. This represents around 3,286 extra dollars in GDP per test carried out.

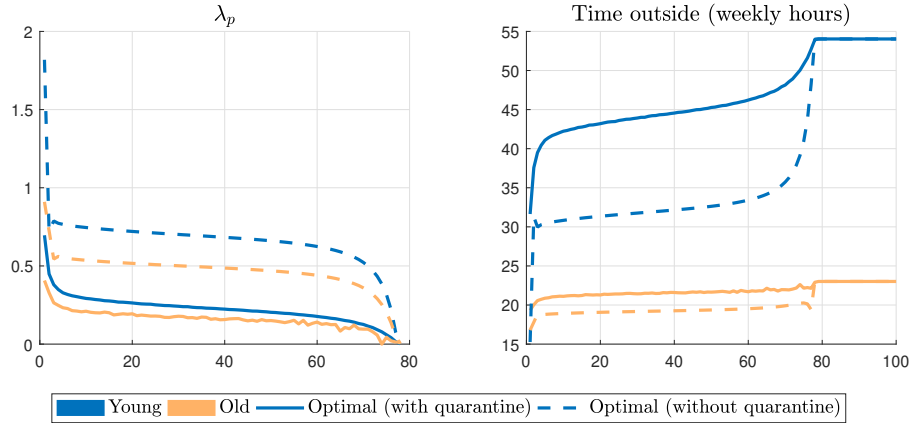
Even though around one eighth of the population catches Covid-19 with a 50% test-and-quarantine policy, agents change their behavior and engage in more risk. At the peak of the disease, the young spend about 11 hours per week longer outside the home compared to the benchmark. Older individuals, who are more affected by the disease, also spend less time at home: seven fewer hours per week relative to the baseline. This risk compensation acts to dampen the effect of the policy.³⁷

Policies that involve quarantines require infected agents to stay longer at home, even if this is against their best interest. However, this can be a welfare-improving policy. The last two rows of Table 8 report the welfare gains for a susceptible person at the outset of the disease, for the young and the old. The welfare for both young and old agents is higher with the quarantine policies in place though they know they might be required to quarantine themselves if they catch the disease.

In sum, testing turns out to be an extremely effective policy, especially when combined with quarantines. The gains in GDP seem much higher than what is needed to cover the costs of the policy. Testing also has non-linear effects: testing half of the exposed population and isolating the positive cases decreases deaths by more than 50%. Finally, age matters for testing as well. Testing the young only (the larger group) decreases the number of tests needed and leads to similar outcomes as a universal testing policy.

³⁷See Greenwood et al. (2019) for a discussion of the quantitative impact of risk-compensation effects on policies that aim to curb the HIV/AIDS epidemic.

Figure 9: Optimal policy with testing and quarantines



7.2 Optimal Lockdown with Test and Quarantine

We now return to the question of what level of confinement a utilitarian planner would choose. Now, however, we ask this question in the presence of testing and quarantines. In particular, suppose testing is available for 50% of those with fever and, if tested positive, the individual is quarantined through an increase in $\lambda_p(i, a)$ that would make young individuals spend an extra 90% of time at home in the no-Covid scenario. Table C4 in Appendix C.5 presents the details for this exercise.

With this extra instrument, like before, the planner aims for a virtual no-Covid strategy and there are very few deaths throughout the pandemic, only 0.02 per 1,000 people. Despite the similar health outcomes between the optimal policy with and without testing and quarantines, the policy, behavior and economic effects are markedly different. With the possibility of quarantining those that are found to be positive, the planner can relax the confinement among the others. The left panel of Figure 9 shows that the policy parameter λ_p is lower for both groups when testing and quarantines are in place. The right panel of Figure 9 shows that this lighter lockdown leads the old and especially the young to spend more time outside compared with the optimal lockdown without testing. With the young being able to spend more time outside and work, GDP is actually 2% higher than in the no-policy baseline and 10% higher than in the optimal policy without testing and quarantines.

8 Other Scenarios

This section contains the results for different scenarios and policy experiments. In particular, we assess the effects of selective mixing (Section 8.1), hospital bed constraints (Section 8.2) and an alternative calibration based on Ferguson et al. (2020) (Section 8.3).

8.1 Selective Mixing

Consider a government that reserves some common spaces for a particular age group only. This could entail that certain hours in supermarkets are reserved for a particular age group, or that some parks or leisure centers are separated in the same manner. This leads to *selective mixing* where agents are more likely to meet those of their own type. Suppose a fraction ζ of people's time outside can be separated into the different age groups and a fraction ϑ_a of the common space is allocated for age group a . Conditional on being outside, some infections now occur only within groups. To generate quantitative results, we use time use surveys to identify which activities could potentially be separated across the age groups and set $\zeta = 0.42$ accordingly. Moreover, assume the common space is divided (ϑ_a) according to the relative sizes of each group. The details are provided in Appendix C.6.

With this selective mixing, the overall death rate falls, both within the first year of the pandemic as well as in the long run, where deaths fall by about 5.5%. This result masks an important heterogeneity across the two groups though: the death rate for the old is reduced by almost 20%, while the young die about 8% more compared with the benchmark. With partially separated spaces, the old interact more among themselves, a group that naturally protects itself more. The opposite happens to the young. Life is now riskier for this group and they behave accordingly. At the peak of the disease, the young spend more time at home compared with the benchmark. The overall effects on the young are small and their output hardly changes, keeping the GDP in the first year of the pandemic only 1% lower than in the no-policy baseline.

8.2 Hospital Bed Constraints

We now explore the effects of scarce hospital resources. The American Hospital Association reports the existence of 84,555 intensive care beds excluding

neonatal units, implying a per-capita number of $Z = 0.032607\%$ ICUs.³⁸ We set $\tilde{\delta}_2(a) = 1$, so that the death for an agent in need of hospitalization that is not given a bed is certain to occur. The results are provided in Appendix C.7. Without behavioral changes (the epidemiological scenario), deaths would more than double. However, agents internalize this effect and change their behavior. At the peak of the disease, the young spend an extra hour at home per week. In the end, the overall death toll is similar to our baseline results as total deaths increase by less than 1%.

8.3 Alternative calibration based on Ferguson et al. (2020)

An influential study was conducted by researchers at the Imperial College London: Ferguson et al. (2020). Their results have been used to calibrate recent models in the economics literature. We recalibrate the disease parameters in our model to reflect those provided by this study. The main differences relative to those from CDC (2020) is that the the young (the old) are less (more) likely to end up in an ICU. On the other hand, the young (the old) are more (less) likely to die, conditional on being in an ICU. The results are qualitatively similar to those obtained with our benchmark calibration; see Appendix C.8. The total death rate is about 20% lower and this is completely explained by fewer deaths among the young; the old end up dying more.

9 Conclusions

This paper provides a framework to study the link between age heterogeneity, incomplete information/testing of the disease status, and the behavioral adjustments that individuals make to protect themselves during the Covid-19 pandemic. These seem to be first order in the spread of the infections and the deadliness of the disease. We embed these elements in an otherwise standard SIR model of disease transmission, calibrate it, and study policy interventions, both in general and targeted to particular age groups.

Especially the old protect themselves during the crisis, which is beneficial because they have a higher chance of dying. An optimal lockdown that weighs the different age groups according to their population shares essentially aims at a no-Covid strategy. It does so by specifically confining the young who contribute more to encounters and infections as they are less intrinsically motivated

³⁸See <https://www.aha.org/statistics/fast-facts-us-hospitals>

to protect themselves. For this strategy to work, the possibility of telework is key. Testing and quarantine are excellent ways of reducing the disease if feasible, even if just concentrated on the young. When testing is available, this contributes to relaxing the strictness of the optimal lockdown, and has a positive effect on GDP.

The model is richer than many existing counterparts, but remains sufficiently tractable to build upon in future work. We explore policies that impose selective mixing but future work could consider individuals who endogenously choose to interact with different groups. The economic costs of lockdowns in the form of closed schools have been particularly large for children, so adding a third age group would be a promising avenue to pursue. Our setup could be used to study other dimensions of heterogeneity such as education, sectors, and gender.³⁹ Another form of heterogeneity may be in the beliefs about the severity of Covid-19. If some people have wrong beliefs about the risk of dying, they have a lower incentive to protect themselves (see Greenwood et al. (2019) for an analysis along these lines in the context of HIV). The existence of such a group of people will have implications for optimal policy. Moreover, the framework can be applied to other pandemics with similar age gradient, such as past and future influenza pandemics. These extensions are part of a broader set of topics that warrant attention in future research.

References

- Acemoglu, Daron, Victor Chernozhukov, Ivan Werning, and Michael D. Whinston. 2020, May. "A multi-risk SIR model with optimally targeted lockdown." Working paper 27102, National Bureau of Economic Research.
- Acemoglu, Daron, Azarakhsh Malekian, and Asu Ozdaglar. 2016. "Network security and contagion." *Journal of Economic Theory* 166:536 – 585.
- Adams-Prassl, Abi, Teodora Boneva, Marta Golin, and Christopher Rauh. 2020. "Work That Can Be Done from Home: Evidence on Variation within and across Occupations and Industries." IZA Discussion Paper No. 13374.
- Alon, Titan, Matthias Doepke, Jane Olmstead-Rumsey, and Michèle Tertilt. 2020a, August. "This Time It's Different: The Role of Women's Employment in a Pandemic Recession." NBER Working Paper 27660.

³⁹Education and income are highly correlated with the ability to work from home (Adams-Prassl et al. 2020) and women are hit more by employment losses (Alon et al. 2020a).

- Alon, Titan, Minki Kim, David Lagakos, and Mitchell VanVuren. 2020b. "How should policy responses to the Covid-19 pandemic differ in the developing world?" *Covid Economics* 22:1–46.
- Alvarez, Fernando, David Argente, and Francesco Lippi. 2020. "A Simple Planning Problem for COVID-19 Lockdown." *American Economic Review: Insights*, vol. forthcoming.
- Andersson, Ola, Pol Campos-Mercade, Fredrik Carlsson, Florian Schneider, and Erik Wengström. 2020. "The Individual Welfare Costs of Stay-At-Home Policies." Lund University Working Paper 2020:9.
- Arias, Elizabeth, and Jiaquan Xu. 2019. "United States life tables, 2017." *National Vital Statistics Reports* 68, no. 7.
- Atkeson, Andrew. 2020, March. "What Will Be the Economic Impact of COVID-19 in the US? Rough Estimates of Disease Scenarios." Working paper 26867, National Bureau of Economic Research.
- Aum, Sangmin, Sang Yoon (Tim) Lee, and Yongseok Shin. 2020, October. "Who Should Work from Home during a Pandemic? The Wage-Infection Trade-off." Working paper 27908, National Bureau of Economic Research.
- Bairoliya, Neha, and Ayşe İmrohoroglu. 2020, November. "COVID-19: Mitigation Measures and the Aftershock in an Overlapping Generations Model." University of Southern California, Unpublished Manuscript.
- Belot, Michèle, Syngjoo Choi, Julian C. Jamison, Nicholas Papageorge, Egon Tripodi, and Eline Van den Broek-Altenburg. 2020. "Unequal Consequences of Covid 19 across Age and Income: Representative Evidence from Six Countries." CEPR Discussion Paper DP14908.
- Berger, David W, Kyle F Herkenhoff, Chengdai Huang, and Simon Mongey. 2020. "Testing and Reopening in an SEIR Model." *Review of Economic Dynamics*, vol. forthcoming.
- Bick, Alexander, Adam Blandin, and Karel Mertens. 2020. "Work from Home After the COVID-19 Outbreak." Working paper 2017, Federal Reserve Bank of Dallas.
- Biggs, Andrew G., and Glenn R. Springstead. 2008. "Alternate Measures of Replacement Rates for Social Security Benefits and Retirement Income." *Social Security Bulletin* 68, no. 2.

- Bognanni, Mark, Douglas Hanley, Daniel Kolliner, and Kurt Mitman. 2020, September. "Economics and Epidemics: Evidence from an Estimated Spatial Econ-SIR Model." Technical Report DP15310, CEPR.
- Boppart, Timo, Karl Harmenberg, John Hassler, Per Krusell, and Jonna Ols-son. 2020, December. "Integrated epi-econ assessment." Technical Report DP15595, CEPR.
- Brotherhood, Luiz, Tiago Cavalcanti, Daniel Da Mata, and Cezar Santos. 2020a. "Slums and Pandemics." CEPR Discussion Paper DP14695.
- Brotherhood, Luiz, Philipp Kircher, Cezar Santos, and Michèle Tertilt. 2020b, May. "An economic model of the Covid-19 epidemic: The importance of testing and age-specific policies." CEPR discussion paper No. DP14695.
- Butler, Rachel, Mauricio Monsalve, Geb W. Thomas, Ted Herman, Alberto M. Segre, Philip M. Polgreen, and Manish Suneja. 2018. "Estimating Time Physicians and Other Health Care Workers Spend with Patients in an Intensive Care Unit Using a Sensor Network." *The American Journal of Medicine* 131 (8): 972.e9–972.e15.
- CDC. 2020. "Severe Outcomes Among Patients with Coronavirus Disease 2019 (COVID-19) — United States, February 12–March 16, 2020." *MMWR Morb Mortal Wkly Rep* 2020 69 (12): 343–346.
- Chan, Tat, Barton Hamilton, and Nicholas Papageorge. 2016. "Health, Risky Behavior and the Value of Medical Innovation for Infectious Disease." *Review of Economic Studies* 83 (3): 1737–1755.
- Chen, Frederick. 2012. "A Mathematical Analysis of Public Avoidance Behavior During Epidemics Using Game Theory." *Journal of Theoretical Biology* 302:18–28.
- Eichenbaum, Martin S, Sergio Rebelo, and Mathias Trabandt. 2020a, March. "The Macroeconomics of Epidemics." Working paper 26882, National Bureau of Economic Research.
- . 2020b, May. "The Macroeconomics of Testing and Quarantining." Working paper 27104, National Bureau of Economic Research.
- Farboodi, Maryam, Gregor Jarosch, and Robert Shimer. 2020, April. "Internal and External Effects of Social Distancing in a Pandemic." Working paper 27059, National Bureau of Economic Research.

- Favero, Carlo, Andrea Ichino, and Aldo Rustichini. 2020. "Restarting the economy while saving lives under Covid-19." CEPR Discussion Paper DP14664.
- Ferguson, Neil M., Daniel Laydon, Gemma Nedjati-Gilani, Natsuko Imai, Kylie Ainslie, Marc Baguelin, ..., and Azra C Ghani. 2020, 19 March. "Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand." Technical Report, Imperial College.
- Fernández-Villaverde, Jesús, and Charles I. Jones. 2020, May. "Estimating and Simulating a SIRD Model of COVID-19 for Many Countries, States, and Cities." Working paper 27128, National Bureau of Economic Research.
- Galeotti, Andrea, and Brian R. Rogers. 2012. "Immunization and Group Structure." *American Economic Journal-Microeconomics*, 5(2): 1-32. 5 (2): 1–32.
- Garibaldi, Pietro, Espen R. Moen, and Christopher A. Pissarides. 2020a. "Modelling contacts and transitions in the SIR epidemics model." *Covid Economics* 5 (April): 1–20.
- . 2020b, November. "Static and Dynamic Inefficiencies in a Model of Epidemics." Technical Report DP15439, CEPR.
- Glover, Andrew, Jonathan Heathcote, Dirk Krueger, and José-Víctor Ríos-Rull. 2020. "Health versus Wealth: On the Distributional Effects of Controlling a Pandemic." Working paper 27128, National Bureau of Economic Research.
- Gollier, Christian. 2020a. "Cost-benefit analysis of age-specific deconfinement strategies." *Journal of Public Economic Theory* 22 (6): 1746–1771.
- . 2020b. "If the Objective is Herd Immunity, on Whom Should it be Built?" *Environmental and Resource Economics* 76:671–683.
- Greenwood, Jeremy, Philipp Kircher, Cezar Santos, and Michèle Tertilt. 2013, March. "An Equilibrium Model of the African HIV/AIDS Epidemic." Working paper 43, National Bureau of Economic Research.
- . 2017. "The Role of Marriage in Fighting HIV: A Quantitative Evaluation for Malawi." *American Economic Review* 117 (5): 158–162 (May).
- . 2019. "An Equilibrium Model of the African HIV/AIDS Epidemic." *Econometrica* 87 (4): 1081–1113 (July).
- Hall, Robert E., Charles I. Jones, and Peter J. Klenow. 2020. "Trading Off Consumption and COVID-19 Deaths." *Minneapolis Fed Quarterly Review* 42, no. 1 (June).

- Heikkinen, Terho, and Asko Järvinen. 2003. "The common cold." *The Lancet* 361 (9351): 51–59.
- Jain, Seema, Wesley H. Self, Richard G. Wunderink, Sherene Fakhran, Robert Balk, Anna M. Bramley, Carrie Reed, Carlos G. Grijalva, Evan J. Anderson, D. Mark Courtney, James D. Chappell, Chao Qi, and Er Hart. 2015. "Community-Acquired Pneumonia Requiring Hospitalization among U.S. Adults." *New England Journal of Medicine* 373 (5): 415–427. PMID: 26172429.
- Keppo, Jussi, Elena Quercioli, Marianna Kudlyak, Lones Smith, and Andrea Wilson. 2020. "The Behavioral SIR Model, with Applications to the Swine Flu and COVID-19 Pandemics." Unpublished Manuscript, Wisconsin.
- Kermack, W. O., and A. G. McKendrick. 1927. "A Contribution to the Mathematical Theory of Epidemics." *Proceedings of the Royal Society A* 114 (772): 700–721.
- Kopecky, Karen A. 2011. "The trend in retirement." *International Economic Review* 52 (2): 287–316.
- Kremer, Michael. 1996. "Integrating Behavioral Choice into Epidemiological Models of AIDS." *The Quarterly Journal of Economics* 111 (2): 549–573 (05).
- Maloney, William, and Temel Taskin. 2020. "Determinants of Social Distancing and Economic Activity during COVID-19: A Global View." World Bank Policy Research Working Paper 9242.
- Manfredi, Piero, and Alberto D’Onofrio, eds. 2013. *Modeling the Interplay Between Human Behavior and the Spread of Infectious Diseases*. Springer Verlag New York.
- McAdams, David. 2020. "Nash SIR: An Economic-Epidemiological Model of Strategic Behavior During a Viral Epidemic." *Covid Economics* 16:115–134.
- Meyer, K. C. 2001. "The role of immunity in susceptibility to respiratory infection in the aging lung." *Respiration physiology* 128, no. 1.
- Mossong, Joel, Niel Hens, Mark Jit, Philippe Beutels, Kari Auranen, Rafael Mikolajczyk, et al. 2008. "Social Contacts and Mixing Patterns Relevant to the Spread of Infectious Diseases." *PLoS Med* 5 (3): e74.
- Piguillem, Facundo, and Liyan Shi. 2020. "The Optimal COVID-19 Quarantine and Testing Policies." EIEF Working Paper Series 2004, Einaudi Institute for Economics and Finance.

- Quercioli, Elena, and Lones Smith. 2006. "Contagious Matching Games." Unpublished Manuscript, University of Michigan.
- Remuzzi, Andrea, and Giuseppe Remuzzi. 2020. "COVID-19 and Italy: what next?" *The Lancet* 395 (10231): 1225–1228 (2020/04/21).
- Rizzo, Caterina, Massimo Fabiani, Richard Amlot, Ian Hall, Thomas Finnie, G. James Rubin, ..., and Andrea Pugliese. 2013. "Survey on the Likely Behavioural Changes of the General Public in Four European Countries During the 2009/2010 Pandemic." In *Modeling the Interplay Between Human Behavior and the Spread of Infectious Diseases*, edited by Piero Manfredi and Alberto D'Onofrio, 23–42. Springer Verlag New York.
- Toxvaerd, Flavio. 2019. "Rational Disinhibition and Externalities in Prevention." *International Economic Review* 60 (4): 1737–1755 (1).
- Verity, Robert, Lucy C Okell, Ilaria Dorigatti, Peter Winskill, Charles Whitaker, Natsuko Imai, ..., and Neil M Ferguson. 2020. "Estimates of the severity of coronavirus disease 2019: a model-based analysis." *The Lancet Infectious Diseases*.
- von Thadden, Ernst-Ludwig. 2020. "A simple, non-recursive model of the spread of Covid-19 with applications to policy." *Covid Economics* 10 (April): 24–43.
- WHO. 2003. Update 49 - SARS case fatality ratio, incubation period. https://www.who.int/csr/sars/archive/2003_05_07a/en/. Date accessed: January 7th 2021.
- . 2021. Preparing for Pandemics. <https://www.who.int/westernpacific/activities/preparing-for-pandemics>. Date accessed: January 7th 2021.
- Wong, J.Y., H. Kelly, D.K. Ip, J.T. Wu, G.M. Leung, and B.J. Cowling. 2013. "Case fatality risk of influenza A (H1N1pdm09): a systematic review." *Epidemiology* 24, no. 6.

Online Appendix - Not for Publication

A Laws of Motion

In the main text, equation (8) describes the overall laws of motion, and (9) describes the sub-part that determines the transitions for susceptible agents. The following contains the transitions for all other types. It also includes the accounting of Covid-deaths and new infections.

The number of fever-susceptible agents who have a fever and are not tested but are truly Covid-negative and susceptible is given by

$$\begin{aligned} M_{t+1}(f_s, a) &= M_t(s, a)\Delta(a)(1 - \xi_{p_t}(a))\pi_f(n_t(s, a) + \ell_t(s, a), \Pi_t(a))\frac{\Pi^*}{\Pi_t(a) + \Pi^*} \\ &+ M_t(f_s, a)\Delta(a)(1 - \xi_{p_t}(a))\pi_f(n_t(f, a) + \ell_t(f, a), \Pi_t(a))\frac{\Pi^*}{\Pi_t(a) + \Pi^*}. \end{aligned} \quad (14)$$

It includes in the first line susceptible people from last period who got fever but were not tested, and are truly Covid-negative and susceptible. The second line again accounts for those in the fever-susceptible state, as they can again catch another fever while truly remaining susceptible.

A similar logic applies to those in the fever-infected state:

$$\begin{aligned} M_{t+1}(f_i, a) &= M_t(s, a)\Delta(a)(1 - \xi_{p_t}(a))\pi_f(n_t(s, a) + \ell_t(s, a), \Pi_t(a))\frac{\Pi_t(a)}{\Pi_t(a) + \Pi^*} \\ &+ M_t(f_s, a)\Delta(a)(1 - \xi_{p_t}(a))\pi_f(n_t(f, a) + \ell_t(f, a), \Pi_t(a))\frac{\Pi_t(a)}{\Pi_t(a) + \Pi^*}. \end{aligned} \quad (15)$$

The total number of individuals in the fever state is then

$$M_{t+1}(f, a) = M_{t+1}(f_s, a) + M_{t+1}(f_i, a) \quad (16)$$

To account for infected people, one counts those who started last period susceptible or fever-susceptible and get infected and tested this period, but also those who started last period infected or fever-infected who neither required

hospitalization nor recovered:

$$\begin{aligned}
M_{t+1}(i, a) = & M_t(s, a)\Delta(a)\xi_{p_t}(a)\pi(n_t(s, a) + \ell_t(s, a), \Pi_t(a)) \\
& + M_t(f_s, a)\Delta(a)\xi_{p_t}(a)\pi(n_t(f, a) + \ell_t(f, a), \Pi_t(a)) \\
& + [M_t(f_i, a) + M_t(i, a)] \Delta(a)(1 - \phi(0, a))(1 - \alpha(a))
\end{aligned} \tag{17}$$

People in hospitals comprise those who entered last period infected or fever-infected and do not recover but instead require hospitalization, as well as those individuals that were already hospitalized in the previous period who neither die nor recover:

$$\begin{aligned}
M_{t+1}(h, a) = & [M_t(f_i, a) + M_t(i, a)] \Delta(a)(1 - \phi(0, a))\alpha(a) \\
& + M_t(h, a)\Delta(a)(1 - \delta_t(a))(1 - \phi(1, a))
\end{aligned} \tag{18}$$

The total number of individuals hospitalized is then

$$M_t(h) = \sum_a M_t(h, a) \tag{19}$$

Recovered and therefore resistant individuals comprise those who were infected or fever-infected and recover, those hospitalized who do not die but recover, and resistant individuals from the previous period:

$$\begin{aligned}
M_{t+1}(r, a) = & [M_t(f_i, a) + M_t(i, a)] \Delta(a)\phi(0, a) \\
& + M_t(h, a)\Delta(a)\phi(1, a) + \Delta(a)M_t(r, a)
\end{aligned} \tag{20}$$

The right hand sides of equations (14) to (20) gives the map T_j for states $j = f_s, f_i, f, i, h, r$.

For accounting purposes, the measure of deceased agents as a result of Covid-19 is given by new Covid deaths and those who died of it in previous periods:

$$M_{t+1}(deceased, a) = M_t(deceased, a) + (1 - \phi(1, a))\delta_t(a)M_t(h, a)\Delta(a),$$

while the number of newly infected people is given by susceptible or fever-

susceptible agents who get infected

$$N_{t+1}(i, a) = M_t(s, a)\Delta(a)\pi(n_t(s, a) + \ell_t(s, a), \Pi_t(a)) \\ + M_t(f_s, a)\Delta(a)\pi(n_t(f, a) + \ell_t(f, a), \Pi_t(a)).$$

B Details on Computations

B.1 Computing Weekly Rates

Let C be the fraction of the population that catches the common cold every year. Then, the weekly infection rate Π^* is given by:

$$\Pi^* = 1 - (1 - C)^{1/52}.$$

Now, consider an agent that is infected with Covid-19. They may recover with probability $\phi(0)$ or require hospitalization with probability α . The following table gives what happens to a measure 1 of agents that are infected right now over the course of the next few weeks.

Week	Frac recovered	Frac still infected	Frac in hospital
1	$\phi(0)$	$(1 - \phi(0))(1 - \alpha)$	$(1 - \phi(0))\alpha$
2	$(1 - \phi(0))(1 - \alpha)\phi(0)$	$[(1 - \phi(0))(1 - \alpha)]^2$	$(1 - \phi(0))(1 - \alpha)(1 - \phi(0))\alpha$
3	$[(1 - \phi(0))(1 - \alpha)]^2 \phi(0)$	$[(1 - \phi(0))(1 - \alpha)]^3$	$[(1 - \phi(0))(1 - \alpha)]^2 (1 - \phi(0))\alpha$
4

Thus, the fraction of people that will need hospitalizations F_h is given by

$$F_h = (1 - \phi(0))\alpha + (1 - \phi(0))(1 - \alpha)(1 - \phi(0))\alpha + [(1 - \phi(0))(1 - \alpha)]^2 (1 - \phi(0))\alpha + \dots \\ = (1 - \phi(0))\alpha [1 + (1 - \phi(0))(1 - \alpha) + [(1 - \phi(0))(1 - \alpha)]^2 + \dots] \\ = (1 - \phi(0))\alpha \frac{1}{1 - (1 - \phi(0))(1 - \alpha)}.$$

Solving out for α gives

$$\alpha = \frac{B\phi(0)}{1 - B(1 - \phi(0))},$$

where $B = F_h/(1 - \phi(0))$. With $\phi(0)$ given by the average time for recovery, one can use the formula above to get α .

We can do something similar for agents in hospitals to figure out at what weekly rate they die. Here is the table:

Week	Frac recovered	Frac still in hospital	Frac dead
1	$\phi(1)$	$(1 - \phi(1))(1 - \delta)$	$(1 - \phi(1))\delta$
2	$(1 - \phi(1))(1 - \delta)\phi(1)$	$[(1 - \phi(1))(1 - \delta)]^2$	$(1 - \phi(1))(1 - \delta)(1 - \phi(1))\delta$
3	$[(1 - \phi(1))(1 - \delta)]^2 \phi(1)$	$[(1 - \phi(1))(1 - \delta)]^3$	$[(1 - \phi(1))(1 - \delta)]^2 (1 - \phi(1))\delta$
4

Using the same steps above and denoting the fraction that die by F_d , we get:

$$\delta = \frac{A\phi(1)}{1 - A(1 - \phi(1))},$$

where $A = F_d/(1 - \phi(1))$.

B.2 Basic Reproduction Number - R_0

The probability that an infected agent leaves such state is: $[\phi(0) + (1 - \phi(0))\alpha] \Delta + (1 - \Delta)$. The squared brackets term is the probability of recovery and the probability that the agent requires hospitalization, conditional on surviving natural causes (probability Δ). The last term is death due to natural causes. Hence, the expected amount of time one stays in state i is:

$$T_i = \frac{1}{[\phi(0) + (1 - \phi(0))\alpha] \Delta + (1 - \Delta)}.$$

The probability that an agent in a hospital leaves such state is: $[\phi(1) + (1 - \phi(1))\delta] \Delta + (1 - \Delta)$. The squared brackets term is the probability of recovery and the death-because-of-Covid probability, conditional on surviving natural causes (probability Δ). The last term is death due to natural causes. Hence, the expected amount of time one stays in state h is:

$$T_h = \frac{1}{[\phi(1) + (1 - \phi(1))\delta] \Delta + (1 - \Delta)}.$$

Now, the probability that one moves from the i state to the h state is given by:

$$P_h = \frac{(1 - \phi(0))\alpha\Delta}{1 - (1 - \phi(0))(1 - \alpha)\Delta}.$$

Note that the expressions above should be functions of one's age a , but we have suppressed this for notational convenience.

Let $\tilde{n}(a)$ denote the amount of time an infected person of age a spends outside. Let $\bar{\ell}$ be the interaction time for people in hospitals. Finally, let \bar{n} be the

average (across ages) amount of time people spend outside. At the outset of the disease, a measure 1 of the population is healthy.

Then, $R_0(a)$ (i.e. for an infected person of age a) is given by:

$$R_0(a) = [\tilde{n}(a)T_i(a) + \bar{\ell}P_h(a)T_h(a)] \bar{n}\Pi_0.$$

This is the average number of people someone infects (for a person of a given age). The economy's R_0 will be the weighted average across ages:

$$R_0 = \sum_a \omega(a)R_0(a),$$

where $\omega(a)$ is the weight of age a in the population.

B.3 Computing CEV

We use a consumption equivalent variation (CEV) measure to evaluate the welfare change of an agent between two scenarios. To do that, we compute the relative change in the consumption good, c , that an agent needs to receive in all periods and states of nature of a given scenario so that their expected discounted utility measured in the first time period as a susceptible agent is equal to that in another scenario. Moreover, we evaluate the value function of the agents with $\lambda_p = 0$, the planner's influence on the individual's utility. This parameter is just a way to operationalize a lockdown. Thus, we evaluate the individual's value function with the distorted allocation and $\lambda_p = 0$. For an example of this welfare measure, Table 6 shows that old agents are indifferent between the optimal policy and having their consumption increased by 6.47% in all periods in the benchmark equilibrium.

C Additional Results

This section provides additional details and results for the following scenarios: a world without altruism (Section C.1), a scenario in which a vaccine arrives only after 15 years (Section C.2), details for the optimal lockdown (Section C.3), a shelter-at-home policy with a late start (Section C.4), details for the optimal lockdown with testing and quarantine (Section C.5), details on the model with selective mixing (Section C.6), a calibration with hospital bed constraints (Section C.7), and an alternative calibration based on different hospitalization and

Table C1: No altruism

	Benchmark	Epidem.
Wks to peak hospitalizations (yng)	14.00	12.00
Wks to peak hospitalizations (old)	11.00	12.00
Hospitalizations p/ 1,000 @ peak (yng)	0.90	12.81
Hospitalizations p/ 1,000 @ peak (old)	0.22	11.11
Dead p/ 1,000 1year (yng)	1.37	4.04
Dead p/ 1,000 1year (old)	4.33	31.40
Dead p/ 1,000 1year (all)	2.00	9.89
Dead p/ 1,000 LR (yng)	1.83	4.04
Dead p/ 1,000 LR (old)	6.15	31.40
Dead p/ 1,000 LR (all)	2.75	9.89
Immune in LR (yng), %	38.57	85.29
Immune in LR (old), %	9.20	45.81
Immune in LR (all), %	32.29	76.84
GDP at peak - rel to BM	1.00	1.17
GDP 1year - rel to BM	1.00	1.10
Hrs @ home (yng) - peak	79.66	57.97
Hrs @ home (old) - peak	105.70	88.99
Hrs @ home (yng) - 6m	77.01	57.97
Hrs @ home (old) - 6m	104.31	88.99

death rates as provided by Ferguson et al. (2020) (Section C.8).

C.1 No Altruism

Table C1 compares the benchmark version of the model with no altruism ($\lambda_a = 0$) versus the epidemiological version with no behavioral adjustments. The qualitative results are very similar to the original baseline results. The main difference is that, without altruism, the death toll is higher. This is due to the fact that, as infected agents do not change their behavior, they end up infecting more people. Note that this world is riskier for healthy people, who behave more cautiously in response. Note that the hours at home for susceptible agents are higher compared to the benchmark with altruism.


C.2 Vaccine arrival in 15 years

In our baseline calibration, a vaccine arrives after 78 weeks. After this, nobody gets infected with Covid-19 anymore. In this section, we explore the effects of a scenario in which a vaccine only arrives after 15 years; i.e. $T^* = 15 * 52 = 780$. As the pandemic lasts longer now, more people get infected; in the long run, more than half the population is immune. The death toll is consequently higher, about twice as much as in the baseline calibration with the arrival of the vaccine in 78 weeks. As discussed in Section 5, as herd immunity is more important

Table C2: Vaccine arrives after 15 years

	Benchmark	Epidem.	Age ext. partial	Age ext. general	No disease
Wks to peak hospitalizations (yng)	14.00	12.00	15.00	17.00	
Wks to peak hospitalizations (old)	12.00	12.00	12.00	12.00	
Hospitalizations p/ 1,000 @ peak (yng)	1.15	12.81	0.24	0.17	
Hospitalizations p/ 1,000 @ peak (old)	0.25	11.11	0.25	0.10	
Dead p/ 1,000 1year (yng)	1.62	4.04	0.42	0.35	
Dead p/ 1,000 1year (old)	4.83	31.40	4.83	2.24	
Dead p/ 1,000 1year (all)	2.31	9.89	1.36	0.76	
Dead p/ 1,000 LR (yng)	2.97	4.04	1.58	2.44	
Dead p/ 1,000 LR (old)	12.55	31.40	12.56	13.40	
Dead p/ 1,000 LR (all)	5.02	9.89	3.93	4.79	
Immune in LR (yng), %	62.77	85.29	33.35	51.53	
Immune in LR (old), %	10.16	23.19	10.16	12.55	
Immune in LR (all), %	51.51	72.00	28.39	43.19	
GDP at peak - rel to BM	1.00	1.12	0.72	0.96	1.14
GDP 1year - rel to BM	1.00	1.07	0.72	0.93	1.10
Cost p/ life saved, million \$	—	—	15.24	2.42	
Hrs @ home (yng) - peak	75.05	57.97	104.95	78.97	57.97
Hrs @ home (old) - peak	105.74	88.99	105.74	95.71	88.99
Hrs @ home (yng) - 6m	72.54	57.97	103.56	78.61	57.97
Hrs @ home (old) - 6m	104.10	88.99	104.10	95.56	88.99

Figure C1: Choices in the optimal policy



without the vaccine, the age externality works differently: as the young become more cautious, the burden of herd immunity falls more heavily on the old. The death toll among the old thus rises.

C.3 Optimal Lockdown - Details

Figure C1 reports the time allocation for young and old individuals under an age-specific optimal lockdown.

Table C3 presents the results of the optimal lockdown without the possibility of telework. Deaths decrease substantially but not as much as in the benchmark with teleworking. Figure C2 reports the policy parameters for the old and the

Table C3: Optimal policy, no teleworking

	Benchmark	Optimal policy
Wks to peak hospitalizations (yng)	12.00	79.00
Wks to peak hospitalizations (old)	11.00	79.00
Hospitalizations p/ 1,000 @ peak (yng)	3.65	0.79
Hospitalizations p/ 1,000 @ peak (old)	0.28	0.17
Dead p/ 1,000 1year (yng)	2.83	0.08
Dead p/ 1,000 1year (old)	4.90	0.47
Dead p/ 1,000 1year (all)	3.27	0.17
Dead p/ 1,000 LR (yng)	3.05	0.49
Dead p/ 1,000 LR (old)	6.30	1.71
Dead p/ 1,000 LR (all)	3.75	0.75
Immune in LR (yng), %	64.39	10.32
Immune in LR (old), %	9.39	2.62
Immune in LR (all), %	52.62	8.67
GDP at peak - rel to BM	1.00	1.00
GDP 1year - rel to BM	1.00	0.67
Hrs @ home (yng) - peak	76.17	74.68
Hrs @ home (old) - peak	110.18	106.35
Hrs @ home (yng) - 6m	65.49	77.93
Hrs @ home (old) - 6m	106.47	97.32
CEV rel. to BM (yng), %	—	0.46
CEV rel. to BM (old), %	—	6.25

young and the resulting time spent outside.

C.4 Shelter-at-home with a Late Start

Figure C3 reports the result of a shelter-at-home policy that is only implemented 8 weeks after the outbreak of the pandemic. This strict policy requires individuals to spend an extra 90% of time at home and lasts for 26 weeks.

Figure C2: Optimal policy, no teleworking

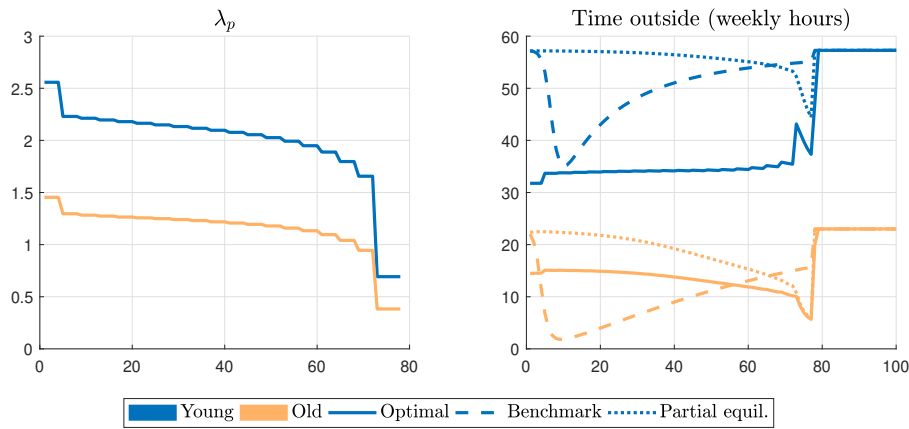
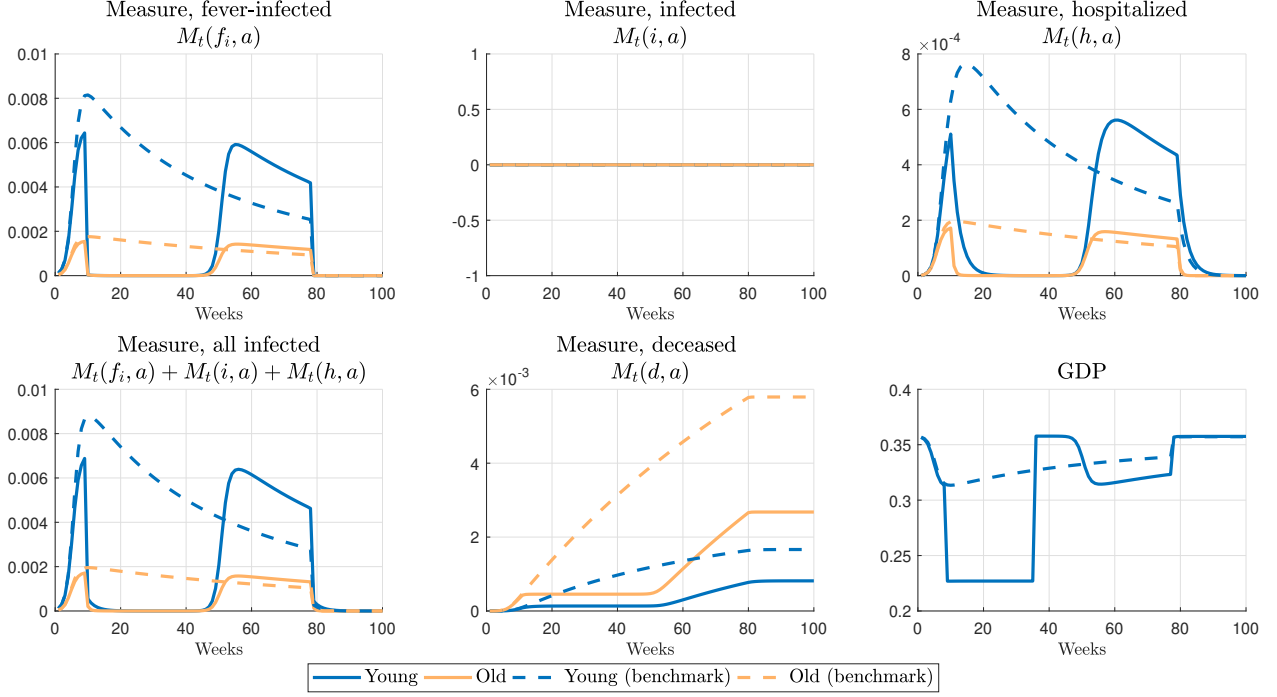


Figure C3: Aggregate variables (late shelter-in-place)



C.5 Optimal Lockdown with Testing and Quarantine - Details

Table C4 reports the results under an optimal policy in the presence of testing and quarantines.

C.6 Selective Mixing - Details

In order to model selective mixing across age groups, assume that only a fraction $1 - \zeta$ of activities remain common and take up a corresponding fraction $1 - \zeta$ of the outside space. But a fraction ζ of activities can be separated into age groups, and for this purpose a fraction ϑ_a of the remaining space is dedicated only to individuals of this specific age group. Those spend $T(\vartheta_a) \leq 1$ of their remaining time there, the rest of the time is lost as individuals arrive in situations where the desired space happens to be dedicated to the other group. Conditional on being in common space where one can get infected, the infection rate

Table C4: Optimal policy with testing and quarantines

	Benchmark	Optimal policy
Wks to peak hospitalizations (yng)	15.00	79.00
Wks to peak hospitalizations (old)	11.00	79.00
Hospitalizations p/ 1,000 @ peak (yng)	0.76	0.01
Hospitalizations p/ 1,000 @ peak (old)	0.20	0.01
Dead p/ 1,000 1year (yng)	1.22	0.01
Dead p/ 1,000 1year (old)	4.03	0.03
Dead p/ 1,000 1year (all)	1.82	0.01
Dead p/ 1,000 LR (yng)	1.66	0.01
Dead p/ 1,000 LR (old)	5.79	0.07
Dead p/ 1,000 LR (all)	2.55	0.02
Immune in LR (yng), %	35.12	0.24
Immune in LR (old), %	8.67	0.10
Immune in LR (all), %	29.46	0.21
GDP at peak - rel to BM	1.00	1.13
GDP 1year - rel to BM	1.00	1.02
Hrs @ home (yng) - peak	76.29	59.28
Hrs @ home (old) - peak	104.44	89.86
Hrs @ home (yng) - 6m	74.66	68.32
Hrs @ home (old) - 6m	103.31	90.57
CEV rel. to BM (yng), %	–	1.33
CEV rel. to BM (old), %	–	6.60

now becomes

$$\begin{aligned}
\hat{\Pi}_t(a) = & (1 - \zeta)\Pi_0 \sum_{\tilde{a}, j \in \{f_i, i, s\}} (n_t(j, \tilde{a}) + \ell_t(j, \tilde{a})) M_t(j, \tilde{a}) \\
& + \zeta T(\vartheta_a)\Pi_0 \sum_{j \in \{f_i, i, s\}} \frac{T(\vartheta_a)}{\vartheta_a} (n_t(j, a) + \ell_t(j, a)) M_t(j, a).
\end{aligned} \tag{21}$$

The first line reflects the $1 - \zeta$ of a person's time spent in the unrestricted space where everything is unchanged from equation (11) in the main text: other individuals spend $1 - \zeta$ of their time across $1 - \zeta$ of the space which leaves the number per area unchanged. The second line reflects the fraction ζ of a person's time spent on age-restricted activities, of which $T(\vartheta_a)$ is lost, where they only meet others of the same age who spend $T(\vartheta_a)\zeta$ of their time on $\vartheta_a\zeta$ parts of the space. The expression reduces to the random mixing rate (11) if selectivity is $\zeta = 0$. It also reduces to random mixing rate independent of selectivity ζ if young and old would be completely identical and space is divided up such according to group size ($\vartheta_y = \sum_j (M(a, j)) / \sum_{a, j} M(a, j)$, say, in steady state) and there is no loss of time due to separation ($T(\vartheta_a) = 1$). If a particular part of time is lost (i.e., $T(\vartheta_a) = \vartheta_a^{1/2}$) then ϑ_a cancels in (21) and the expression reduces to "preferred matching" in Jacquez et al (1988) and Kremer (1996), only that infected agents are here weighted by their activity outside the house.

Table C5: Selective mixing

	Benchmark	Sel. mix.
Wks to peak hospitalizations (yng)	15.00	14.00
Wks to peak hospitalizations (old)	11.00	11.00
Hospitalizations p/ 1,000 @ peak (yng)	0.76	0.85
Hospitalizations p/ 1,000 @ peak (old)	0.20	0.16
Dead p/ 1,000 1year (yng)	1.22	1.33
Dead p/ 1,000 1year (old)	4.03	3.31
Dead p/ 1,000 1year (all)	1.82	1.75
Dead p/ 1,000 LR (yng)	1.66	1.79
Dead p/ 1,000 LR (old)	5.79	4.72
Dead p/ 1,000 LR (all)	2.55	2.41
Immune in LR (yng), %	35.12	37.74
Immune in LR (old), %	8.67	7.06
Immune in LR (all), %	29.46	31.17
GDP at peak - rel to BM	1.00	0.98
GDP 1year - rel to BM	1.00	0.99
Cost p/ life saved, million \$	–	3.51
Hrs @ home (yng) - peak	76.29	77.99
Hrs @ home (old) - peak	104.44	102.34
Hrs @ home (yng) - 6m	74.66	76.12
Hrs @ home (old) - 6m	103.31	100.92
CEV rel. to BM (yng), %	–	-0.13
CEV rel. to BM (old), %	–	1.64

The cost of selective mixing is that at some times some of the common space is no longer available to a person who would like to use it. That is, there is a time loss of $(1 - T(\vartheta_a))\zeta > 0$ per time unit spent outside, which acts as a tax on both labor time n and outside leisure time ℓ , which are reduced by this as this time is lost and neither brings income nor utility. The simplest version is $T(\vartheta_a) = \vartheta_a$, where each unit of space dedicated to the other group is lost.

To generate quantitative results for the selective-mixing scenario, we must pick values for ζ and ϑ . According to the American Time Use Survey (ATUS), 42% of non-work time outside the home by the average American is spent purchasing goods and services and in organizational, civic, and religious activities. We take these activities as those that can be separated by age groups. We thus set $\zeta = 0.42$. Moreover, we set $\vartheta_o = 0.214$ and $\vartheta_y = 0.786$. The latter divides the space up according to the relative sizes of each group in the overall population. As explained above, this leaves a priori no reason for efficiency costs (i.e., if the groups were identical, the infections would not change). The results from this experiment are reported in Table C5.

Table C6: Benchmark results (hospital constraints)

	Benchmark	Epidem.	Age ext. partial	Age ext. general
Wks to peak hospitalizations (yng)	12.00	11.00	11.00	17.00
Wks to peak hospitalizations (old)	9.00	12.00	9.00	12.00
Hospitalizations p/ 1,000 @ peak (yng)	0.41	6.45	0.27	0.12
Hospitalizations p/ 1,000 @ peak (old)	0.16	9.39	0.16	0.07
Dead p/ 1,000 1year (yng)	0.84	19.02	0.56	0.24
Dead p/ 1,000 1year (old)	3.48	33.36	3.48	1.62
Dead p/ 1,000 1year (all)	1.41	22.09	1.18	0.54
Dead p/ 1,000 LR (yng)	1.36	19.02	0.90	0.38
Dead p/ 1,000 LR (old)	5.34	33.36	5.35	2.48
Dead p/ 1,000 LR (all)	2.21	22.09	1.85	0.83
Immune in LR (yng), %	28.66	82.87	9.27	8.10
Immune in LR (old), %	8.02	44.56	8.02	3.71
Immune in LR (all), %	24.24	74.67	9.01	7.16
GDP at peak - rel to BM	1.00	1.16	0.74	0.99
GDP 1year - rel to BM	1.00	1.10	0.74	0.96
Cost p/ life saved, million \$	–	–	48.26	2.10
Hrs @ home (yng) - peak	78.56	57.97	101.86	79.28
Hrs @ home (old) - peak	101.85	88.99	101.84	94.76
Hrs @ home (yng) - 6m	76.33	57.97	101.38	79.08
Hrs @ home (old) - 6m	101.07	88.99	101.06	94.65

C.7 Hospital Bed Constraints

The American Hospital Association reports the existence of 84,555 intensive care beds excluding neonatal units, implying a per-capita number of $Z = 0.032607\%$ ICUs.⁴⁰ We set $\tilde{\delta}_2(a) = 1$, so that the death for an agent in need of hospitalization and that does not receive a bed is certain to occur. The results are provided in Table C6.

C.8 Details on Calibration based on Ferguson et al. (2020)

We take the age-dependent transition rates from Ferguson et al. (2020) and aggregate to our two age groups (20-64 and 65-plus) using the relative population weights for the US. This yields the following weekly probability of recovery from mild symptoms: $\phi(0, y) = 0.991$ and $\phi(0, o) = 0.893$. The age-dependent weekly death rates are: $\delta(y) = 0.373$ and $\delta(o) = 0.371$. Table C7 compares these parameters with the ones from the benchmark calibration. Table C8 provides the results for this calibration.

⁴⁰See <https://www.aha.org/statistics/fast-facts-us-hospitals>

Table C7: Parameters - baseline calibration and Ferguson et al. (2020)

Parameter	Bench. calibration	Ferguson et al. (2020)
$\phi(0, y)$	0.983	0.991
$\phi(0, o)$	0.954	0.893
$\phi(1, y)$	0.284	0.284
$\phi(1, o)$	0.284	0.284
$\delta(y)$	0.065	0.373
$\delta(o)$	0.738	0.371

Table C8: Alternative calibration - based on Ferguson et al. (2020)

	Benchmark	Epidem.
Wks to peak hospitalizations (yng)	13.00	12.00
Wks to peak hospitalizations (old)	11.00	12.00
Hospitalizations p/ 1,000 @ peak (yng)	0.66	3.42
Hospitalizations p/ 1,000 @ peak (old)	0.27	13.24
Dead p/ 1,000 1year (yng)	0.63	1.05
Dead p/ 1,000 1year (old)	4.95	36.53
Dead p/ 1,000 1year (all)	1.55	8.65
Dead p/ 1,000 LR (yng)	0.71	1.05
Dead p/ 1,000 LR (old)	6.66	36.53
Dead p/ 1,000 LR (all)	1.98	8.65
Immune in LR (yng), %	57.14	85.22
Immune in LR (old), %	8.38	44.84
Immune in LR (all), %	46.70	76.58
GDP at peak - rel to BM	1.00	1.09
GDP 1year - rel to BM	1.00	1.03
Hrs @ home (yng) - peak	70.70	57.97
Hrs @ home (old) - peak	109.39	88.99
Hrs @ home (yng) - 6m	65.78	57.97
Hrs @ home (old) - 6m	106.68	88.99