
# Optimal Targeted Lockdowns in a Multi-Group SIR Model

Daron Acemoglu, Victor Chernozhukovz, Iván Werning, Michael D. Whinston

**ERSA working paper 826**

**July 2020**

# Optimal Targeted Lockdowns in a Multi-Group SIR Model*

Daron Acemoglu†
Victor Chernozhukov‡
Iván Werning §
Michael D. Whinston¶

May 2020

We study targeted lockdowns in a multi-group SIR model where infection, hospitalization and fatality rates vary between groups—in particular between the "young", "the middle-aged" and the "old". Our model enables a tractable quantitative analysis of optimal policy. For baseline parameter values for the COVID-19 pandemic applied to the US, we find that optimal policies differentially targeting risk/age groups significantly outperform optimal uniform policies and most of the gains can be realized by having stricter lockdown policies on the oldest group. Intuitively, a strict and long lockdown for the most vulnerable group both reduces infections and enables less strict lockdowns for the lower-risk groups. We also study the impacts of group distancing, testing and contract tracing, the matching technology and the expected arrival time of a vaccine on optimal policies. Overall, targeted policies that are combined with measures that reduce interactions between groups and increase testing and isolation of the infected can minimize both economic losses and deaths in our model.

*Rebekah Anne Dix and Tishara Garg provided excellent research assistance. For useful conversations, comments and suggestions we thank Fernando Alvarez, Alyssa Bilinski, Samantha Burn, Arup Chakraborty, Joe Doyle, Glenn Ellison, Zeke Emanuel, Ruth Faden, Eli Fenichel, Michael Greenstone, Simon Johnson, Angela McLean, Jessica Metcalf, Simon Mongey, Robert Shimer, and Alex Wolitzky. We thank Sang Seung Yi for providing us with the Korean case and mortality data. All remaining errors are our own.

†MIT Economics Department
‡MIT Economics Department
§MIT Economics Department
¶MIT Economics Department and Sloan School of Management

# 1 Introduction[1]

The principle of targeting plays an important role in economic analyses of government policy. Applying this well-respected principle is another matter, one that requires showing substantial benefits on a case-by-case basis. In many epidemics, the risk of infection or serious health complications varies greatly between different demographic groups. The cost of preventing economic activity through lockdowns is also typically heterogeneous within the population. The COVID-19 pandemic, which has claimed the lives of more than 360,000 people worldwide (as of May 29, 2020) and led to the largest global recession of the last nine decades, is no exception. It is distinguished by a very steep mortality risk with respect to age: for those over 65 years of age mortality from infection is about 60 times that of those aged 20-49. Differences of this magnitude merit examining the benefits of targeted policies.

In this paper we develop a multi-group version of the epidemiological SIR population-based model and undertake a quantitative analysis applied to COVID-19.[2] We focus on identifying the benefits arising from optimal targeted policies that lock down the various groups differentially. To do so, we solve an optimal control problem and examine how the possibility of targeting improves the tradeoff between lives lost and economic losses. We find that the benefits of targeting are significant. We believe the model and analysis we develop could apply to the study of future pandemics the world might need to prepare for.

We start with the special case of our model consisting of three groups—young (20-49), middle-aged (50-64) and old (65+) and where the only differences in interactions between the three groups come from differential lockdown policies. We base our parameter choices on the COVID-19 pandemic and characterize different types of optimal policies. Consistent with other works on the pandemic, when the menu is restricted to

[1] Visit our online MR-SIR simulator (GUI) to explore the model, the effect of simple policies, and optimal policies, for various parameter values: https://mr-sir.herokuapp.com/

[2] The SIR model was originally proposed by Kermack, McKendrick and Walker (1927) and remains an important reference in the epidemiological literature. It belongs to a wider class of deterministic population-based compartmental models that are a workhorse in epidemiological analyses, alongside agent-based and network models. The benefit of the SIR framework and its extensions is that it is relatively tractable, and amenable to our optimal control analysis.

Extensions of the model that allow for differences across groups are referred to in varying ways in the epidemiological literature, such as multi-group, and when focusing on age, age-structured or age-stratified. We prefer to refer to our model as a "multi-group" one because this terminology is familiar for both economists and epidemiologists, and because we believe that the general principles our analysis elucidates are applicable beyond the heterogeneities across age groups that arise in the context of COVID-19.

Other important extensions of this framework include SEIR models that allow for a non-infectious exposed (E) state, which is relevant for COVID-19. We use the benchmark SIR model as our baseline and show that our results are very similar with a similar multi-group version of the SEIR model.

uniform policies that treat all groups symmetrically, there are difficult trade-offs facing policy-makers. When the priority is to save lives (a "safety-focused" approach), the economy will have to endure a lengthy lockdown and sizable declines in GDP. For example, in order to keep the mortality rate in the (adult) population below 0.2%,[3] policy-makers will have to impose a full or partial lockdown of the economy for almost one year and a half and put up with economic costs equivalent to as much as 37% of one year's GDP. Conversely, policy-makers prioritizing the economy (employing an "economy-focused approach) and attempting to keep economic damages to less than 10% of one year's GDP may be forced to put up with a mortality rate over 1%.

Our main result is that this policy trade-off can be significantly improved with targeted policies that apply differential lockdowns on the various risk groups. To make this point, we focus on the ("Pareto") frontier between economic loss and loss of life, which represents the aforementioned trade-off facing policy-makers and is depicted by the solid curve in Figure 1.1. The frontier is upward sloping after a certain point, indicating that the absence of any mitigation policies will lead to both greater economic loss and more lives lost. This is because economic damages include lost productivity due to illness and the forgone productivity contributions of those who die because of the virus. Most importantly, the dashed frontier for targeted policies is much closer to the "bliss point" represented by the origin than is the frontier for uniform policies. This figure in addition helps us understand why targeted policies can save a significant number of lives—moving horizontally from the uniform policy frontier to the targeted policy frontier keeps the economic loss the same but substantially reduces fatalities. For example, we show that compared to the economy-focused uniform policy, targeting enables mortality in the (adult) population to be reduced from above 1% to around 0.5%, saving over 1.3 million lives compared to the benchmark of optimal uniform policy. Alternatively, with the safety-focused objective of 0.02% mortality, targeting reduces economic damages from around 37% to 25%. Naturally, the exact gains depend on the initial reference point on the frontier and whether the gains from moving to targeted policies are taken as a reductions in mortality or economic losses, or some combination. Our approach based on comparing entire frontiers has the benefit of sidestepping the difficult choice of a particular point on the frontier.[4]

We will also see that, for our COVID-19-based parameters, almost all of the gains from

[3] Throughout, by "population" we refer to the adult population (over 20 years old).

[4] A common alternative is to pick a point along the frontier by solving for the optimal policy using a "value of life" parameter that trades off deaths versus economic losses. However, the considerable disagreement about the right value of life, or even about the validity of this concept in practice, complicates applications of this approach.

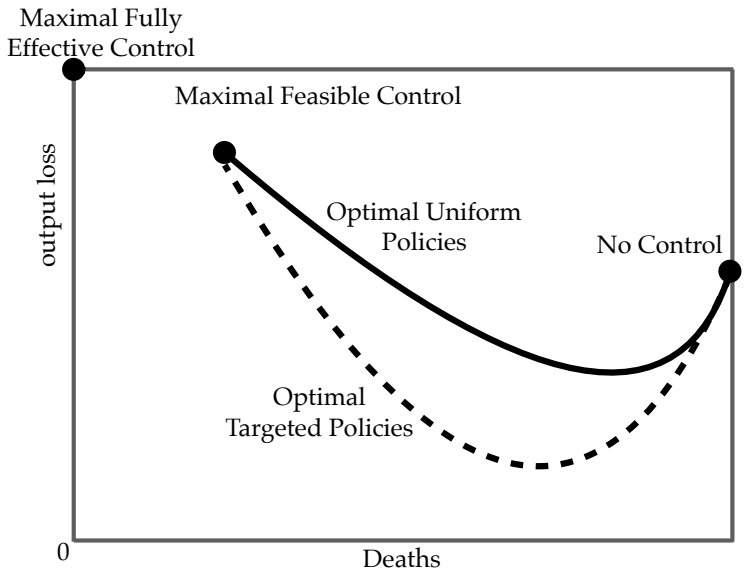


**Figure 1.1:** Frontier: economic vs. lives lost.

targeting can be achieved without the need to resort to complicated targeting policies. Rather, a "semi-targeted" policy that simply treats the most vulnerable (older) age group differently than the rest of the population performs nearly as well as "fully-targeted" policies (which also treat the young and the middle-aged differentially).

Our model also enables an analysis of a richer menu of policy options. One promising set of policies (in practice implemented both by explicit regulations and norms) is those that reduce interactions between the most vulnerable age group and the rest of the population. These policies, which we call "group distancing", turn out to be very powerful in reducing mortalities, because they complement targeted lockdowns. In our model, the mortality rate of the older age group can still be relatively high even under optimal targeted policies because, as in the real world, lockdowns are imperfect and older individuals still come into contact with and become infected by the young and the middle aged.[5] Group distancing partially rectifies this and enables both better health outcomes and shorter lockdowns (since, when the older group is further isolated from the younger ones, the latter's lockdown can be eased more rapidly). For example, with group distancing and semi-targeted policies, economic losses resulting from the safety-focused objective of no more than 0.2% mortality rate can be reduced to about 16% of one year's GDP (or, alternatively, mortality can be reduced further).

[5]This problem is highlighted by the plight of the nursing home populations in the United States, who were exposed to the virus via visitors and staff, and to date account for more than one-third of all COVID-19-related deaths in the country (e.g. https://www.nytimes.com/interactive/2020/05/09/us/coronavirus-cases-nursing-homes-us.html).

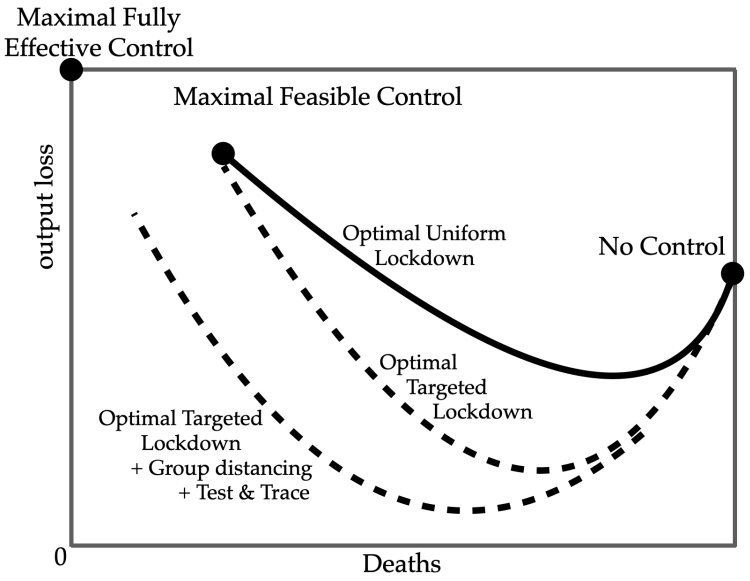


**Figure 1.2:** Frontier: economic vs. lives lost with additional policies.

Another set of policies that significantly improve the trade-offs facing policy-makers is testing and (contact) tracing. By identifying and isolating infected individuals, these policies provide better protection against the virus and lead to much lower economic and public health costs, especially when combined with semi-targeted policies. Interestingly, when both group distancing and testing-tracing policies are adopted and we consider semi-targeted optimal policies, the trade-off between lives lost and economic damages improves radically as illustrated Figure 1.2. For example, policy-makers can keep economic damages to about 7% of one year's GDP while achieving a 0.2% mortality.

We also investigate the implications of the returns to scale in the"matching technology" between susceptible and infected individuals for the dynamics of the pandemic and optimal policy. At one extreme one may assume a quadratic specification, which implies that new infections depend on the product of susceptible and infected populations (called "density-dependent" or "mass action" in the epidemiological literature, e.g., McCallum et al. (2001)). At the other extreme, one may postulate a constant returns specification, which captures the notion that each person will be involved with some fixed number of interactions and what matters is the fraction of infected people they contact (called frequency-dependent in the epidemiological literature).With a constant returns to scale matching technology, the recovered offer greater protection to the susceptibles, but counteracting this, lockdowns themselves are less effective (because the "double benefit" generated by lockdowns under the quadratic matching technology is weakened). Despite these differences, we find that optimal semi-targeted policies are broadly similar

under different matching technologies and continue to significantly outperform optimal uniform policies.

We stress that there is much uncertainty about many of the key parameters for COVID-19 (Manski and Molinari, 2020) and any optimal policy, whether uniform or not, will be highly sensitive to these parameters (e.g., Atkeson, 2020a, Avery et al., 2020, Stock, 2020). Nonetheless, while the specific numbers on economic and public health costs are sensitive to parameter values, our general conclusion that targeted policies bring sizable benefits appear very robust (as we document as well).

Within the incipient economics-epidemiological literature, Atkeson (2020b) and Stock (2020) provide an introduction to the SIR framework and its implications for COVID-19 in the US. Fernández-Villaverde and Jones (2020) fit a standard SIR model to multiple regions (countries, states and cities) and uses the model to infer unobservables (such as number of recovered) and create forecasts. Closer to our paper, a number of recent papers have started incorporating economic trade-offs and conducting optimal policy analysis within the SIR framework (e.g. Rowthorn and Toxvaerd, 2020, Eichenbaum, Rebelo and Trabandt 2020a, Alvarez, Argente and Lippi 2020, Jones, Philippon and Venkateswaran, 2020, Farboodi, Jarosch and Shimer, 2020 and Garriga et al., 2020).[6] All of the papers undertaking an optimal control analysis have worked with single-group models.

Several recent papers independently investigate the role of age-dependent hospitalization and fatality rates in SIR models (Gollier, 2020, Favero, Ichino and Rustichini, 2020, Rampini, 2020, Bairoliya and İmrohoroğlu, 2020, Brotherhood et al. (2020) and Glover et al. (2020)). The main differences between these papers and ours are: (1) our general treatment of dynamics of infection in an SIR model with multiple risk groups, different interaction structures and potentially imperfect testing and tracing, and (2) more importantly, our analysis of optimal policy. For example, our results showing that semi-targeted policies can significantly improve over optimal uniform policies and achieve the great majority of the gains of optimal fully-targeted policies have no counterparts in these papers. Brotherhood, Kircher, Santos and Tertilt (2020) and Glover, Heathcote, Krueger and Ríos-Rull (2020), in particular, study infection and economic dynamics in settings with additional economic choices (labor supply and consumption choices under incomplete information about infection status in the former, and sectoral choice in the latter). Their focus and main results are different and complementary to ours. For example, Brotherhood et al. (2020) focus on younger individuals' risk-taking behavior and the implications

[6] In addition, Eichenbaum, Rebelo and Trabandt (2020a), Jones, Philippon and Venkateswaran (2020), Farboodi, Jarosch and Shimer (2020), Kudlyak, Smith and Wilson (2020) and Garibaldi, Moen and Pissarides (2020) are recent papers endogenizing economic behavior in basic SIR models. Early related contributions include Geoffard and Philipson (1996) and Fenichel (2013).

of this for testing and conditional quarantining, while Glover et al. (2020)'s main emphasis is on the conflict between the young and the old about mitigation policies. Although Glover et al. (2020) consider optimal policy, this policy is chosen from a parametric family and their main emphasis is on the contrast between this policy and those preferred by the young and the old. Finally, recent work by Baqaee et al. (2020) uses a model where policy is targeted according to age and sector to investigate alternative reopening scenarios (but consider on policies where policy-makers link activity to the unemployment rate and whether deaths are rising or high).

The rest of the paper is organized as follows. The next section outlines the main elements of our multi-group SIR model, presenting the continuous-time laws of motion for infectious, susceptible and recovered populations by group, as well as the economic and mortality outcomes that our model with lockdown policies can generate. Section 4 describes our parameter choices and numerical methods. Our main results are presented in Section 5, which also contains a number of robustness exercises. Section 6 contains our conclusions.

## 2 Multi-group SIR model

Our multi-group SIR model is set in continuous time $t \in [0, \infty)$. Individuals are partitioned into groups $j = 1, \ldots, J$ with $N_j$ initial members.[7] The total population is normalized to unity so that $\sum_j N_j = 1$.

At any point in time $t$, individuals in group $j$ are subdivided into those susceptible (S), those infected (I), those recovered (R) and those deceased (D),

$$S_j(t) + I_j(t) + R_j(t) + D_j(t) = N_j.$$

Agents move from susceptible to infected, then either recover or die.[8,9] We write $S(t) = \{S_j(t)\}_j$ and similarly for $I(t)$, $R(t)$ and $D(t)$. Groups interact with themselves as well as with each other, as described below.

Before describing the details, we anticipate one of our key equations. In the canonical

[7] See Heesterbeek and Roberts (2007) and Bayham, Kuminoff, Gunn and Fenichel (2015) and the references therein for a discussion of age- or stage-structured compartmental epidemiological models.

[8] Given relatively short time horizon of this pandemic, we abstract from other sources of deaths and births.

[9] The most common extensions of the SIR model include either the "Exposed" who have not yet turned infectious (SEIR) or the "Asymptomatic" who may still infect others (SAIR). See, for example, Brauer, van den Driessche and Wu (2008). We believe our main conclusions and results are robust to such extensions and plan to investigate this in the future.

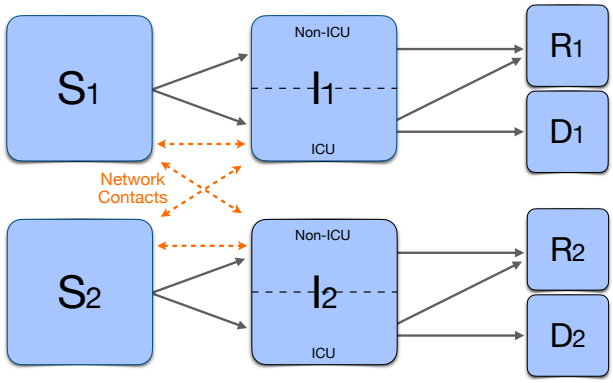


**Figure 2.1:** MG-SIR: Multiple-Risk Susceptible Infected Recovered Model. Solid lines show the flows from one state to another. Dashed lines emphasize interactions that take place across risk groups.

single-group model, the key evolution equation is quadratic:

$$\text{new infections} = \beta SI.$$

In our model, absent lockdowns and isolations, we have

$$\text{new infections in group } j = \beta S_j \frac{\sum_k \rho_{jk} I_k}{\left(\sum_k \rho_{jk}(S_k + I_k + R_k)\right)^{2-\alpha}},$$

where $\{\rho_{jk}\}$ are parameters that control the contact rate between group $j$ and $k$. Here $\alpha \in [1,2]$ allows us to control the returns to scale in matching: when $\alpha = 1$ we have constant returns: infections double if $S$, $I$ and $R$ double; when $\alpha = 2$ we obtain the quadratic specification that, with a single group, boils down to the canonical SIR model. Below we develop the full model, complementing and extending this basic equation to include testing, isolation, lockdowns, hospital capacity, the arrival of a vaccine and other considerations.

## 2.1 Model Assumptions

Here we discuss the basic elements of our model and then turn to the dynamic equations describing the evolution of the state variables.

**Infection, ICU, Fatality and Recovery.** Susceptible individuals may become infected by coming into contact with infected individuals. Those infected may or may not require

"ICU care", a catch all label that we use to capture the need for ventilators and other specialized medical care. We suppose that the need for ICU care is immediately realized upon infection. Let $\iota_j$ denote the constant fraction of infected people of type $j$ needing ICU care. With Poisson arrival $\delta_j^r$ an ICU patient of type $j$ recovers. Non-ICU patients do not die, and recover with Poisson arrival $\gamma_j$.

Only those needing ICU care may die. While in the ICU, patients die with Poisson arrival $\delta_j^d(t)$. We allow this death rate to be a function of total ICU needs relative to capacity, which may vary over time. We assume that

$$\gamma_j = \delta_j^d(t) + \delta_j^r(t).$$

These assumptions imply that the proportions of ICU and non-ICU patients among the infected in group $j$ do not change over time.[10]

Let $H_j(t)$ denote the number of type $j$ individuals needing ICU care at time $t$, so that $H_j(t) = \iota_j I_j(t)$. Total ICU needs are $H(t) = \sum_j H_j(t)$. We assume that the death probability conditional on ICU is a non-decreasing function of the number of patients

$$\delta_j^d(t) = \psi_j(H(t)),$$

for some given nondecreasing function $\psi_j$.[11]

**Detection: Infection and Immunity.** Detection of infected individuals as well as their isolation is assumed to be imperfect. To avoid additional state variables, for each infected individual it is determined immediately upon infection whether detection and isolation is possible. We denote by $\tau_j$ the constant probability that an infected individual of type $j$ not needing ICU care becomes isolated. Isolation may occur merely because the individual becomes too sick to work and socialize, perhaps needing hospitalization short of the ICU. It may also occur because the individual is detected as infected by showing symptoms or by testing and tracing efforts. Similarly, we let $\phi_j$ denote the probability that an individual of type $j$ needing ICU care is detected and isolated. While it may be reasonable to set $\phi_j = 1$, we allow $\phi_j < 1$ to nest earlier work, such as Alvarez et al. (2020), that did not model ICU needs explicitly (by setting $\tau_j = \phi_j$). Summing up, the probability that an

[10] Relaxing these assumptions is possible, but requires an additional state variable to keep track of ICU vs. non-ICU infected individuals.

[11] One may extend the model to incorporate capacity that increases over time by allowing $\psi_j(H(t), t)$ to depend on $t$. This could capture the efforts undertaken at the start of the epidemic to improve medical resources for COVID-19 patients.

infected person fails to be detected and isolated is given by

$$\eta_j \equiv 1 - (\iota_j \phi_j + (1 - \iota_j)\tau_j).$$

We assume that recovered agents are immune, at least for the remaining duration of the pandemic, until the vaccine and cure arrive.[12] Individuals that recover can be released from lockdown at no risk to themselves or others. However, due to imperfect testing, we suppose that only a fraction $\kappa_j$ of recovered agents are identified as such and allowed to work freely. The remaining fraction is not identified andthus treated identically to those that are not deemed recovered. Once again, to simplify we assume this status is determined immediately after recovery.[13]

**Lockdown and Social Distancing.** We now consider lockdown and social distancing measures that affect the rate of transmission of infections. To simplify our discussion we label all of these as "lockdown" policies, although they could represent other social distancing behavior and policies as well.

Individuals in group $j$ produce $w_j$ when they are not in lockdown and $\xi_j w_j$ during lockdown, where $\xi_j \in [0, 1]$. The parameter $\xi_j$ captures the relative productivity of home vs. market production. If $\xi_j > 0$, then some productive work can be done from home.

A full lockdown creates a loss for each member of group $j$ equal to $(1 - \xi_j)w_j$.[14] For each group, lockdowns are not necessarily all or nothing. Let $L_j(t)$ denote the share of group $j$ that is in "lockdown" with $L_j(t) = 1$ designating full lockdown and $L_j(t) = 0$ no lockdown. Intermediate values of $L_j(t) \in (0, 1)$ represent less extreme situations. One interpretation is that at each instant $t$ a random fraction $L_j(t)$ is drawn independently (of all past events) from group $j$ and sent into lockdown.[15]

[12] Although many experts agree that this is currently the leading hypothesis for COVID-19, it is important to stress that, due to the recency of the pandemic, at this time this hypothesis is not backed by conclusive evidence. Indeed, the World Health Organization has stated "There is currently no evidence that people who have recovered from COVID-19 and have antibodies are protected from a second infection." (April 24, 2020, briefing: https://www.who.int/news-room/commentaries/detail/immunity-passports-in-the-context-of-covid-19).

[13] If the only constraint were detection, it might be reasonable to suppose $\kappa_j \in [\iota_j \cdot \phi_j + (1 - \iota_j)\tau_j, 1]$ capturing the notion that we do not forget those that were identified as being infected. However, we do not require this condition.

[14] Note that $w_j$ summarizes both the wage and employment level of group $j$. In the data, as people age wages rise but labor market participation falls. More generally, $w_j$ may also capture more than production, such as utility benefits experienced by consumers or leisure outside of lockdown. We focus on the narrow interpretation of this variable to facilitate our quantitative analysis, though we return to this issue in our robustness checks.

[15] This implies that intermediate lockdowns are not selecting the same workers to be locked down persistently. These types of policies can be incorporated into our framework by splitting identical workers into

Full lockdown may not be feasible, so we impose $L_j(t) \leq \bar{L}_j \leq 1$. This may capture that some goods are deemed essential, so their production cannot be shut down. Alternatively, it may be hard to monitor and prevent some people from going to their workplace. We also assume that even if full lockdown were feasible, so that nobody could work, it would not necessarily eliminate all human interactions and contagion. In particular, we assume that lockdown $L_j(t)$ reduces actual work by $L_j(t)$, but reduces the presence of group $j$ in infectious interactions only by a factor $1 - \theta_j L_j(t)$ where $\theta_j \leq 1$. This may be because people are still allowed on the streets and transmission occurs when they cross paths, or it may be because people disobey lockdowns or quarantines, or cheat at the margin by visiting each other socially. It may also capture transmission that occurs without direct person-to-person contact, as when someone touches an object recently touched by an infected person. Basically, for a number of reasons, lockdown is not fully effective at removing transmission and $\theta_j < 1$ is a measure of this inefficiency.

**Vaccine and Cure.** We assume a vaccine becomes available at some date $T$. To simplify, we assume that a cure for all those currently infected also becomes available at the same date.[16] In the case of COVID-19, experts currently estimate the time for developing, testing and rolling out a vaccine to be between one and two years, although extraordinary efforts are underway to speed up vaccine developments.

## 2.2 Dynamics in MG-SIR

Before the vaccine and cure, infections for group $j$ evolve according to the differential equation for all $t \in (0, T)$

$$\dot{I}_j = M_j(S, I, R, L)\beta(1 - \theta_j L_j)S_j \sum_k \rho_{jk}\eta_k(1 - \theta_k L_k)I_k - \gamma_j I_j,$$

for nonnegative $\beta$ and contact coefficients $\{\rho_{jk}\}$ and where

$$M_j(S, I, R, L) \equiv \left( \sum_k \rho_{jk}[(S_k + \eta_k I_k + (1 - \kappa_k)R_k)(1 - \theta_j L_k) + \kappa_k R_k] \right)^{\alpha - 2}.$$

different groups that can be treated differently.

[16] Anti-viral drugs to treat an infection and vaccines involve two different medical advances and implementations. In practice, the possibility of a vaccine being discovered in the absence of a cure makes little difference for optimal policy as long as the infections at date $T$ will be low, so that the value of a cure for the currently infected is small relative to the benefit of preventing further infections. Discovery of a perfect and immediate cure would, in principle, be as good as having a vaccine.

Note that $M_j(S, I, R, L) = 1$ in the quadratic case in which $\alpha = 2$.

The rest of the laws of motion for $t \in (0, T)$ are

$$\begin{aligned}
\dot{S}_j &= -\dot{I}_j - \gamma_j I_j, \\
\dot{D}_j &= \delta_j^d(t) H_j, \\
\dot{R}_j &= \delta_j^r(t) H_j + \gamma_j (I_j - H_j),
\end{aligned}$$

where, again, $H_j = \iota_j I_j$ denote the number of ICU patients in group $j$.

After the vaccine and cure arrive at $T$, every individual that is alive is placed in the recovered category: $S(t) = I(t) = 0$ and $R(t) = S(T_-) + R(T_-)$ and $D(t) = D(T_-)$ for $t \geq T$. Naturally, we set lockdowns to zero $L(t) = 0$ for $t \geq T$.

**Discussion.** To gain further insight into the law of motion for infections, consider the single group case and set $\alpha = 2$ and $\rho = \eta = 1$. Then we obtain

$$\dot{I}_j = \beta S I (1 - \theta L)^2 - \gamma I,$$

which is the standard quadratic SIR specification.

With multiple groups, the coefficients $\{\rho_{jk}\}$ allow for different contact rates across groups. For instance, it may be natural to suppose that $\rho_{jj} > \rho_{jk}$ for $k \neq j$, so that individuals of a given group tend to have a higher contact rate within their own group (i.e., the young may work and socialize with each other more than with the old).

Although the majority of the economic-epidemiology literature focuses on the quadratic matching case with $\alpha = 2$, we allow a more general formulation. One way to see the difference is to think about scale. The matching technology with $\alpha = 1$ exhibits constant returns to scale: doubling $S$, $I$ and $R$ (while fixing the lockdown policy $L$) doubles the number of infections, leaving the rate of growth in infections unchanged. In contrast, with quadratic matching, it would quadruple it. Intuitively, the constant returns case with $\alpha = 1$ assumes that the number of contacts for each individual are not affected when the total number of people not in lockdown doubles. This is referred to as "frequency dependence" in epidemiology. In contrast, the quadratic case with $\alpha = 2$ assumes the number of contacts doubles. The quadratic matching technology assumes that a more densely occupied environment leads to more contacts in strict proportion to density. This is referred to as "density dependence" in epidemiology.

In the absence of lockdowns, the differences between $\alpha = 2$ and $\alpha = 1$ are small in practice. This is because, we can adjust the value of $\beta$ to compensate for initial size

differences in population. Then, as long as mortality is never too high, so that the value of those alive $S + I + R$ is not too affected, the two specifications behave very closely. However, in the presence of lockdowns there are nontrivial differences between $\alpha = 2$ and $\alpha = 1$. To see this, consider as an example the single group case again, but this time with $\alpha = 1$:

$$\dot{I}_j = \beta S I \frac{1 - \theta L}{S + I + R + \frac{\theta L}{1-\theta L}\kappa R} - \gamma I$$

Setting $\kappa = 0$ and $S + I + R \approx 1$ first, we see that lockdowns affect new infections by a factor $1 - \theta L$ instead of the quadratic $(1 - \theta L)^2$. Thus, halving the population interacting by way of lockdowns, i.e., $(1 - \theta L) = \frac{1}{2}$, has the effect of reducing new infections by a forth with $\alpha = 2$ and by a half with $\alpha = 1$. Turning to the denominator, when $\kappa > 0$, a further difference emerges due to due to the presence of $\frac{\theta L}{1-\theta L}\kappa R$. With non-zero lockdown $L > 0$ and fixing $S + I + R$, we see that larger numbers of recovered individuals $R$ contribute to lower infections, which is a strong type of herd immunity effect.

In practice, some social interactions are closer to the quadratic kind, while others to constant returns. Walking along a street or shopping a supermarket generate interactions that do increase with density. However, interactions at work or with friends may be closer to constant returns to scale, since each person will make time and seek out a given number of co-workers or friends to interact closely with.

**An Aggregation Result.** Our MG-SIR model displays a useful aggregation property, behaving like a single group SIR model in special cases when lockdowns are uniform.

Suppose that effective contact rates and resolution rates out of infection the same across groups, so that $\beta_{jk} = \beta$ and $\gamma_j = \gamma$, and consider lockdown policies that are uniform across groups as well, so that $L_j(t) = L(t)$ for all $j$. Suppose further that infection rates are initially identical across groups, so that $S_j(0)/N_j$, $I_j(0)/N_j$ and $R_j(0)/N_j$ are independent of $j$. Then it is straightforward to see that, despite differences in fatality rates, the evolution of infections within each group, and hence aggregate infections, is identical to that of a single group SIR model. The same is not true for deaths—these are different across groups, but do not affect the evolution of infections.

This aggregation result is verified in our simulations when we present uniform optimal lockdowns.

# 3 Efficient Frontier

Our planner controls lockdown for each group $\{L_j(t)\}_j$ for all $t \in [0, T)$ taking into account the effect this has on the dynamical system described above. We will consider cases where lockdown is restricted to be uniform across groups $L_j(t) = L(t)$ as well as cases where it is differential across groups.

There are two objectives to consider, and the trade-off between the two gives our frontier. To simplify the exposition we ignore discounting.[17] The first objective is to minimize total (excess) deaths during the pandemic, which is simply

$$\text{Lives Lost} = \sum_j D_j(T).$$

Recall that there will be no (excess) deaths after the arrival of the vaccine and cure, that is for $t > T$.[18]

The second goal is to minimize the total economic loss, which we represent as

$$\text{Economic Losses} = \int_0^T \sum_j \Psi_j(t)\, dt,$$

where the flow cost for group $j$ is given by

$$\begin{aligned}\Psi_j(t) = (1-\xi_j)w_j S_j(t)L_j(t) + (1-\xi_j)w_j I_j(t)(1-\eta_k(1-L_j(t))) \\ + (1-\xi_j)w_j(1-\kappa_j)R_j(t)L_j(t) + w_j\Delta_j\iota_j\delta_j^d(t)I_j(t)\end{aligned}$$

The first term represents the economic loss of lockdown from susceptible individuals.

[17] Discounting can be easily added, but note that nominal interest rates were reduced to near zero in April 2020 so this has no significant effects on our results.

[18] Allowing the vaccine arrival time to be stochastic is simple. Suppose $T$ has cumulative distribution $F(T)$ and assume there is no information about the vaccine's time of arrival before its actual arrival. The first objective is the expected number of deaths, which can be integrated by parts to yield,

$$\int_0^\infty f(T)\left(\sum_j D_j(T)\right)dT = \int_0^\infty (1-F(t))\left(\sum_j \iota_j\delta_j^d(t)I_j(t)\right)dt$$

Similarly, the economic objective becomes

$$\int_0^\infty (1-F(t))\sum_j \Psi_j(t)dt$$

Previously $1-F(t) = 1$ for $t < T$ and $1-F(t) = 0$ for $t \geq T$. A convenient choice assumed in some papers (e.g. Alvarez et al., 2020; Eichenbaum et al., 2020b) is a Poisson distribution with constant arrival rate.

The second term is the economic loss if lockdown from infected individuals, which now captures the fact that even in the absence of a lockdown ($L_j(t) = 0$) a fraction $1 - \eta_j$ finds itself not working because it is in isolation. The third term represents the lockdown cost from recovered individuals, which is zero if all recovered agents are immediately sent back to work, i.e., $\kappa_j = 1$. The last term represents the economic cost of deaths in group $j$, given by the product of the death flow rate $\iota_j \delta_j^d(t) I_j(t)$ and the economic loss $w_j \Delta_j$ per death, where $\Delta_j$ captures the present discounted value of a group $j$ memeber's remaining employment time until retirement that is lost due to death.[19]

**Frontier and Value of Life.** Our approach is to characterize the frontier between economic losses and lives. A common alternative is to weight the latter by some value, say $\chi$ and minimize

$$\int_0^T \sum_j \Psi_j(t)\, dt + \chi \sum_j D_j(T).$$

To see the relation with our approach, note that for any $\chi$, minimizing this expression will lead to a particular point on the frontier. As long as the frontier is convex, one can trace out the entire frontier by varying $\chi$. Geometrically, each value of $\chi$ corresponds to a point on the frontier and $\chi$ equals the the tangent slope at this point.

Under this interpretation, $w_j \Delta_j + \chi$ denotes the total value of life, or penalty for death, comprised of an economic component $w_j \Delta_j$ and a non-pecuniary component $\chi$.

The main distinction between our approach and one centered around the weighted objective using $\chi$ is that we characterize the entire frontier and describe points on the frontier by the number of deaths and the economic losses, rather than invoking a value of life or mapping it to a parameter such as $\chi$.

This difference may seem subtle, but it offers a distinct perspective: we do not purport to determine an optimum, selecting a point along the frontier. Rather, by characterizing the frontier we describe the menu of difficult choices society (policy-makers) faces. Our results showing the benefits of targeting policies acquire the following interpretation: Suppose society is using uniform lockdowns and finds itself somewhere along the uniform policy frontier, or worse. Then our analysis can shed light on how to improve this situation, by either reducing the number of deaths or lowering the economic losses, or some combination. In contrast, an approach centered around a choice for the value of life or $\chi$ is less attractive, in part because there is great disagreement about the right values for this parameter.

[19] To simplify, we take $\Delta_j$ to be constant. An alternative is to let it depend on time of death, so that we have $\Delta_j(t)$. This would makes little difference for the range of values of $T$ we consider.

**Marginal Value of Vaccine Innovations.** Suppose we have a marginal improvement in the arrival of the vaccine. From $T$ to $T - dT$ with $dT > 0$ denoting an improvement. Applying the Envelope Theorem, we compute the marginal change in the two objectives to be, deaths and economic losses:[20]

$$-dT\iota_j\delta_j^d(t)I_j(t) \qquad \text{and} \qquad -dT\sum_j \Psi_j(T).$$

As usual, the planner will react to a change in $T$ to achieve a new optimum, but the Envelope Theorem allows us to evaluate the marginal value of small changes without computing the new solution. Intuitively, when the vaccine arrives earlier, the distribution of $F(t)$ shifts up, lowering the cost because the future flow costs from lockdown, represented by integrand, are no longer incurred.[21]

From the above calculations, the marginal value of an earlier vaccine/cure depends crucially on lockdown and infection rates. In particular, if both are vanishingly small at $T$, then there is only a vanishingly small improvement. For example, if the original solution was "going for herd immunity" (see below) so that lockdowns were zero and infections were near zero at $T$, then a slightly earlier vaccine arrival is of little value.

**Many Herd Immunities and Waiting for the Vaccine.** Although the exact form of the optimal lockdown is complex, it is useful to discuss some broad concepts and aspects of the alternative strategies available to reduce fatalities and economic losses. First, one can either hold out for the vaccine or alternatively go for "herd immunity." Second, there are many ways to reach herd immunity and different policies can steer the pandemic toward different herd immunity outcomes. To simplify the discussion, we consider the two-group case and assume that the two groups, labeled the old and the young, have equal size with $\rho = \eta = 1$ and $\alpha = 2$.

Figure 3.1 shows, for this case, the time path for the pair $(S_y(t), S_o(t))$ over the course

[20] This calculation presumes the solution is continuous in the vaccine arrival c.d.f. $F$ at the original optimal point. Because the problem is not convex, solutions may be discontinuous. If the solution is discontinuous, a similar formula holds for the directional derivative.

[21] In the stochastic arrival case, we have a distribution $F(t)$, and consider a change in this distribution $\delta F$, which changes deaths by

$$-\int_0^\infty \delta F(t)\left(\sum_j \iota_j\delta_j^d(t)I_j(t)\right)dt$$

and economic losses by

$$-\int_0^\infty \delta F(t)e^{-rt}\sum_j \Psi_j(t)dt$$

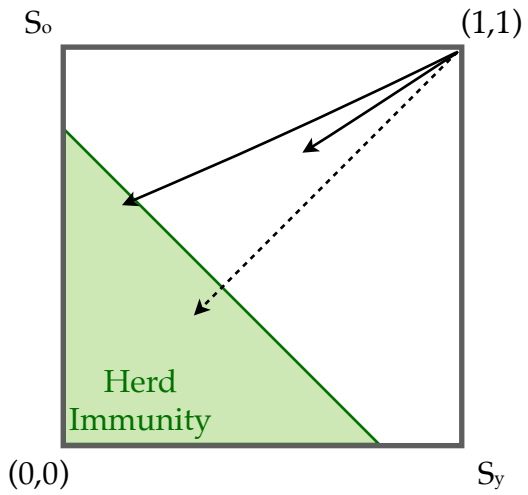


**Figure 3.1:** Illustrative herd immunity region and different time paths for the pandemic with two groups, old and young.

of the pandemic for $t \in [0, T]$ until the arrival of the vaccine. The pandemic starts near $(1, 1)$ with few infections and travels down and to the left, as more people get infected.

The shaded area represents the region of herd immunity, where the size of the susceptible population is sufficiently low that, once we enter this region, the pandemic comes to an end quickly (but not immediately as we discuss below).[22] In single-group SIR models, this region corresponds to an interval of the form $S \in [0, \bar{S}]$, and within this interval we have $\dot{I} < 0$. With multiple groups this same concept defines a region for the pair $(S_y, S_o)$. When $\rho = 1$ and the two groups have equal sizes, this region is symmetric, with slope $-1$, as shown in the figure.

Without any mitigation the disease follows the dashed 45-degree line, starting from an initial condition where almost nobody has been sick and reaching a situation where the majority in both groups have been infected at some point. In this case, the pandemic goes beyond the frontier for herd immunity—a phenomenon referred to as "overshooting" in the epidemiology literature. This occurs if there are a significant number of infections when crossing the threshold. Note that although the pandemic travels along the 45-degree line so that the same fraction of young and old get infected, mortality may be significantly higher for the old.

[22] More formally, we can define the region of herd immunity as the set of points $(s_y, s_0)$ with the property that, in the absence of lockdowns $L_y(t) = L_o(t) = 0$, the dynamic system starting from $(S_y(0), S_0(0)) = (S_y, S_0)$ and small initial infections $(\bar{I}_y(0), \bar{I}_o(0))$ converges to points near $(S_y, S_0)$. In other words $(S_y(\infty), S_0(\infty))$ is continuous in $(I_y(0), I_o(0))$ so that for the limit point $(S_y(\infty), S_0(\infty))$ we have $(S_y(\infty), S_0(\infty)) \to (S_y, S_0)$ as $(I_y(0), I_o(0)) \to (\bar{I}_y(0), \bar{I}_o(0))$. One can express this property as a condition on the largest (dominant) eigenvalue of the linearized dynamical system.

We can classify any policy by whether the path it induces reaches the herd immunity region before time $t = T$. Those that do not do so can be said to be "waiting for the vaccine", while those that reach this region are going for herd immunity.

Any uniform policy sends $(S_y, S_o)$ along the 45-degree line by virtue of our aggregation result. Uniform policies with sufficiently strict lockdowns will keep infections low enough to wait for the vaccine, and so will be on the segment of the dashed line outside of this region at time $t = T$. Other uniform policies may reach herd immunity but avoid a significant overshoot. Moreover, as our analysis below will reveal, even optimal policies that go for herd immunity will seek to "flatten the curve" in order to avoid the higher mortality due to limited hospital capacity.

More targeted mitigation policies open up new possibilities. The top solid line locks down the old more aggressively than the young, leading to lower infections among old relative to young. The path of the pandemic may then reach the region of herd immunity at an angle, with a higher fraction of infected among the young than the old, reducing the excess mortality coming from the old. With targeted policies, too, the planner may opt to hold out for the vaccine as with the lower solid line, but will also adopt policies that reduce infections among the old.

# 4 Specification and Calibration

In our analysis in the remainder of the paper, we focus on targeting policies based on age.[23] We consider a setting with three groups, the "young" (y) who are ages 20-49, the "middle-aged" (m) who are 50-64, and the "old" (o) who are 65 and above. We do not include those under 20 in our analysis. We take the population share of these three groups among those over 20 years of age from BLS population data for 2019, setting $N_y = 0.53$, $N_m = 0.26$, and $N_o = 0.21$. In our baseline, we assume equal earnings per capita for the young and middle-aged groups, which we normalize by setting $w_y = w_m = 1$.[24] Based on the fact that only 20% of those over 65 years of age are employed, and their average full-time weekly earnings, we set $w_o = 0.26$. We set $\xi = 0.3$, which implies that working from home results, on average, in a 70% loss of productivity. This number is chosen so that a full-year lockdown of all non-essential workers would lead to a roughly 50% decline in GDP. It is also in the ballpark of the numbers suggested in Dingel and Neiman (2020).

[23] Another factor that targeted policies could depend on is the presence of co-morbidities, which have been shown to lead to significantly higher mortality and ICU needs.

[24] From BLS statistics, the full-time employed middle-aged have 12% higher weekly earnings, but are 13% less likely to be employed than the young. The share of workers who are employed full-time versus part-time is roughly equal in the two groups.

As in Alvarez, Argente and Lippi (2020), in our baseline calibration we set $\bar{L} = 0.7$ when we consider uniform policies, reflecting the need for essential services, and set $\bar{L}_o = 1$ and $\bar{L}_j = 0.7$ for the other groups when considering targeted policies. In our baseline, we assume that there is a 10% chance that an infected individual is isolated, a number that is a few percentage points higher than the 6.6% probability of hospitalization for a COVID-19 case implied for the US by the assumptions in Ferguson et al. (2020). Thus, we set $\phi_j = \tau_j = 0.1$, implying that the probability of failing to isolate an infected individual is $\eta_j \equiv \eta = 0.9$ for all groups. (In Section 5.3, we consider the effect of decreases in $\eta$ thanks to identification and isolation of the infected by symptoms, testing, or tracing.) As in Alvarez et al. (2020), for our baseline, we assume that there is perfect identification of those individuals who have recovered and are allowed to go back to work avoiding lockdowns, so that $\kappa_j = 1$. This case is of interest because issuing such "immunity card" permits for work is clearly efficient in the model. However, it may not be realistic if testing is limited, as well as the fact that there are many ethical and practical issues to resolve. Thus, we will also consider the opposite assumption, so that the recovered are not identified or treated differently, and are subject to the same lockdowns as everyone else, $\kappa_j = 0$.

We choose $\gamma = 1/18$ so that a COVID-19 case reaches a conclusion, with the individual either recovering or dying, in 18 days on average.[25] We set $\beta$ equal to 0.134, reflecting a basic reproductive value of $R_0 = 2.4$ without social distancing and isolation measures,[26] which is the baseline value used in Ferguson et al. (2020).

For the nature of interactions we start by assuming that there is a single pool in which all of those who are not effectively locked down (share $(1 - \theta L_j)$ of each group $j$) interact. We set $\theta = 0.75$ in our baseline and examine lower values of $\theta$ in our robustness analysis. This value of $\theta$ implies that a full lockdown reduces interactions by 75%. For the contact matrix $\{\rho_{ij}\}$ we start with a conservative benchmark and assume $\rho_{ij} = 1$ for all $i, j$, so that all age groups interact equally with each other. This is not meant to be realistic, but it diminishes the benefit of targeting, for it implies that the old will be more exposed to the infected among the younger groups. In Section 5.6 we use data from Klepac et al. (2020), based on the BBC pandemic project, for interaction patterns across different age groups in the UK and show that this richer interaction structure has little effect on our

[25] Setting $\gamma = 1/5$ or $1/7$ to match the length of time during which an individual is infectious and then recalibrating $\beta$ so that $R_0 = 2.4$ as in our baseline leads to essentially identical results. In our robustness exercises in Section 5.6 we show that explicitly incorporating a distinction between the length of the infectious state and the total duration of the infection using an SEIR model leads to essentially identical results.

[26] $R_0 = \beta/\gamma$ can be thought of as the number of new infections generated in a day by a single infectious individual when the rest of the population is susceptible.

| Age Group | Fatality Rate |
|---|---|
| 20-49 | 0.001 |
| 50-64 | 0.01 |
| 65+ | 0.06 |

**Table 1:** Infection Fatality Rate from COVID-19.

main results.[27]

We take from Ferguson et al. (2020) the case fatality rates for the three age groups, conditional on infection and ICU services being available, which we denote by $(\overline{\delta}_y^d, \overline{\delta}_m^d, \overline{\delta}_o^d)$ and summarize in Table 1. For the young and middle-aged groups, these numbers closely match mortality rates we derived from recent South Korea data, a country with ample ICU capacity relative to needs and widespread testing (and thus hopefully relatively little selection of the more seriously ill among those tested).[28] For those over 70, however, the South Korean data give a higher fatality rate than that used by Ferguson et al. (2020). Given even lower fatality rates for older cohorts from the *Diamond Princess* cruise ship, we use mortality rates close to Ferguson et al. (2020) for those ages 65+ as our baseline and then verify the robustness of our results to Korean rates. To match the case fatality rates conditional on infection shown in Table 1, we set base daily mortality rates conditional on ICU need for the different groups at: $\underline{\delta}_y^d = 0.001\gamma/\iota_j$, $\underline{\delta}_m^d = 0.01\gamma/\iota_m$, and $\underline{\delta}_o^d = 0.06\gamma/\iota_o$.

Finally, we model the effect of the population infection rate on mortality due to limited ICU capacity by assuming that ICU needs are proportional to the mortality rates in Table 1, letting $\iota_j = \sigma\overline{\delta}_j^d$ for some parameter $\sigma > 0$. We set $\sigma = 0.0076$ based on the fraction of infections requiring ICU care by age used in Ferguson et al. (2020), adjusted for the structure of the US population, and the assumption that 10 days of the 18-day average case duration for those infected individuals is spent in the ICU.[29] Hence, ICU needs at time $t$ are $H(t) = (0.0076)\sum_k \overline{\delta}_k^d I_k(t)$, and we specify mortality rates as a function of

[27] This data set is more recent, more detailed and more systematically collected (using information all locations from 36,000 volunteers using the BBC pandemic project smart phone app) than other available sources such as POLYMOD (Hedengren et al., 2014; Prem et al., 2017). The POLYMOD data, in fact, show a stronger decline in contacts with age than the BBC pandemic project data.

[28] We used age-specific deaths reported on April 11 and divided by the total number of age- specific cases reported 18 days earlier. The data are available in the Korean language press releases of the Korean Central Disease Control Headquarters & Central Disaster Management Headquarters: `http://ncov.mohw.go.kr/tcmBoardList.do?brdId=3`

[29] Both Ferguson et al. (2020) and the Korean data, which report the numbers of current "critical" and "severe" active cases, provide support for the assumption that ICU needs for different age groups are proportional to case fatality rates.

$H(t)$ as follows:

$$\delta_j^d(t) \;=\; \underline{\delta}_j^d \cdot [1 + \lambda H(t)].$$

It is difficult to know how high mortality rates would go if ICU needs massively exceeded available capacity. In our baseline, we set $\lambda = \left[(0.0076)\sum_k \overline{\delta}_k^d N_k\right]^{-1}$, implying that when there is a uniform 10% infection rate for the three groups, mortality rates are 1.1 times the base mortality rates. We examine a higher mortality penalty in our robustness checks. In Section 5.6 we show that our results are robust to the extreme case in which ICU capacity is a hard constraint that cannot be exceeded.

We set the present discounted value of the lost work-life of the three groups upon death, $(\Delta_y, \triangle_m, \triangle_o)$, by assuming a retirement age of 67.5 years, so that there are 32.5 remaining work-years for the young, 10 years for the middle-aged and 2.5 years for the old. We also choose the interest rate as 1%.

Finally, in our baseline we treat the arrival time of the vaccine, $T$, as deterministic. We have experimented with some specifications including uncertainty, such as Poisson arrival rates with mean arrival times of 1 or 1.5 years, and the results are very similar. We prefer deterministic arrival times as our baseline, since these make it easier to interpret our solution. For example, with deterministic $T$ it is easier to judge whether the solution is attempting to avoid infections and hold out for the vaccine or giving up on this and going for herd immunity before its arrival. Specifically, we suppose that a vaccine will arrive in one and a half years, and so set $T = 548$ days.

**Computational Method.** The planning problem amounts to an optimal control problem over the horizon $t \in [0, T]$. There are nine state variables: $(S_j(t), I_j(t), R_j(t))$ for $j = y, m, o$. In the case of uniform policies there is one control, $L_j(t) = L(t)$, while with semi-targeted policies there are two controls $L_o(t)$ and $L_y(t) = L_m(t)$, and with fully targeted policies, three controls, $L_o(t)$,$L_m(t)$ and $L_y(t)$.

We discretize the problem to weekly time intervals. To generate our frontiers, we minimize a weighted sum of economic costs and deaths subject to the laws of motion of our model.

The solution is found by using a nonlinear programming method called IPOPT (Wächter and Biegler, 2006), which implements an interior point solution algorithm. We rely on an APMonitor-Gekko interface to formulate and pass the optimization problem to the IPOPT (Hedengren et al., 2014; Beal et al., 2018). As with any non-fully exhaustive numerical solution technique, one must consider the possibility that the solution converges to local

maxima that are not global maxima. To explore this possibility, we initiated the routines with different initial control sequences, but never encountered situations where the solution depended on these initial conditions.

# 5 Optimal Policies

In this section we present our main quantitative results for the baseline parameter values described in the previous section. Throughout, our main focus is on the comparison between the trade-offs presented by optimal uniform policies (where all three age groups are treated symmetrically), optimal semi-targeted policies (where the oldest age group is treated differently than the young and middle-aged groups) and optimal fully-targeted policies (where all three age groups are treated differentially). We summarize these trade-offs with the frontiers between lives lost and economic damages (as in Figure 1.1 in the Introduction). The focus on frontiers (rather than a single "optimal" policy) is motivated by disagreement among policy-makers, scholars, and members of the public about the exact value of life relative to other economic and social objectives, and because frontiers summarize the best policy options and provide a better sense of the range of improvements targeting enables. The main message from our baseline results in the next subsection is that targeting has major benefits in terms of both lives saved and reduced economic damages, and interestingly, most of these benefits can be achieved with the semi-targeted policies. The rest of the section turns to other (non-pharmacological) interventions, whose analysis is made possible or enriched by our model with multiple risk groups. First, we introduce and study the implications of "group distancing" (which involves additional social distancing between age groups). Next, we look at the effects of testing and tracing, with or without group distancing. We also discuss the role of the matching technology and how the arrival time of the vaccine changes trade-offs. We end the section with a battery of robustness checks, confirming the broad patterns we have so far emphasized—in particular, the substantial improvements in public health and economic outcomes owing to targeting.

## 5.1 Baseline Results

Figure 5.1 depicts the frontier between lives lost and economic damages under different policies. We use the baseline parameter values introduced in the previous section. Namely, we set $\alpha = 2$, so that the matching technology is quadratic; $\rho = 1$ so that, absent differential lockdowns, individuals match with members from their group and different

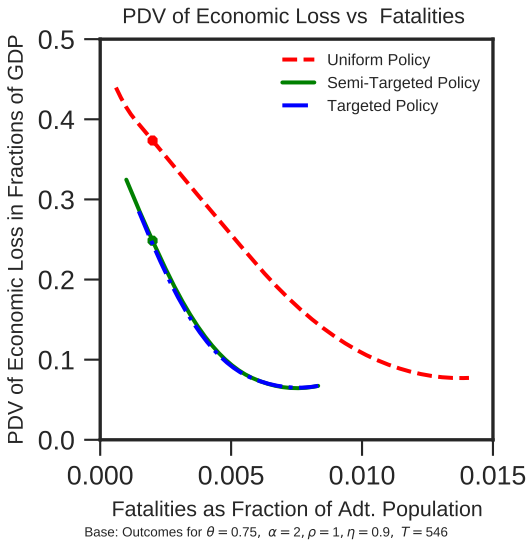


**Figure 5.1:** Frontiers of output loss vs. death for baseline specification. The three frontiers represent different levels of targeting.

groups at the same rate; $\theta = 0.75$ for all groups so that a full lockdown will be disobeyed a quarter of the time; $\xi = 0.3$, so that workers on average lose 70% of their productivity when under lockdown; $\eta = 0.9$, so that only the 10% of infected individuals end up being isolated in the baseline; and $T = 548$, so that the vaccine will arrive in one and a half years. Throughout, we take the initial conditions of the dynamical system to be 98% susceptible, 1% infected and 1% recovered within each group. We then trace the frontier between total deaths from the pandemic and its total economic damages.[30] As in Figure 1.1 in the Introduction, the bliss point in this figure is the origin, where there are no (excess) lives lost and no economic damages. Each curve in the figure represents the frontier resulting from a different class of policies: the top (red) frontier is for uniform policies, then below it we have the (green) frontier for semi-targeted policies, and slightly below this (in blue) is the frontier for fully-targeted policies. The convex shape of the frontiers represents diminishing returns to pursuing one objective at the expense of the other.

The main message from the red curve in Figure 5.1 is clear. The trade-off facing policy-makers when the menu of options is limited to uniform policies is quite grim. For example, policy-makers prioritizing saving lives could aim for keeping total mortality from COVID-19 to less than 0.2% of the (adult) population.[31] This "safety-focused" optimal

[30] In practice, we compute optimal policies for eleven points along the frontier and interpolate between them.

[31] A mortality rate of 0.2% is still very high. However, as we will see, without systematic testing and tracing, even this policy leads to very sizable economic losses. This motivates our choice of 0.2% as the benchmark for safety-focus.

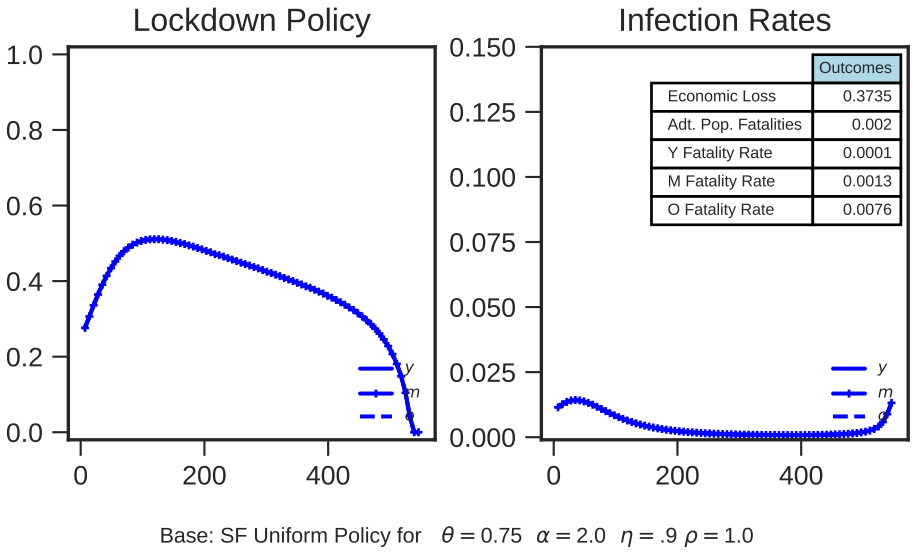


**Figure 5.2:** Optimal uniform policy for baseline parameters that achieves the "safety-focused" objective of limiting the population mortality rate to no more than 0.2%.

uniform policy is depicted in Figure 5.2 and would necessitate a lockdown remaining in effect in some form until the vaccine's arrival.[32] This lengthy lockdown has significant economic costs. The economic damages from these optimal uniform policies amount to 37.3% of one year's GDP (36.3% of this loss is in terms of current losses and the remaining 1% are due to the forgone productive contributions of those who die due to the pandemic). The economic loss and deaths that result from this policy are also represented by the dot that we show on the uniform policy frontier. Several points about the form of optimal policy are noteworthy in this case. First, consistent with our aggregation result (and because of the assumption that infection rates are symmetric across risk groups under uniform policies), the infection rates for the three age groups are on top of each other in the second panel of Figure 5.2. Nevertheless, the table on the top right of the figure shows that mortality rates are much higher for the older group, reflecting their greater vulnerability to the infection. Second, the time path of the infection rate follows an inverse U shape, typical in SIR models, peaking in about one and a half months and

[32]This policy would be optimal, alternatively, if we assigned a "value of statistical life" (cost of a death)—including both economic and non-pecuniary (psychic/emotional) costs—equal to $2.8 million. Given the 8.66 average life-years lost from a death under a uniform policy, this cost of a death implies a cost per life-year lost of $306,000. For comparison, the Environmental Protection Agency's *Guidelines for Preparing Economic Analyses* suggests a value of life equal to $9.4 million when updated to 2020 dollars. At a 3% annual real interest rate and a 40.3 year life expectancy for the median person in the US, this value of life implies a value per life-year of $402,000. In contrast, the US military appears to place a lower value, paying a total death benefit of $100,000 (https://myarmybenefits.us.army.mil/Benefit-Library/Federal-Benefits/Death-Gratuity-).

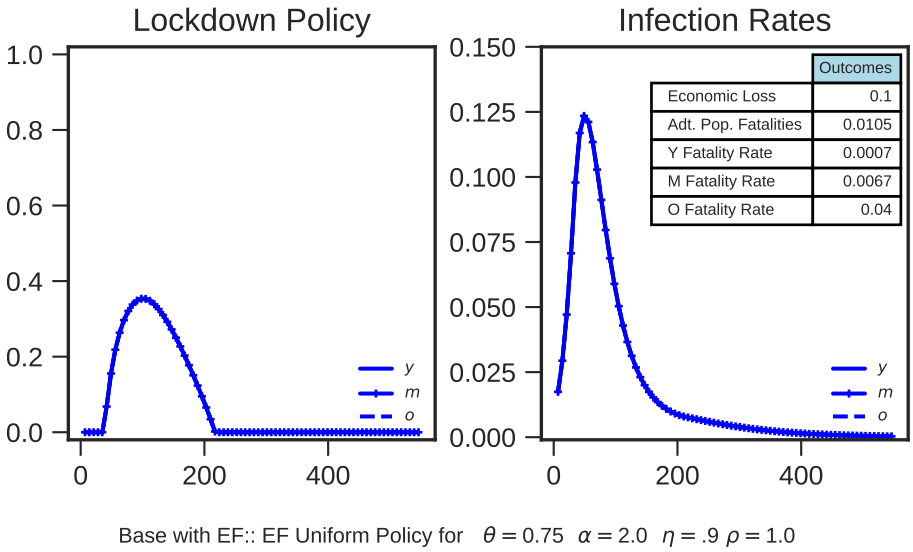


**Figure 5.3:** Optimal uniform policy for baseline parameters that achieves the "economy-focused" objective of limiting economic losses to no more than 10% of one year's GDP.

declining slowly thereafter. Third, the behavior of the infection rate reveals that optimal policy in this case is "waiting for the vaccine" as it does not lead to herd immunity against the virus. This can be seen from the fact that when the lockdown is lifted shortly before the vaccine's arrival, infections start increasing immediately (only to be brought under control by the vaccine). This last point is further illustrated in Figure A.1 in the Appendix, which plots the evolution of the share of susceptibles in the population and the reproduction rate of the virus under the safety-focused optimal uniform policy.

We can contrast this safety-focused optimal policy with another point on the frontier with very different priorities—an "economy-focused" optimal uniform policy, limiting economic damages to no more than 10% of one year's GDP.[33] This policy is shown in Figure 5.3. In this case, a significantly higher fraction of the population, about 1.05%, will perish because of the disease.[34] It is also worth noting that, differently from the safety-focused optimal uniform policy, the economy-focused policy goes for "herd immunity", with a shorter lockdown that nevertheless significantly flattens the curve of infections (which is beneficial for avoiding overwhelming ICU capacity as discussed in the previous section). Infections now peak at a higher level, about 9%, but they also decline to zero and

[33]The value of a life-year that would justify the economy-focused policy, computed in the same way as in the previous footnote, is $139,000 compared to $306,000 for the safety-focused policy.

[34]We stress that this is different from the "no lockdown" policy, which would lead to even more deaths, perhaps as much as 6.5% of the population, though this number is most probably an overestimate because, when there are no mandated lockdowns, individuals would adopt a range of voluntary social distancing measures, reducing the reproduction rate of the virus and infections.

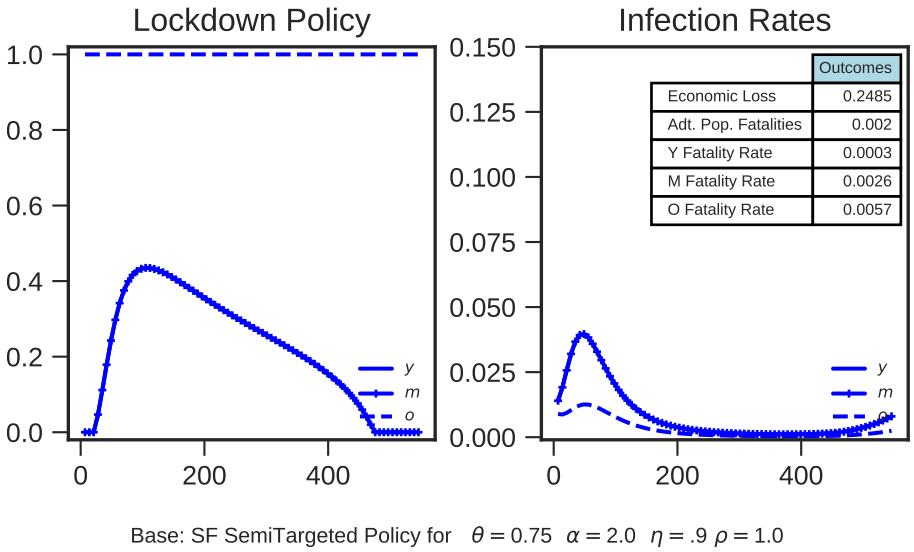


**Figure 5.4:** Optimal semi-targeted policy for baseline parameters that achieves the "safety-focused" objective of limiting the population mortality rate to no more than 0.2%.

never show a further uptick. This is because the faster spread of the virus leads to herd immunity even before large-scale vaccination.

Our main result can be gleaned by comparing the frontier for optimal uniform policies to the one for optimal semi-targeted policies, shown in green in Figure 5.1. For example, for the safety-focused objective which aims to keep total mortality from the virus to less than 0.2%, a semi-targeted policy can reduce economic losses from the 37.3% mentioned above to 24.9% (23.3% of this coming in the form of a decline in current GDP). The form of the safety-focused semi-targeted optimal policy is depicted in Figure 5.4 and has a number of noteworthy differences from the safety-focused optimal uniform policy. Most importantly, the lockdown is very strict on the older group and much less strict on the rest of the population, whose lockdown declines more rapidly and ends sooner than than in the optimal uniform safety-focused policy. Also notable is the time path of infections. The safety-focused optimal semi-targeted policy waits for the vaccine for the older group (who are in lockdown until the vaccine's arrival) but only partially so for the rest of the population (whose curve is again flattened so much that by the time the vaccine arrives, there is still no population-wide herd immunity, as can be seen from the uptick of the infections just before the vaccine). These points are also illustrated by Figure A.1 in the Appendix, which shows that the reproduction rate of the virus is above one for several months before the vaccine's arrival, but is everywhere less than under the optimal uniform policy depicted in the upper panels. Finally, compared to the pattern under optimal uniform policies, the infection rate of the 65+ group reaches a smaller peak, because

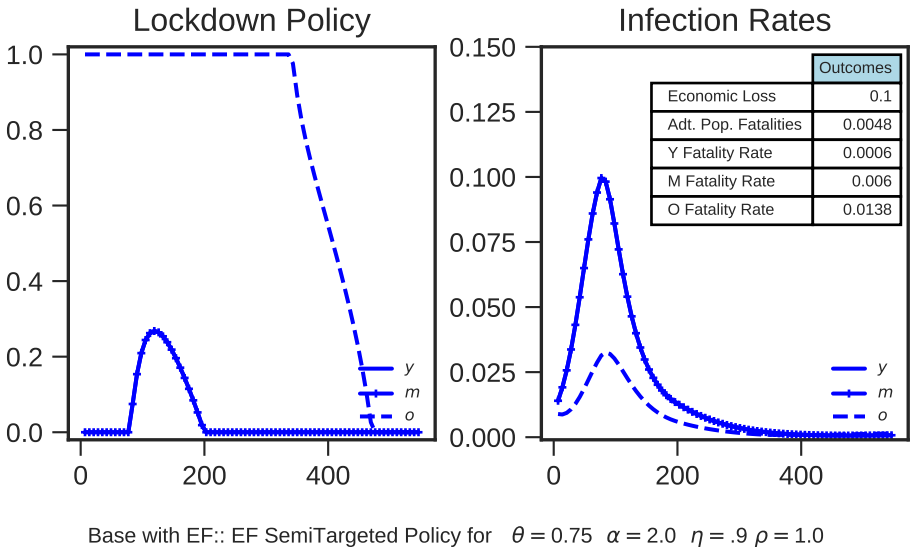


**Figure 5.5:** Optimal semi-targeted policy for baseline parameters that achieves the "economy-focused" objective of limiting economic losses to 10% of one year's GDP.

they are protected by their more strict lockdown. Notably, however, they are still being infected by the young and the middle-aged because our parameter choice of $\theta = 0.75$ implies that they are in not-too-infrequent contact with these younger groups. This is, in fact, the reason why the optimal semi-targeted policy in this case keeps the young and the middle-aged under a relatively long lockdown—as a way of protecting the old.[35]

Figure 5.5 turns to the optimal semi-targeted policy for achieving the economy-focused objective of keeping economic losses to less than 10% of one year's GDP. This policy achieves much better public health outcomes than the economy-focused optimal uniform policy, which led to an adult mortality rate of 1.05%. Instead, we now have a lower mortality rate—0.48%—because the stricter lockdown on the older group partially protects them from infections. Put differently, semi-targeted policies in this case can save 1.37 million (=(0.0105-0.0048)×241,000,000) additional lives relative to optimal uniform policies while achieving the same economic loss. It is worth noting that, like economy-focused uniform policies, optimal semi-targeted policies in this case also go for herd immunity, but with a nuance—this herd immunity is achieved primarily with the infections of the young and the middle-aged, while the more vulnerable older group is kept under lockdown. Herd immunity also explains why the older group is allowed to come out of its lockdown gradually starting in about a year.

[35] An even more "safety-prioritizing" strategy would be to take all the gains from targeting in the form of reducing mortality (keeping economic damages at 37.3%). If policy-makers pursued this strategy, semi-targeted policies would enable mortality to be reduced to 0.036%.

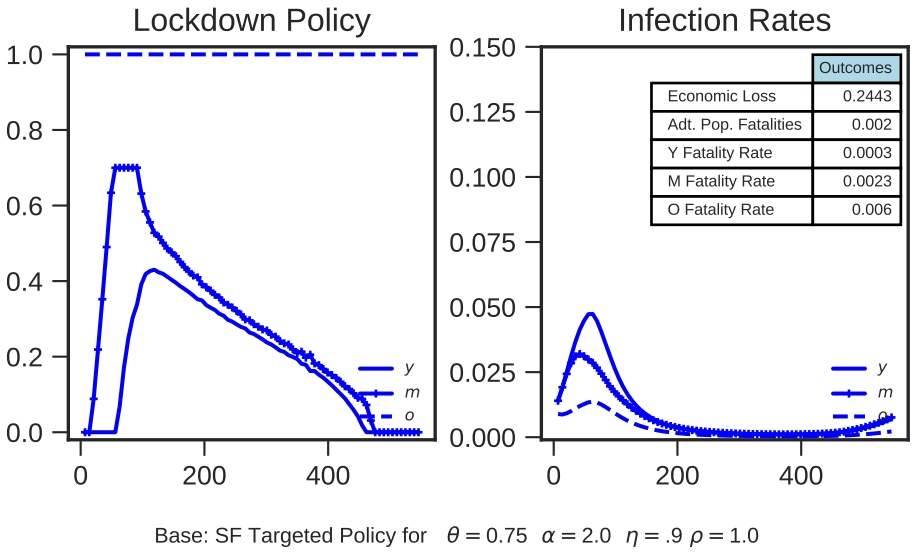


**Figure 5.6:** Optimal fully-targeted policyfor baseline parameters that achieves the "safety-focused" objective of limiting the population mortality rate to 0.2%.

A surprising result, at least relative to our initial expectations, is that fully-targeted policies that treat the young and the middle-aged differently perform essentially as well as semi-targeted policies. This can be seen from the fact that the blue curve in Figure 5.1 is nearly indistinguishable from the green curve (for semi-targeted policies). The reason is that the asymmetric treatment of the young and the middle-aged generates much smaller gains than those coming from protecting the most vulnerable with the strict lockdown on the older age group and with the moderate flattening of the curve of infection among the rest of the population. The form of optimal fully-targeted policies are somewhat different from those of optimal semi-targeted policies, however. This is shown in Figure 5.6 for the safety-focused fully-targeted policies and highlights that the middle-aged, who have higher mortality rates from the virus than the young, are put under a stricter and longer lockdown. The numbers in the table at the upper right corner of the figure show that there is some improvement over the optimal semi-targeted safety-focused policy in this case, but this improvement is small.

One question is whether the strict lockdowns on the old under targeted policy is primarily a result of their higher mortality rate or their lower economic participation. To shed light on this question, we considered an extended model with four groups. In particular, we separated the old into two groups, the old-retired, who have no economic contribution (making up about 80% of the old), and the old-workers, who have the same productivity as the young and the middle-aged (making up the remaining approximately 20% of the old). We took the mortality rates of the old-retired and old-workers to be

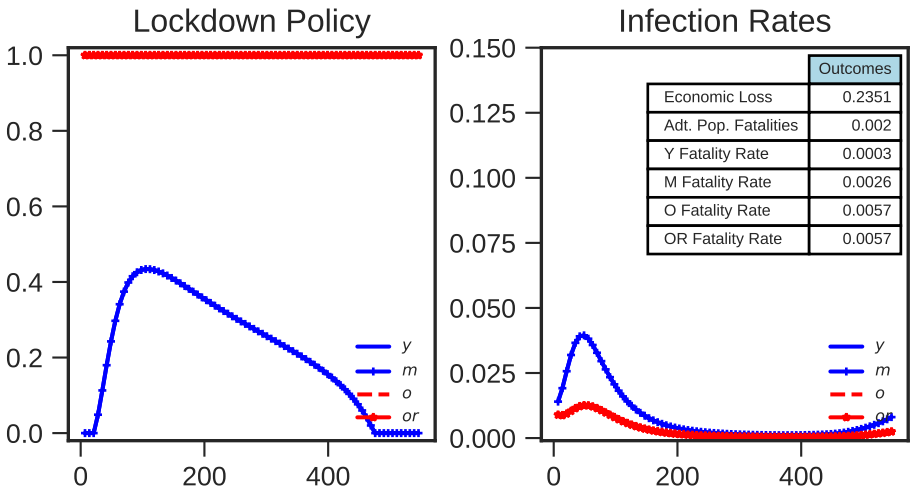


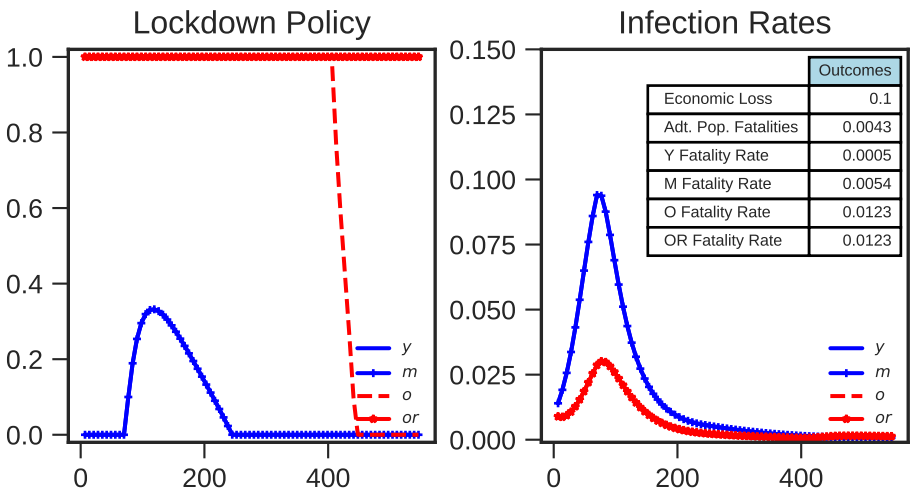


**Figure 5.7:** Optimal semi-targeted policies with separate old-working and old-retired groups – "safety-focus" (top) and "economy-focus" (bottom).

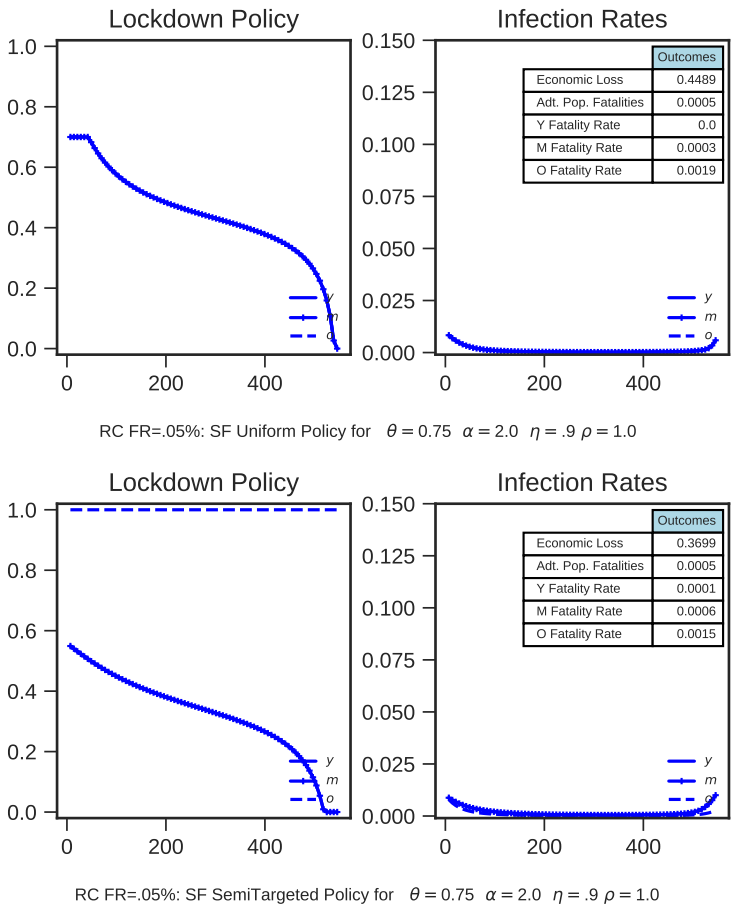


**Figure 5.8:** Optimal uniform and semi-targeted policies for a "safety-first" policy that achieves the objective of limiting the population mortality rate (with baseline parameters) to no more than 0.05%.

the same in order to make the two groups identical except for their economic opportunities (and thus clarify the source of our results in a more transparent manner). Figure 5.7 depicts the safety-focused and economy-focused optimal semi-targeted policies, which treats the young and the middle-aged symmetrically but applies differential lockdowns to the old-retired and old-workers, and thus is most useful for clarifying the form of our semi-targeted optimal policies. The figure confirms that the main source of the asymmetric treatment of the older and the younger groups is their differential susceptibility to the virus—the old-working, who have the same economic opportunities as the young and the middle-aged, are kept under complete lockdown like the old-retired in the safety-focused policy and under a relatively long lockdown, coming to an end only after one year, in the economy-focused policy.

Next, we explore another point on the frontier, which prioritizes saving lives more strongly than our safety-focused strategy. This strategy, which we label "safety-first",

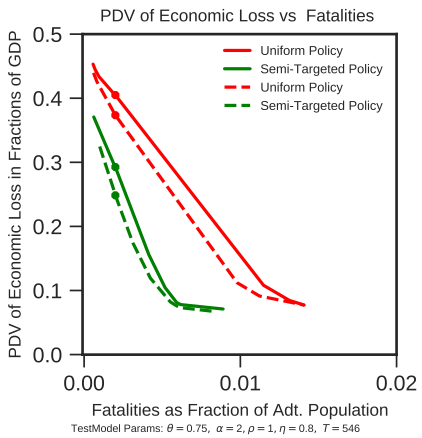


**Figure 5.9:** Optimal uniform and semi-targeted frontiers when recovered cannot prove immunity or avoid lockdowns (solid) and baseline where they can (dashed).

aims to keep overall (adult) mortality below 0.05%. Figure 5.8 shows the associated uniform and semi-targeted policies. Lockdowns are now much more severe, remaining at high levels and falling slowly up until the end, just before the vaccine and cure arrive. With uniform policies, the lower infection and mortality rates come at the higher economic cost of 45% (as a fraction of one year's GDP). Once again, however, targeted policies can significantly improve these outcomes. For example, with optimal semi-targeted policy, the same mortality objective can be reached with economic damages amounting to 37% of one year's GDP.

Finally, we explore the implications of relaxing the assumption that the previously infected and now recovered individuals are identified as such and allowed to work, circumventing any lockdowns. Figure 5.9 shows the baseline parameterization but setting $\kappa_j = 0$, so the recovered agents are treated like the rest of the population and cannot avoid lockdowns. This makes lockdowns more costly, especially later in the epidemic. Figure 5.9 confirms that the frontiers in this case (shown with solid curves) are further away from the origin than our baseline frontiers (shown with dashed lines) where the recovered could prove their immunity. Nevertheless, the figure also confirms that semi-targeted policies still significantly outperform optimal uniform policies.

Overall, our main results show that the trade-offs between lives lost and economic damages from the pandemic are significantly improved when we consider targeted policies and most of the gains can be achieved with simple semi-targeted policies that apply more strict lockdowns on the older, more vulnerable group. We next show that there are even larger gains when semi-targeted lockdown policies are combined with other non-pharmacological interventions.

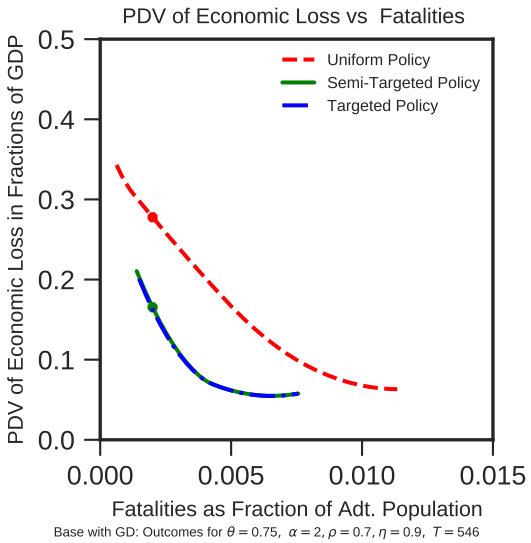


**Figure 5.10:** Frontier of output loss vs. death with greater group distancing, $\rho = 0.7$.

## 5.2 The Role of Group Distancing

Our model with multiple risk groups enables an investigation of the implications of policies that reduce inter-group interactions. These types of "group distancing" measures are particularly important in conjunction with targeted policies, because, as we have emphasized, older individuals get infected from their interactions with the young and the middle-aged, even when they are under strict lockdown. Recall that our baseline parameterization assumed the rate of social contact and infection, absent lockdowns, to be the same across groups—summarized with the infection rate $\beta$. A natural alternative is to assume that matches between groups can be reduced, say by a fraction $1 - \rho$, so that

$$\rho_{ij} = \begin{cases} 1 & i = j \\ \rho & i \neq j. \end{cases}$$

Our baseline then corresponds to the special case of this formulation where $\rho = 1$, and we now consider the implications of reducing $\rho$. There are various policy tools for achieving such reductions, including norm-based interventions (so that people visit their elderly relatives less often) or law-based interventions (e.g., designating elderly-only hours at supermarkets and pharmacies or restricting who can visit and work in nursing homes).[36]

[36]With our baseline quadratic matching technology, any change in between-group matching will influence the total number of matches and do so in ways that depend on group size. With the interventions we have in mind, we believe this is the right type of variation to consider, though it should be borne in mind that reduced number of matches will directly decrease infection rates as well.

Figure 5.10 depicts the frontier between lives lost in economic damages under uniform, semi-targeted and fully-targeted policies when $\rho = 0.7$, which implies a 30% reduction in the interactions between the older age group and the rest of the population relative to the baseline. The main message is the same as in the previous subsection: there is a significant improvement when going from uniform to semi-targeted policies, evidenced by the sizable shift-in of the frontier from the red (uniform) to green (semi-targeted). There is essentially no further improvement from full targeting.

This reconfirmation of our main message notwithstanding, there are also notable differences from the baseline. First, adding group distancing increases the gains from targeted policy, as depicted in Figure 5.11: safety-focused semi-targeted policies with group distancing, for example, can achieve a 0.2% mortality with 16.5% economic damages—an almost one-third reduction in economic loss compared to the optimal safety-focused semi-targeted policy without group distancing. This figure also shows that group distancing improves uniform policies, but does not change their overall structure. In particular, in this case it is still optimal to wait for the vaccine, but now the lockdown is less severe because group distancing enables the young and the middle-aged to go back to work with less impact on the old. The corresponding semi-targeted policy changes in notable ways, however. First, the lockdown on the younger age groups is now much less severe, peaking at just about 30% of these workers, and is completely lifted in just over a year. Second, the older group is now released from lockdown before the vaccine's arrival, which is thanks to the fact that group distancing further protects them from the infected among the young and the middle-aged.

The economy-focused optimal policies, both uniform and targeted, are similar to those in our baseline case, but lead to lower mortality. For example, as Figure 5.12 shows, mortality falls from 0.48% to 0.32% in the optimal semi-targeted policy.

The overall message from this subsection is that, if it is feasible to reduce interactions between high-risk groups and the rest of society with policies similar to those used for lockdown, then mortality rates or economic damages, or both, can be reduced considerably.

## 5.3 The Effects of Testing and Tracing

We next investigate how the ability to test and isolate infected individuals affects optimal policies and outcomes. Recall that testing and tracing were incorporated into our model in Section 2, but our baseline scenario assumed that these were not being used, setting $\eta = 0.90$ to reflect only that 10% of the infected would be isolated because of illness mak-

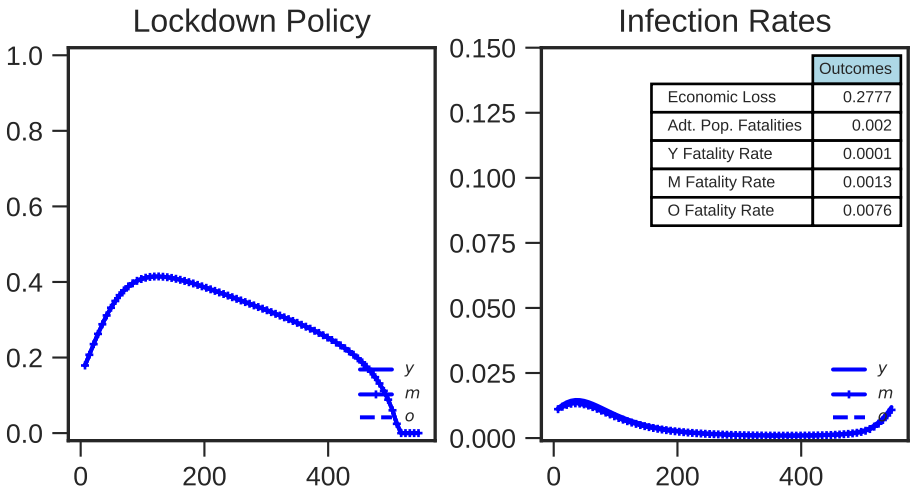


**Figure 5.11:** Optimal uniform and semi-targeted "safety-focused" policies with greater group distancing $\rho = 0.7$.

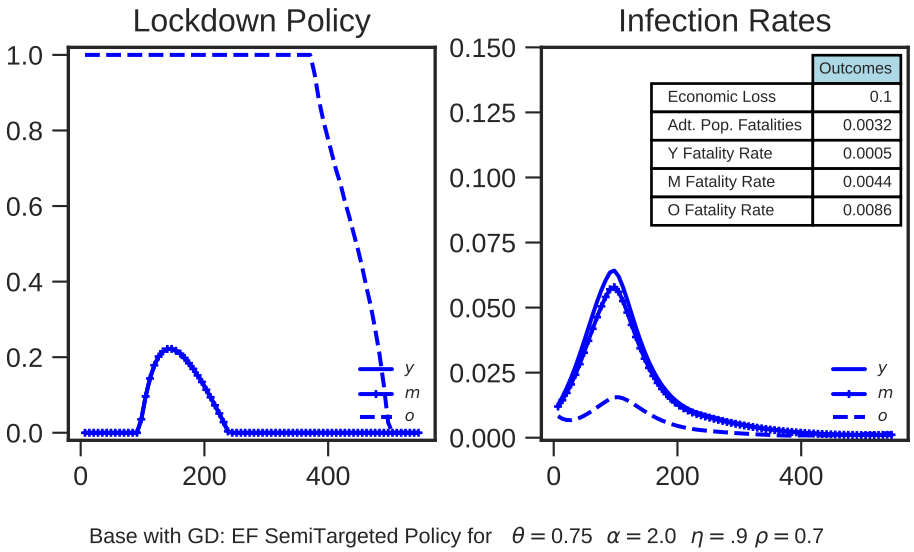


**Figure 5.12:** Optimal semi-targeted "economy-focused" policy with greater group distancing ($\rho = 0.7$).

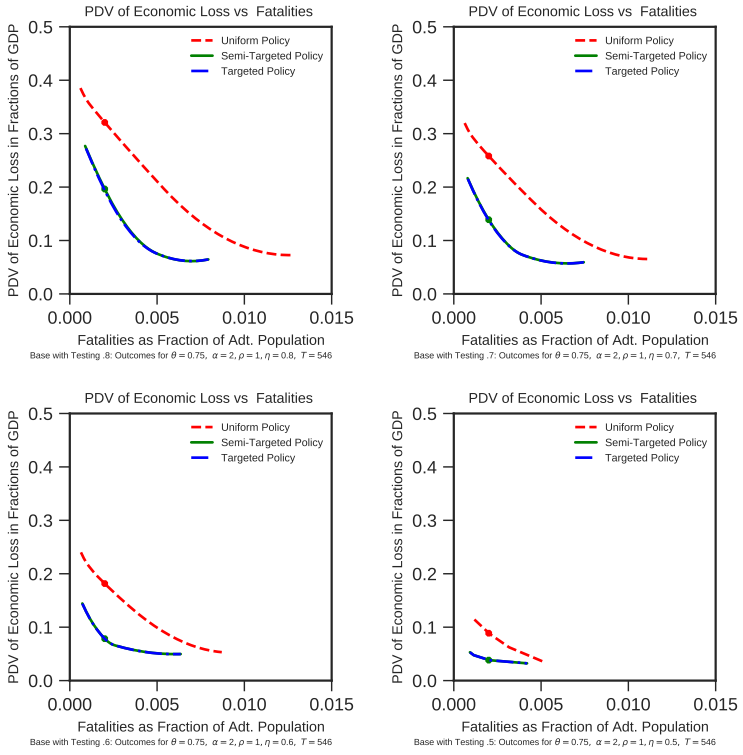


**Figure 5.13:** Frontiers of output loss vs. death with improved testing and isolation.

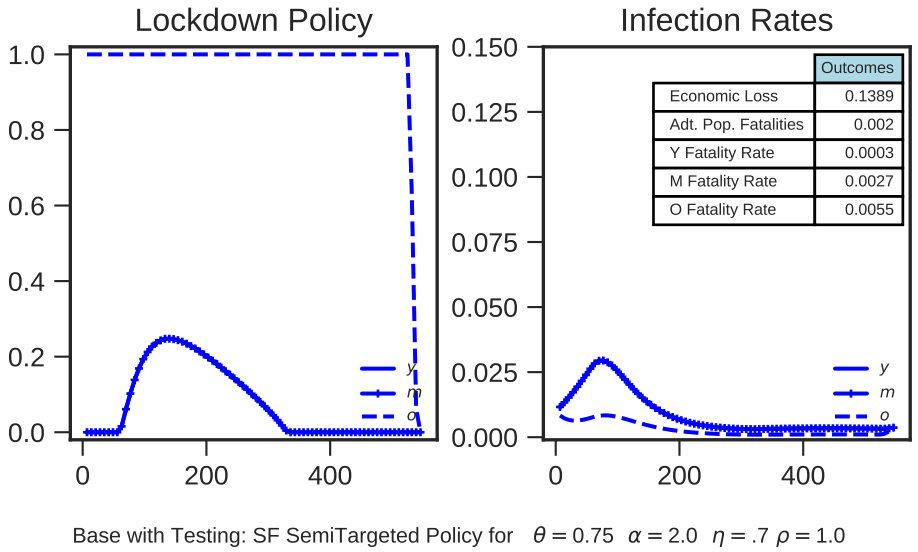


**Figure 5.14:** Optimal "safety-focused" semi-targeted policy with improved testing and isolation ($\eta = 0.7$).

ing them unable to work or socialize. We now introduce testing and tracing and, given the uncertainties about the exact effectiveness of these measures and the US capacity to implement them, we consider four cases in Figure 5.13: $\eta = 0.80$, $\eta = 0.70$, $\eta = 0.60$ and $\eta = 0.50$. For example, with a policy in which everybody who is symptomatic is isolated, we would end up with $\eta = 0.60$.[37] In practice, the limited number of tests and implementation difficulties may imply a lower probability of isolating infected individuals (and thus higher value such as $\eta = 0.70$ or even $\eta = 0.80$) and successful contract tracing might lead to a higher probability of identifying and isolating infections (e.g., $\eta = 0.50$). We keep all other parameters at the levels in our base specification.

Figure 5.13 presents the uniform and semi-targeted frontiers for these four values of $\eta$. ($\eta = 0.80$ is in the top left corner of the figure; other values proceed clockwise.) As the ability to identify and isolate the infected increases, the trade-offs improve remarkably. When $\eta = 0.60$ or 0.50, Figure 5.13 shows that mortality rates can be kept at very low levels at relatively small economic costs.

Figure 5.14 explains why this is so for the case in which the probability of identifying and isolating the infected individual is equal to 0.30 or $\eta = 0.70$: a short and limited lockdown on the young and the middle-aged, combined with a still fairly strict lockdown on the older group, is enough to keep infections very low and the overall mortality rate at 0.2%, and has an economic cost of just 13.9% of GDP.

[37] This number is based on the following reasoning: approximately 50% of infected individuals are assume to be asymptomatic and it takes five days for the remaining 50% to exhibit symptoms.

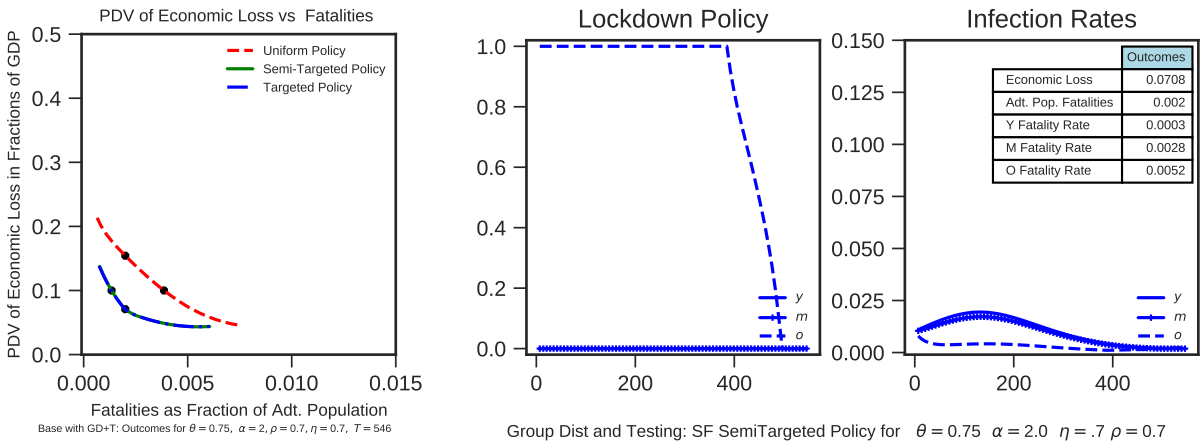


**Figure 5.15:** Optimal Semi-Targeting with group distancing and testing.

This conclusion is further reinforced in Figure 5.15, which presents the frontiers and the form of the optimal semi-targeted policy (for the safety-focused case) when there is group distancing as well as testing-tracing sthat yields $\eta = 0.70$. Thanks to the combination of group distancing and testing-tracing, semi-targeted policies now achieve the 0.2% mortality objective at an economic cost of just 7% of one year's GDP and requires no lockdown of the young and middle-aged groups.

Overall, measures to combat the pandemic become much more powerful when targeted policies are combined with group distancing and improved testing and tracing strategies. With these tools, relatively short lockdowns are sufficient to protect public health and naturally lead to much more limited economic damages.

## 5.4 The Role of the Matching Technology

As noted in the Introduction, our baseline model assumes quadratic matching. To highlight the role of the matching technology we now present optimal policies when the matching technology has a more limited degree of increasing returns to scale, in particular, $\alpha = 1.5$. In Figure 5.16 we show the frontiers for optimal uniform, semi-targeted and fully-targeted policies for $\alpha = 1.5$, together with the frontiers for $\alpha = 2$ for comparison. The contrast of the frontiers shows that the trade-off facing policy-makers is even worse under more limited increasing returns in the matching technology. Interestingly, this worsening is more pronounced when the priority is saving lives rather than saving the economy. Figure 5.17, which plots optimal uniform and semi-targeted policies for the safety-focused case, explains why. Uniform policies now necessitate an economic loss of over 45% (again expressed as a fraction of one year's GDP) to achieve a 0.2% overall

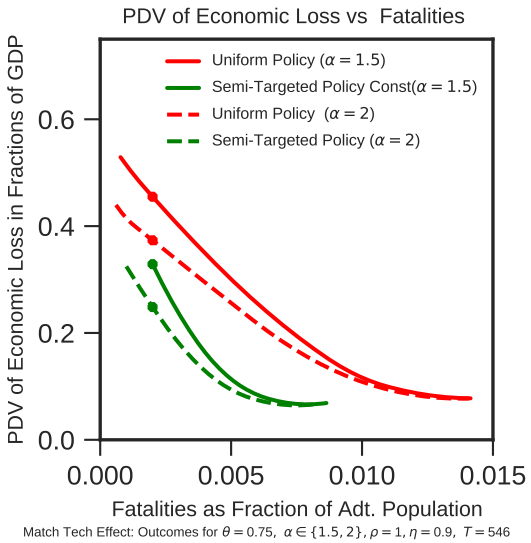


**Figure 5.16:** Frontier of output loss vs. death with quadratic and non-quadratic matching technologies: $\alpha = 2$ vs. $\alpha = 1.5$.

mortality. As usual, targeted policies do much better, but still lead to an economic cost of 32.8%.

This result may at first appear surprising, since herd immunity is easier to achieve when there are more limited increasing returns in the matching technology and this might make us expect to see lower economic costs. Going against this force, however, is that lockdowns are less effective when the extent of increasing returns is less than quadratic, because withdrawing a fraction of the population from the matching pool does not lead to a quadratic decline in the number of matches. Our results show that the second effect dominates, and as a result, the trade-off facing policy-makers worsens when $\alpha = 1.5$.

The situation is somewhat different when the objective is economy-focused, as depicted in Figure 5.18. In this case, the additional loss relative to the quadratic matching technology is small—with semi-targeted policies, 0.0054% mortality as compared to 0.0048% before. This is because the economy-focused policies go for herd immunity, in which case the effects of the less powerful lockdowns under $\alpha = 1.5$ are lessened and the benefits of more limited increasing returns in matching are realized. This highlights that more limited increasing returns to scale in matching tends to favor strategies that go for herd immunity.

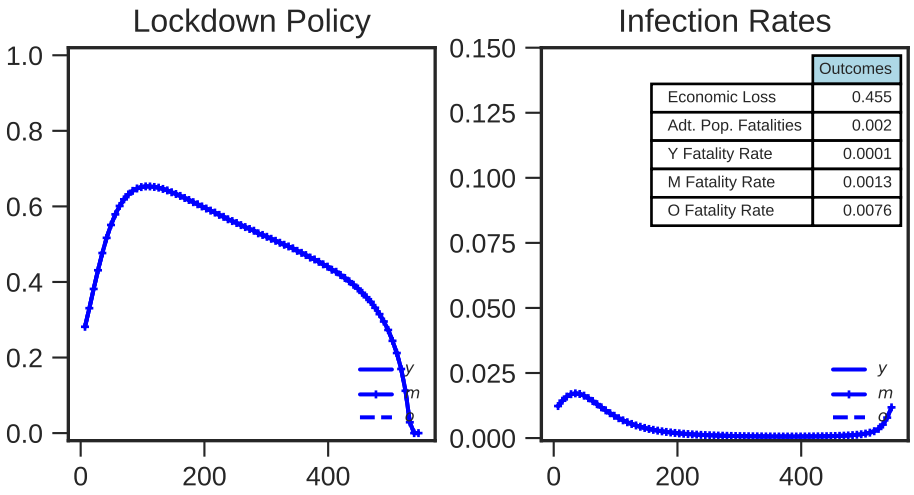


**Figure 5.17:** Optimal uniform and semi-targeted policies with $\alpha = 1.5$ to achieve the "safety-focused" objective of limiting population mortality to 0.2%.

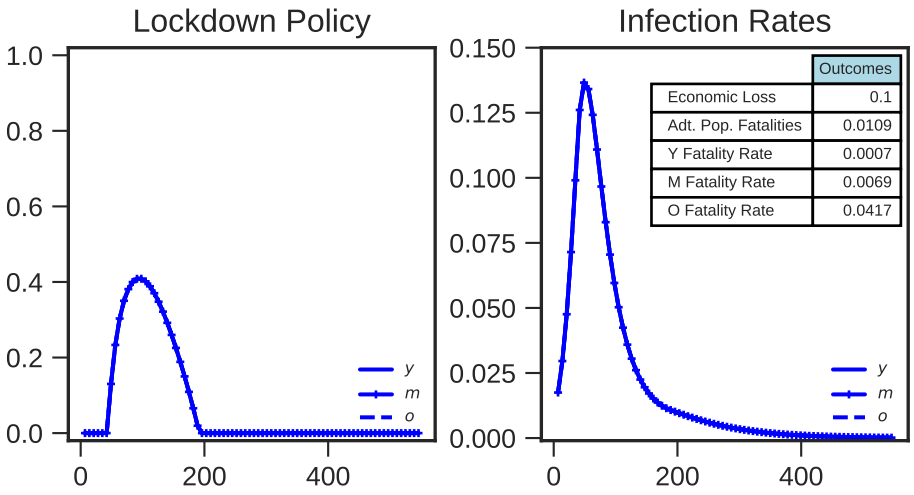


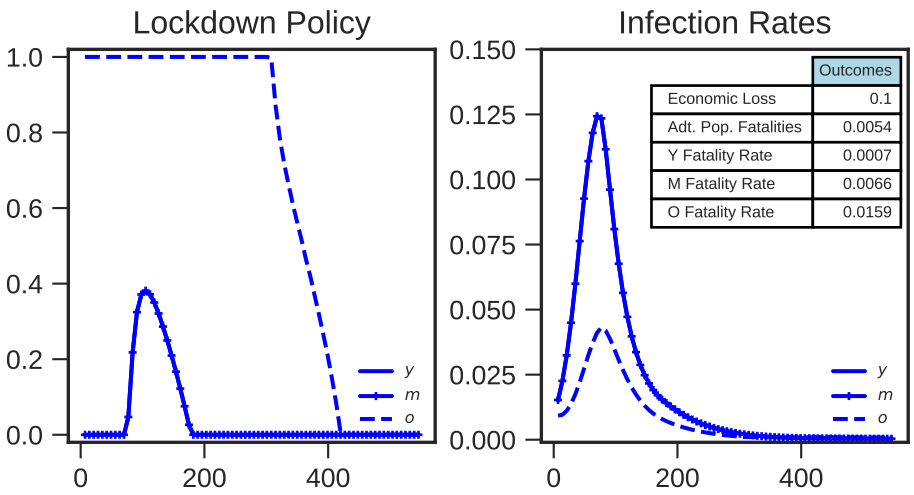


**Figure 5.18:** Optimal uniform and semi-targeted policies with $\alpha = 1.5$ to achieve the "economy-focused" objective of limiting economic losses to 10%.

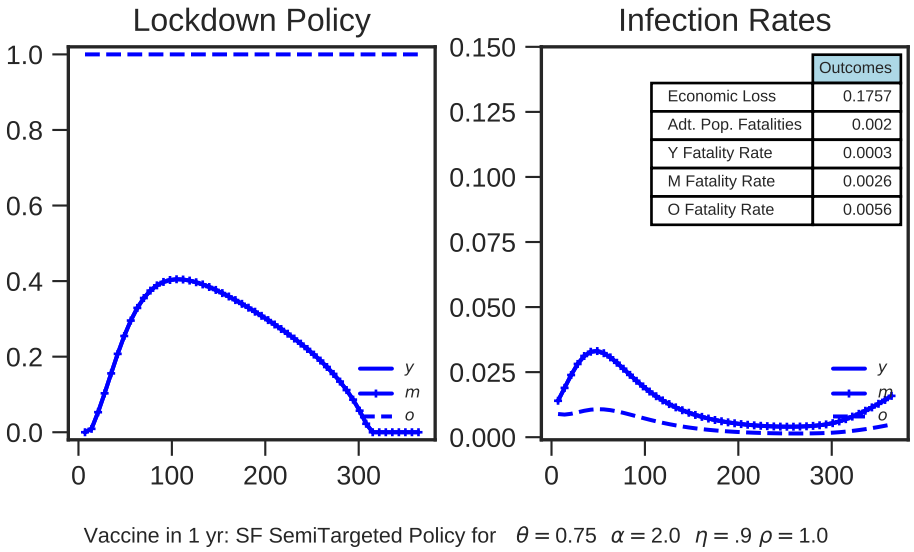


**Figure 5.19:** Optimal semi-targeted ""safety-focused" policy with earlier vaccine arrival $(T = 365)$.

## 5.5 The Promise of a Vaccine

In this subsection, we consider the implications of an earlier arrival date for the vaccine. Namely, we reduce $T$ from one and a half years to one year. This has little impact on the form of the safety-focused semi-targeted policy, which is shown in Figure 5.19 and again keeps the old undergoing a lockdown until the vaccine arrives and has a fairly lengthy lockdown on the younger age groups. The economic costs, however, are now lower (for example only 17.6% of one year's GDP for the safety-focused case shown in the figure), because lockdowns end sooner with the earlier arrival of the vaccine: the earlier vaccine arrival saves 7.3% in economic costs (equivalent to $1.8 trillion) while achieving the same 0.2% mortality rate .

## 5.6 Other Robustness Exercises

As already noted in the Introduction, there is considerable uncertainty about both the relevant parameters for the COVID-19 virus and for the economic damages from lockdowns. In this subsection, we confirm that our qualitative conclusions on significant gains from semi-targeted policies, both in terms of lives saved and reduced economic damages, are robust to a broad range of variations in parameters. To conserve space, all of the relevant figures are presented in the Appendix, and here we here confine ourselves to a brief discussion of the results and the implications of these robustness checks.

First, we follow some of the epidemiology literature on COVID-19 policy and im-

pose a hard constraint on ICU capacity (because over-running ICU capacity would be extremely costly and unacceptable).[38] In practice this means that only policies that ensure $H(t) \leq 0.000133$ at all times are feasible. This constraint implies that under a uniform policy, ICU capacity would be reached with a population infection rate of 2%.[39] Figure A.2 shows the frontiers in this case both with and without imposing the constraint.[40] The ICU constraint binds along the uniform policy frontier, worsening achievable outcomes, for economic losses less than 0.34.[41] Semi-targeting helps avoid hitting ICU capacity, so that the ICU constraint binds below an economic loss of 0.20. To illustrate how optimal policies respond to the need to avoid exceeding available ICU capacity, Figure A.3 shows the economy-focused uniform policy that results in an economic loss of 0.3%. We see in this figure that optimal policy now ensures that we remain just against the ICU constraint for over six months, and thereafter the lockdown is tightened as we wait for the vaccine. It is worth stressing that there is no analogue of the repeated toggling on and off of lockdown measures that emerge in studies such as Ferguson et al. (2020).

We also investigate the implications of increasing the mortality rate of the oldest group following infection from the Ferguson et al. (2020) numbers to the higher South Korean numbers. This naturally impacts the overall mortality numbers under a fixed policy, but has little effect on the gaps between the three frontiers as illustrated in the first panel of Figure A.4. The table inside the third panel indicates that, in this case, a 0.2% overall (adult) mortality can be achieved at the slightly higher economic cost of 30.7% than in the baseline semi-targeted policy.

Our next exercise reduces the basic transmission rate of the disease, which corresponds to the parameter $\beta$ in our model decreasing from 0.134 (which implied $R_0$ =2.4) to 0.1 (which implies $R_0 = 1.8$). Figure A.5 shows that, not surprisingly, semi-targeted policies now achieve a 0.2% mortality rate in the adult population at a much lower economic cost: 12% of one year's GDP, with semi-targeted policies. Nevertheless, the significant gain from targeting remains regardless of the exact value of $\beta$.

[38] We also experimented with a higher mortality penalty, $\lambda$, to capture the extreme scenario in which nearly those who need but are denied ICU die. The results are similar to our baseline, but of course lead to higher mortality rates, especially under the economy-focused policies. The value of semi-targeted policies are even greater in this scenario.

[39] Prior to the emergence of COVID-19, the US had roughly 32,000 available ICU beds. We suppose that available beds can be surged by 15%. Following Ferguson et al. (2020), we assume that 1.58% of adult infections in the US require ICU care and that, for those cases, 10 days of the 18-day average case duration will have to be in the ICU.

[40] We no longer impose a mortality penalty, so the frontiers without the constraint differ slightly from those in our baseline case.

[41] We also note that it is impossible to achieve an economic loss below 0.25% without violating the ICU capacity constraint.

We next change the initial state to reflect what may be the current situation in some areas of the US with COVID-19: all three groups start with 15% of their members already recovered, 84% susceptible, and 1% infected. This leads to large reductions in both economic losses and mortality for all policies, which can be seen from the positions of the frontiers in Figure A.6 relative to the baseline in Figure 5.1. This is a direct consequence of starting with a smaller susceptible population, which directly reduces infections and also alters the nature of the safety-focused optimal policies. Indeed, Figure A.6 shows that the optimal semi-targeted policy gets closer to (but does not quite achieve) herd immunity for the younger groups as can be seen from the smaller uptick in infection rates (compared to our baseline) just before the vaccine's arrival as well as from the fact that the older group is released from lockdown before the vaccine. The implications of this change are quite significant for the economic costs of the safety-focused semi-targeted policy, which are now a 13.9% economic cost (compared to the 24.8% economic damages in the baseline).

We next turn to a number of variations concerning social and economic parameters. We start with the implications of a lower value of $\theta$, 0.5, rather than our baseline of 0.75. This variation is important for verifying that the benefits from the differential treatment of the more vulnerable groups remain even when there is greater slippage from the lockdowns. Figure A.7 shows that, in this case, both safety-focused optimal uniform and optimal semi-target policies become much more strict in response to lockdowns becoming less effective. This might have been expected to make optimal policy go for herd immunity, as it has now become very difficult to protect the vulnerable elderly. However, the safety-focused policies, which aim to keep mortality in the population to below 0.2%, are forced to go for longer lockdowns to achieve this objective. The result is much larger economic costs, 52.7% under uniform policies and 42.7% under targeted policies. Nevertheless, as these numbers suggest, the significant gains from targeting remain even in this case. Figure A.8 shows the economy-focused uniform and semi-targeted policies with this lower value of $\theta$. Confirming this interpretation, the economy-focused optimal policies that still go for herd immunity respond to the lower effectiveness of lockdowns by opening the economy sooner than in the baseline. The lockdown on the younger groups is completely ended by about six months and even the old are released from their lockdown completely before the end of the first year. In both cases, targeted policies continue to improve outcomes significantly.

In Figure A.9 we consider a different formulation for "working from home". Our baseline assumption was that all workers could work from home but this entailed a productivity loss. We assume that only a fraction 30% of the population can work from home

and can do so with full productivity. This improves the effectiveness of both uniform and targeted policies. But the gains are more pronounced with targeted policies because now the young and middle-aged workers who can work from home productively can be locked down at no cost, while the remainder are prioritized for release given their lower productivity at home. As a result, the economic cost of achieving 0.2% mortality in the population is now 19.7% compared to 24.8% in the baseline. This important difference notwithstanding, the comparison of the different frontiers is very similar to our baseline results.

We next consider the implications of increasing the value of the old being out of lockdown. Our baseline value of this parameter, $w_o = 0.20$, was motivated by the relative earnings of the over 65 group. There may, however, be additional gains from having the older group out of lockdown, because they may better enjoy leisure or contribute to economic activity via their spending. To verify whether our results are robust to incorporating these considerations, we increase $w_o$ to 0.50. Figure A.10 shows that this variation has little impact on our main results. The gap between the uniform and semi-targeted frontiers remains sizable. In addition, the safety-focused optimal semi-targeted policy now exhibits a slight non-monotonicity for the lockdowns on the old. The intuition for this behavior is as follows: optimal policy starts releasing the old from their lockdown after about nine months to benefit from the now higher value of their economic and social activity. However, once the old join the matching pool, infections in the population start increasing with the old also taking part in economic and social activity, and after three months, the lockdown on this group is re-tightened. It is also interesting to see that at the moment the old are back in lockdown, there is an inflection in the lockdown time-path for the younger groups. This is because, once the old are protected via their own lockdown, the release of the younger groups can be accelerated, since the concern that their infections will be transmitted to the more vulnerable elderly is lessened.

We next investigate the implications of incorporating a contact matrix based on Klepac et al. (2020), as discussed in Section 4, rather than our baseline uniform contact matrix. This contact matrix, depicted in Figure A.11, implies that age groups are more likely to interact within themselves than with other age groups and the older group have fewer contacts.[42] The results are presented in Figure A.12, which shows that our main results and even the quantitative gains from targeting are similar to our baseline findings. For ex-

[42]Specifically, we now have$\rho_{yy} = 1$, $\rho_{mm} = .6$, $\rho_{oo} = 0.5$, $\rho_{ym} = 0.5$, $\rho_{yo} = \rho_{mo} = 0.4$. As already noted in Section 4, the alternative POLYMOD (Hedengren et al., 2014; Prem et al., 2017) data show an even stronger decrease in contacts with age and an even lower level of contacts between the old and other age groups. We prefer to use the more recent and more detailed evidence. If the old were further isolated from the younger groups, the benefits of targeting would be greater.

ample, semi-targeted safety-focused policies reduce economic damages from 38% to 27%. But there are also two noteworthy differences. First, fully-targeted policies now generate somewhat greater gains than semi-targeted policies. Second, the lockdown on the old is again non-monotonic. The intuition is that the asymmetric contact matrix introduces a strong form of group distancing, which allows the old to be released from their lockdown for a while, despite the higher levels of infections among the younger groups. But once infections within this group start increasing, a second lockdown for them becomes necessary and is maintained until the vaccine arrives. This pattern recurs with fully-targeted policies in the bottom panels.

Finally, we turn to an investigation of whether the fact that in our model we did not include exposed, but not yet non-infectious individuals mattered for our conclusions. This is important, since COVID-19 infections do not immediately lead to infectiousness, which emerges after about 5 days. Incorporating this feature necessitates generalizing the classic SIR model to its SEIR variant.[43] We do this in Figure A.13, where we assume that, on average, an infected person is non-infectious for five days before becoming infectious for an another five days. We recalibrate $\beta$ so that the initial reproduction rate of the virus, $R_0$, is still 2.4. The results from this exercise are remarkably similar to our baseline. For example, the safety-focused optimal targeted policies save about 12% in terms of economic damages relative to safety-focused optimal uniform policies. Figure A.14 shows that incorporating the contact matrix from Klepac et al. (2020) with this SEIR extension leads to very similar results to those in Figure A.12, including the somewhat larger gains from fully- targeted policies and the non-monotonic lockdown pattern for the old.

Overall, the various exercises in this subsection demonstrate that our qualitative conclusions on the significant gains that can be obtained by targeting are quite robust.

[43] Mathematically, this involves modifying our equations, so that now

$$\dot{E}_j = M_j(S, I, R, L)\beta(1 - \theta_j L_j)S_j \sum_k \rho_{jk}\eta_k(1 - \theta_k L_k)I_k - \gamma_j^E E_j$$
$$\dot{I} = \alpha_j E_j - \gamma_j^I I_j$$
$$\dot{S}_j = -\dot{E}_j - \gamma_j^E E_j$$

and the rest is unchanged. Here $1/\gamma_j^E$ is the average time individuals are exposed and non-infectious, while $1/\gamma_j^I$ is the average time that individuals remain infectious.

# 6 Conclusions

In this paper, we took a first step in analyzing the role of optimal targeted lockdowns in a multi-group extension of the standard SIR model. This generalization is important in the context of the COVID-19 pandemic, since existing evidence demonstrates very large differences in hospitalization and fatality rates between age groups. After providing a basic analysis of the dynamics of infections in this multi-group setting, we proceeded to a quantitative investigation of optimal policy, focusing in particular on the frontier summarizing the trade-offs facing policy-makers between saving lives and improving economic outcomes. Optimal uniform policy, which treats different demographic groups symmetrically, as in other works on COVID-19, underscores the grim choices confronting policy-makers. In our parameterizations, if policy-makers go for safety-focus and try to limit the mortality rate in the adult population to no more than 0.2%, then they have to put up with economic losses amounting to 37.3% of one year's GDP. If they prioritize the economy and attempt to limit economic losses to be less than 24.8% of one year's GDP, then society suffers a mortality rate as high as 1.05%.

Our main result, however, is that better social outcomes are possible with targeted policies. Differential lockdowns on groups with differential risks can significantly improve policy trade-offs, enabling large reductions in economic damages or excess deaths or both. We also find that the majority of these gains can be achieved with a simple targeted policy that applies an aggressive lockdown on the oldest group and treats the rest of the population uniformly. These qualitative conclusions are quite consistent across different parameterizations of our model and are the main take-away message from the paper. For our baseline parameterization, semi-targeted policies could limit economic losses to 24.8% of one year's GDP while maintaining an adult mortality rate of 0.2%, or limit mortality among the adult population to 0.48% while keeping economic losses at 10% of one year's GDP.

We hasten to emphasize that there is considerable uncertainty about some of the key parameters governing the spread of COVID-19, including how long-lasting immunity is, its basic reproduction rate, the base mortality rates by group, the additional mortality penalty from ICU overcapacity and how effective new cures and vaccines will be, as well as on the exact economic damages caused by lockdowns (in part because neither the extent to which work from home can substitute for workplace interactions nor the knock-on effects of current measures on supply chains and worker-firm relations are yet well understood). There is also no consensus on the value of life relative to the social and economic damages of lockdowns. These are the primary reasons why we emphasize the qualitative

conclusions from our work—that the trade-off between lives lost and economic damages improves substantially with targeted policies. We also document that our main results are robust to a range of changes in parameters.

It is noteworthy that the gains from targeted policies can be substantially increased if we combine them with additional measures that our multi-group SIR model highlights. For example, increasing the "social distance" between the oldest group and the rest of the population—by norms that temporarily reduce visits to older relatives or regulations that better protect nursing homes and segregate the times when different demographic groups can go to grocery stores and pharmacies—can reduce fatalities to as little as 0.02% of the population and cut economic damages to about 7% of one year's GDP. Another promising policy, which is also synergistic to targeting, is testing and tracing. Improved testing and tracing combined with targeted policy can reduce fatalities to 0.02% of the (adult) population and economic damages to about 7%. If targeted policies, group distancing and testing and tracing were combined, the trade-offs facing policy-makers would be radically improved, and our baseline parameterization suggests that the pandemic could be brought under control very rapidly and with minimal deaths and economic costs.

There is a sense in which our analysis understates the gains from targeting, because we have focused only on targeting by age. The mortality rates of COVID-19 also vary significantly by pre-existing co-morbidities, and targeting lockdown and protection policies to co-morbidities can multiply the benefits from targeting significantly. We leave a detailed investigation of this issue to future work.

One issue we did not address is how lockdown policies can be implemented, especially when they are heterogeneous by group and also involve various between-group social distancing elements. The first aspect of this question is that voluntary behavioral changes may already achieve some of the objectives of optimal policy (e.g., by voluntary social distancing and individuals substituting work from home). A second aspect concerns the "mechanism design" question—how these regulations can be effectively implemented. We note here that semi-targeted policies may be easier to implement because the strictest lockdowns are for older individuals and can be interpreted as a form of "protective custody" for that group, meaning that it is mostly to protect the group itself not to reduce the externalities they create on others. The same applies to measures to reduce interactions between this group and the rest of the population. Despite the importance of questions surrounding endogenous behavioral change and implementation, we leave these issues for future work as well.

We view our paper as a first step in enriching the SIR model, which has become a workhorse tool for understanding and combating the COVID-19 pandemic, by bringing

in economic effects and trade-offs that depend on differential risks in the population. As already noted, the results from our study must be taken as illustrative and interpreted with caution—and hence our greater emphasis on the qualitative patterns and the comparison of frontiers. There is much to be done to incorporate richer forms of economic and epidemiological interactions within and between different groups—especially to study how changes in economic and social incentives impact the dynamics of infections. Our conclusion that targeted policies can significantly improve social welfare also highlights the need for more granular data on COVID-19 in particular, and other diseases and epidemics more generally.

# References


**Alvarez, Fernando, David Argente, and Francesco Lippi**, "A Simple Planning Problem for COVID-19 Lockdown," Working Paper 26981, National Bureau of Economic Research April 2020.

**Atkeson, Andrew**, "How Deadly Is COVID-19? Understanding The Difficulties With Estimation Of Its Fatality Rate," Working Paper 26965, National Bureau of Economic Research April 2020.

_ , "What Will Be the Economic Impact of COVID-19 in the US? Rough Estimates of Disease Scenarios," Working Paper 26867, National Bureau of Economic Research March 2020.

**Avery, Christopher, William Bossert, Adam Clark, Glenn Ellison, and Sara Fisher Ellison**, "Policy Implications of Models of the Spread of Coronavirus: Perspectives and Opportunities for Economists," Working Paper 27007, National Bureau of Economic Research April 2020.

**Bairoliya, Neha and Ayşe İmrohoroğlu**, "Macroeconomic Consequences of Stay-At-Home Policies During the COVID-19 Pandemic," April 2020. mimeo.

**Baqaee, David, Emmanuel Farhi, Michael J Mina, and James H Stock**, "Reopening Scenarios," Working Paper 27244, National Bureau of Economic Research May 2020.

**Bayham, Jude, Nicolai V. Kuminoff, Quentin Gunn, and Eli P. Fenichel**, "Measured voluntary avoidance behaviour during the 2009 A/H1N1 epidemic," *Proceedings of the Royal Society B: Biological Sciences*, 2015, *282* (1818), 20150814.

**Beal, L.D.R., D.C. Hill, R.A. Martin, and J.D. Hedengren**, "GEKKO Optimization Suite.„" *Processes*, 2018, *6* (106).

**Brauer, Fred, Pauline van den Driessche, and J. Wu**, *Mathematical Epidemiology*, Springer-Verlag Berlin Heidelberg, 2008.

**Brotherhood, Luiz, Philipp Kircher, Cezar Santos, and Michèle Tertilt**, "An economic model of the Covid-19 epidemic: The importance of testing and age-specific policies," 2020.

**Dingel, Jonathan and Brent Neiman**, "How Many Jobs Can be Done at Home?," April 2020. White Paper.

**Eichenbaum, Martin, Sergio Rebelo, and Mathias Trabandt**, "The Macroeconomics of Epidemics," Working Paper 26882, National Bureau of Economic Research March 2020.

**_ , _ , and _** , "The Macroeconomics of Testing During Epidemics," April 2020. mimeo.

**Farboodi, Maryam, Gregor Jarosch, and Robert Shimer**, "Internal and External Effects of Social Distancing in a Pandemic," Working Paper 27059, National Bureau of Economic Research April 2020.

**Favero, Carlo A., Andrea Ichino, and Aldo Rustichini**, "Restarting the Economy While Saving Lives Under COVID-19," April 2020.

**Fenichel, Eli P**, "Economic considerations for social distancing and behavioral based policies during an epidemic.," *J Health Econ*, Mar 2013, *32* (2), 440–451.

**Ferguson, NM, D. Laydon, G. Nedjati-Gilani, N. Imai, K Ainslie, M. Baguelin, S. Bhatia, A. Boonyasiri, Z. Cucunubá, G. Cuomo-Dannenburg, and A. Dighe**, "Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand," March 2020. Imperial College COVID-19 Response Team.

**Fernández-Villaverde, Jesús and Charles J. Jones**, "stimating and Simulating a SIRD Model of COVID-19 for Many Countries, States, and Cities," 2020.

**Garibaldi, Pietro, Espen R. Moen, and Christopher A Pissarides**, "Modelling contacts and transitions in the SIR epidemics model," in Charles Wyplosz, ed., *Covid Economics Vetted and Real-Time Papers*, CEPR, April 2020.

**Garriga, Carlos, Rody Manuelli, and Siddhartha Sanghi**, "Optimal Management of an Epidemic: An Application to COVID-19. A Progress Report," April 2020. mimeo.

**Geoffard, Pierre-Yves and Tomas Philipson**, "Rational Epidemics and Their Public Control," *International Economic Review*, 1996, *37* (3), 603–624.

**Glover, Andrew, Jonathan Heathcote, Dirk Krueger, and José-Víctor Ríos-Rull**, "Health versus Wealth: On the Distributional Effects of Controlling a Pandemic," Working Paper 27046, National Bureau of Economic Research April 2020.

**Gollier, Christian**, "Cost-benefit analysis of age-specific deconfinement strategies," April 2020. presentation slides.

**Hedengren, John D., Reza Asgharzadeh Shishavan, Kody M. Powell, and Thomas F. Edgar**, "Nonlinear modeling, estimation and predictive control in APMonitor," *Computers and Chemical Engineering*, 2014, *70*, 133 – 148.

**Heesterbeek, J.A.P. and M.G. Roberts**, "The type-reproduction number T in models for infectious disease control," *Mathematical Biosciences*, 2007, *206* (1), 3 – 10. Alcala Special Issue.

**Jones, Callum J, Thomas Philippon, and Venky Venkateswaran**, "Optimal Mitigation Policies in a Pandemic: Social Distancing and Working from Home," Working Paper 26984, National Bureau of Economic Research April 2020.

**Kermack, William Ogilvy, A. G. McKendrick, and Gilbert Thomas Walker**, "A contribution to the mathematical theory of epidemics," *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 1927, *115* (772), 700–721.

**Klepac, Petra, Adam J Kucharski, Andrew JK Conlan, Stephen Kissler, Maria Tang, Hannah Fry, and Julia R Gog**, "Contacts in context: large-scale setting-specific social mixing matrices from the BBC Pandemic project," *medRxiv*, 2020.

**Kudlyak, Marianna, Lones Smith, and Andrea Wilson**, "Avoidance Behavior at the COVID19 Breakout in an SI3R Model," April 2020. mimeo.

**Manski, Charles F and Francesca Molinari**, "Estimating the COVID-19 Infection Rate: Anatomy of an Inference Problem," Working Paper 27023, National Bureau of Economic Research April 2020.

**McCallum, Hamish, Nigel Barlow, and Jim Hone**, "How should pathogen transmission be modelled?," *Trends in Ecology and Evolution*, 2001, *16* (6), 295 – 300.

**Prem, Kiesha, Alex R. Cook, and Mark Jit**, "Projecting social contact matrices in 152 countries using contact surveys and demographic data," *PLOS Computational Biology*, 09 2017, *13* (9), 1–21.

**Rampini, Adriano A**, "Sequential Lifting of COVID-19 Interventions with Population Heterogeneity," Working Paper 27063, National Bureau of Economic Research April 2020.

**Rowthorn, Robert and Flavio Toxvaerd**, "The Optimal Control of Infectious Diseases via Prevention and Treatment," Technical Report 2013, Cambridge-INET Working Paper 2020.

**Stock, James H**, "Data Gaps and the Policy Response to the Novel Coronavirus," Working Paper 26902, National Bureau of Economic Research March 2020.

**Wächter, A. and L. Biegler**, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming.," *Math. Program*, 2006, *106*, 25–57.

# A Appendix: Additional Figures

## A.1 Figures for Baseline Simulation

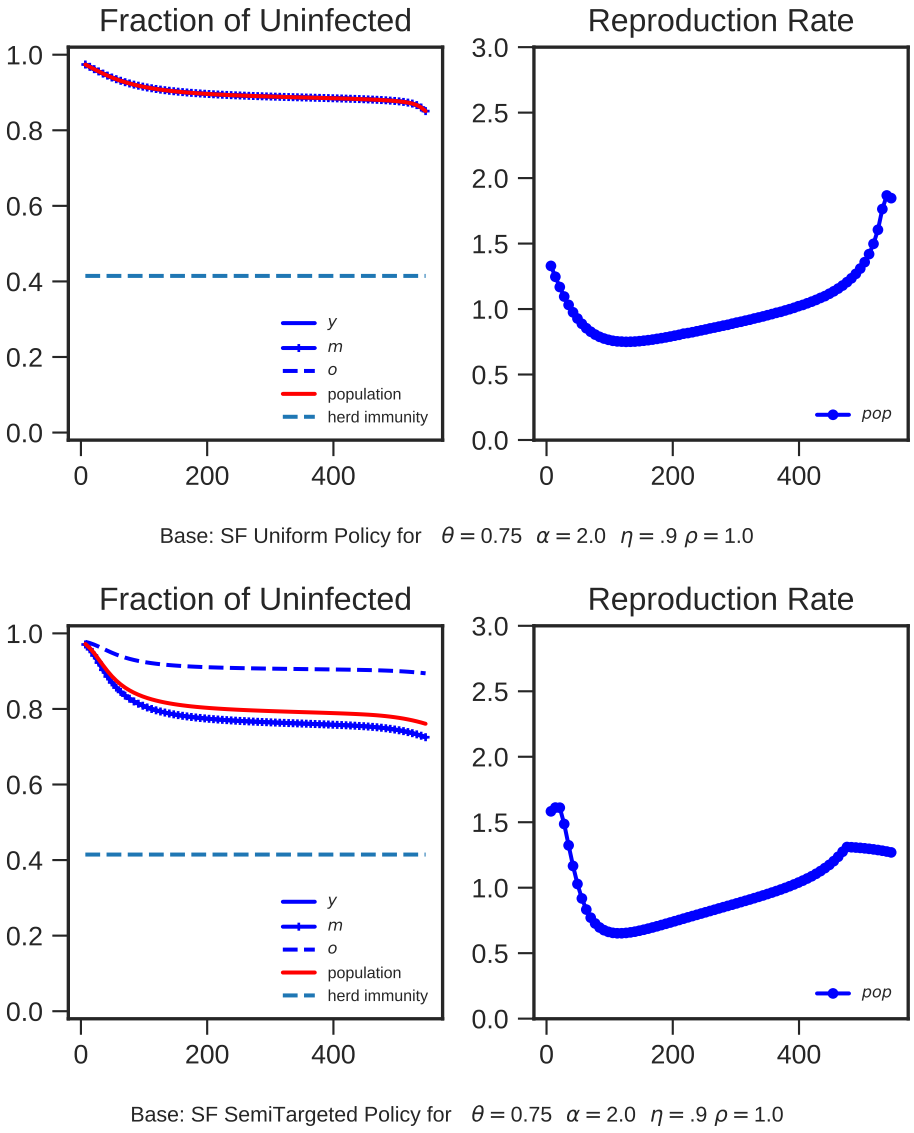


**Figure A.1:** Fraction susceptible (left panel) and the reproduction rate $R(t)$ (right panel) for baseline parameterization under uniform and semi-targeted optimal policies.

## A.2 Figures for Robustness Exercises

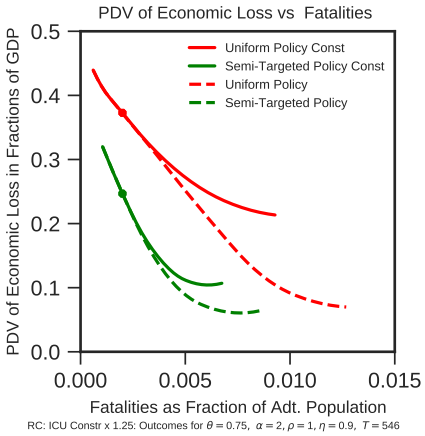


**Figure A.2:** Uniform and semi-targted frontiers when imposing a hard constraint on ICU capacity.

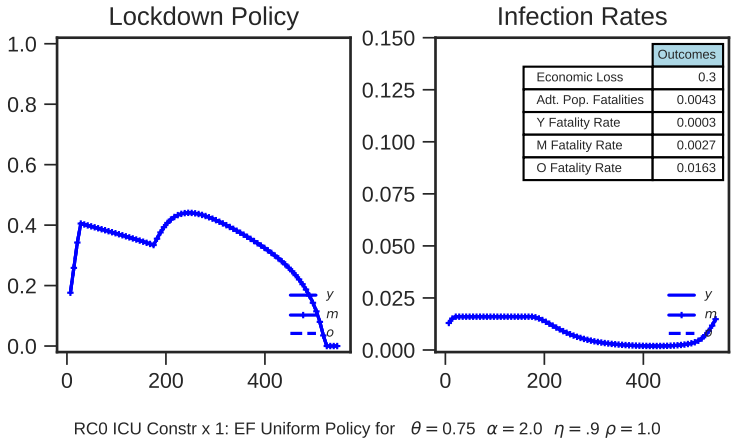


**Figure A.3:** Imposing a hard constraint on ICU capacity, optimal "economy-focused" uniform policy .

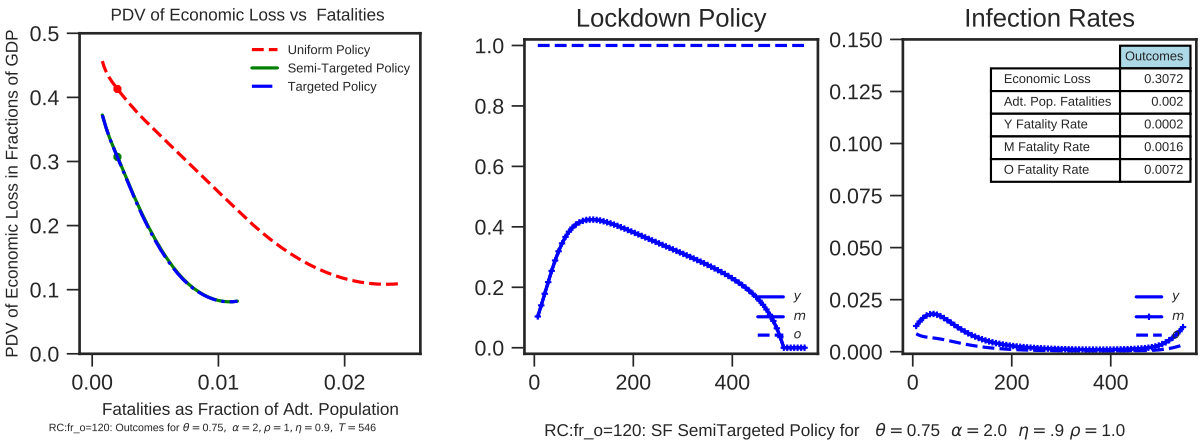


**Figure A.4:** Higher mortality rate for the elderly based on South Korean data– frontiers and optimal "safety-focused" semi-targeted policy.

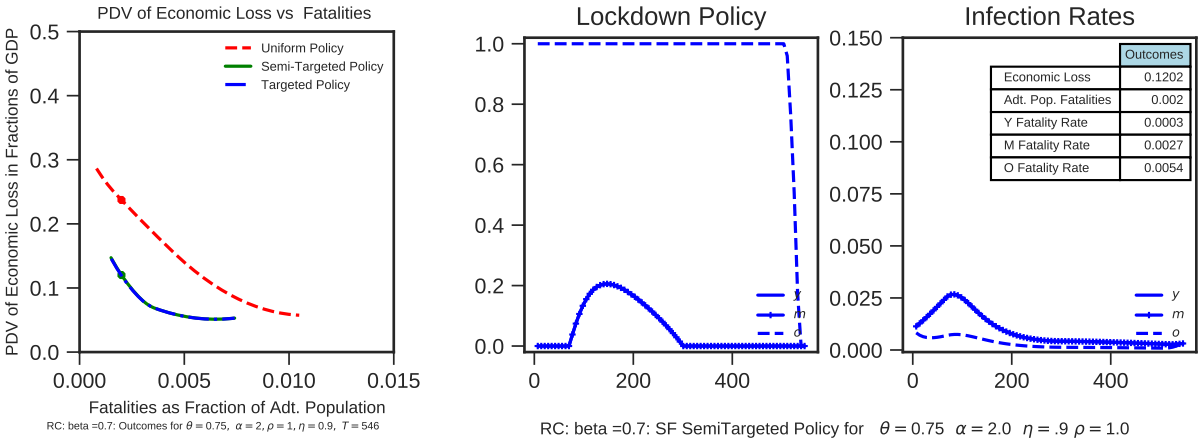


**Figure A.5:** Lower value for transmission rate $\beta$, optimal uniform and semi-targeted policies for safety-first point on the frontier.

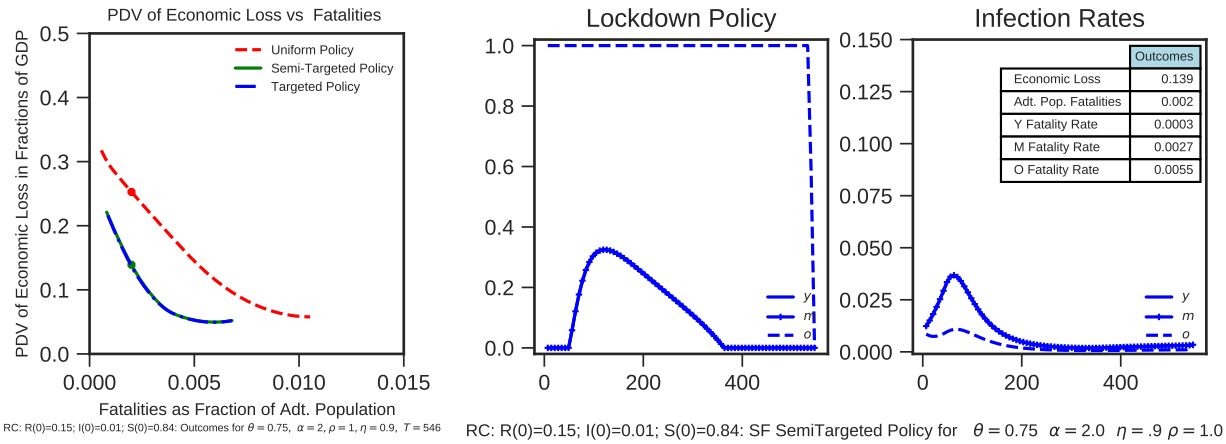


**Figure A.6:** Alternative initial condition with higher number of recovered individuals – frontiers and optimal "safety-focused" semi-targeted policy.

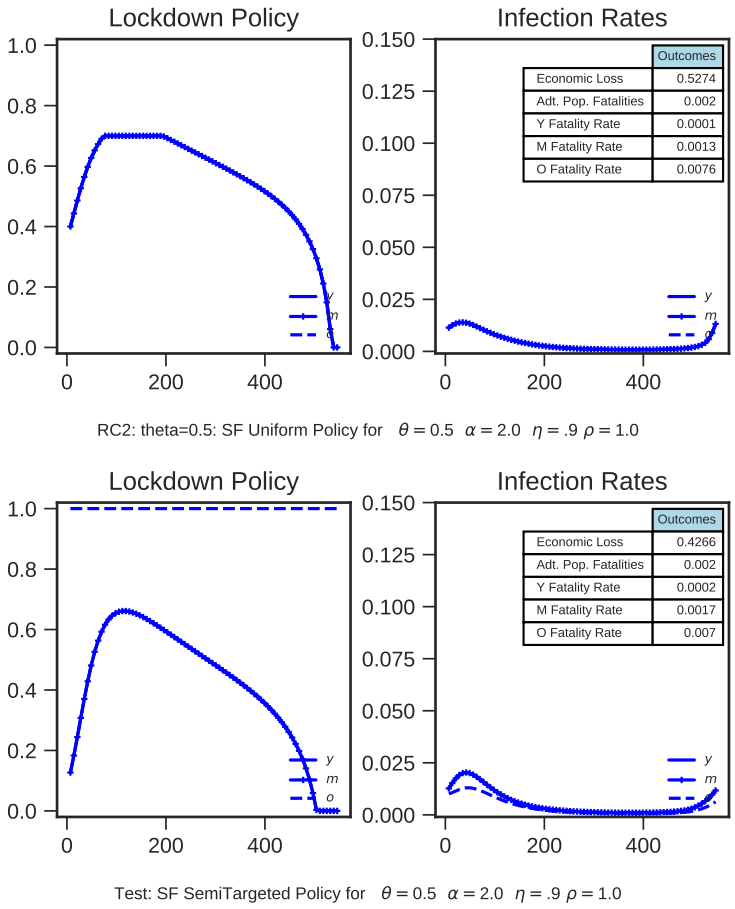


**Figure A.7:** Lower effectiveness of lockdowns, $\theta = 0.5$ – frontiers and optimal "safety-focused" semi-targeted policy.

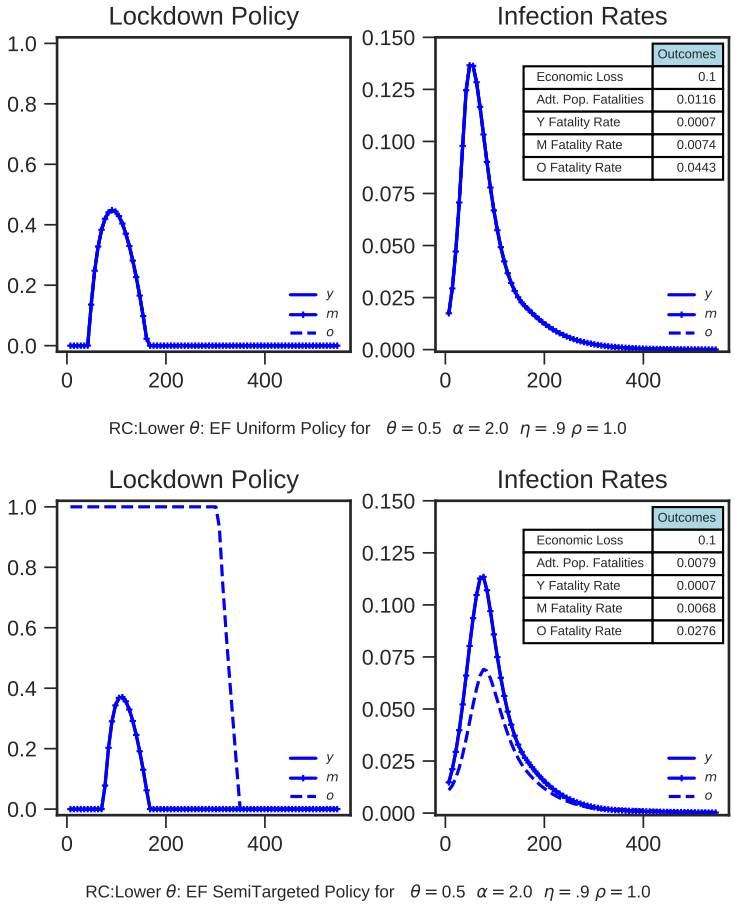


**Figure A.8:** Lower effectiveness of lockdowns, $\theta = 0.5$ – optimal "economy-focused"uniform and semi-targeted policies.

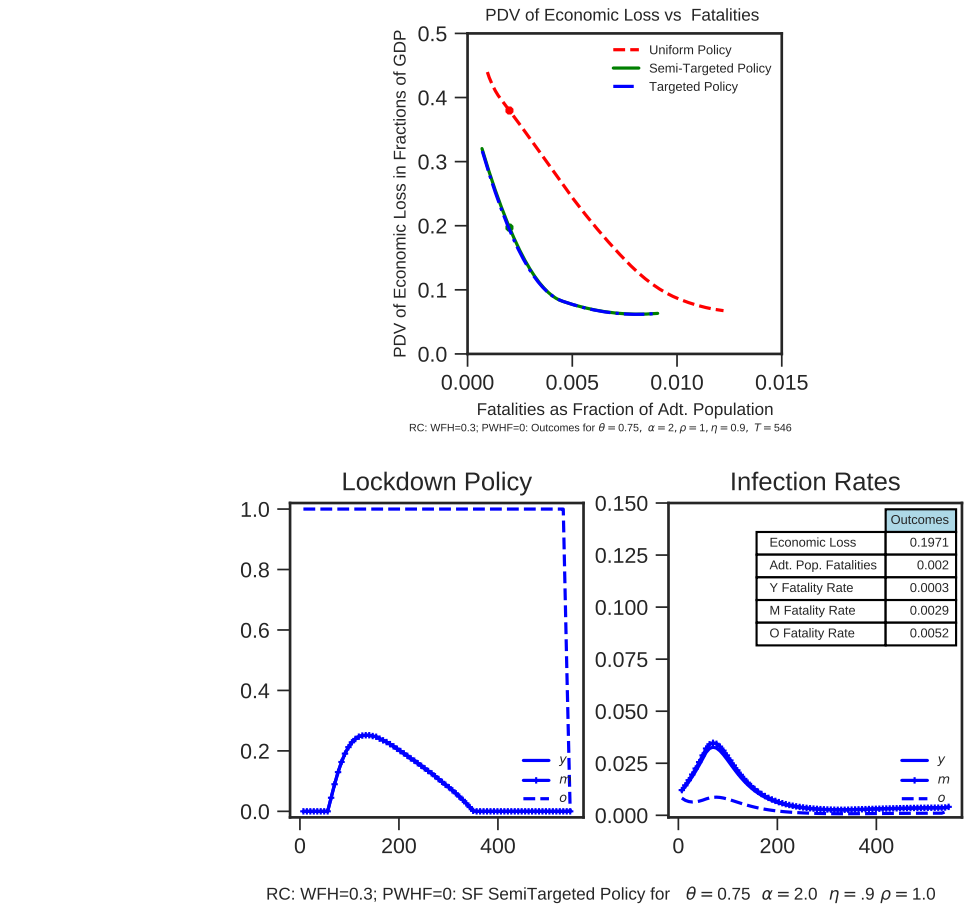


**Figure A.9:** Alternative formulation for working from home – frontiers and optimal "safety-focused" semi-targeted policy.

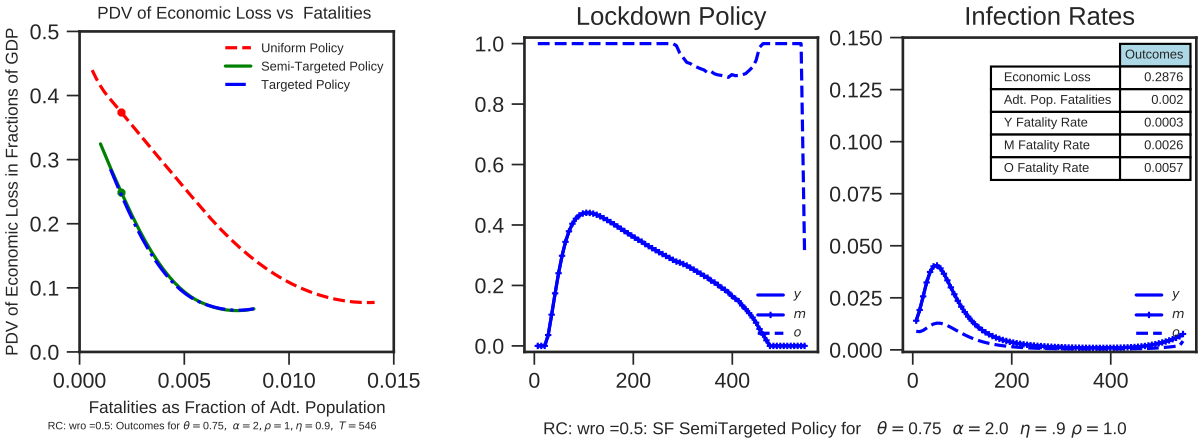


**Figure A.10:** Higher value of the old being out of lockdown – frontiers and optimal "safety-focused" semi-targeted policy.

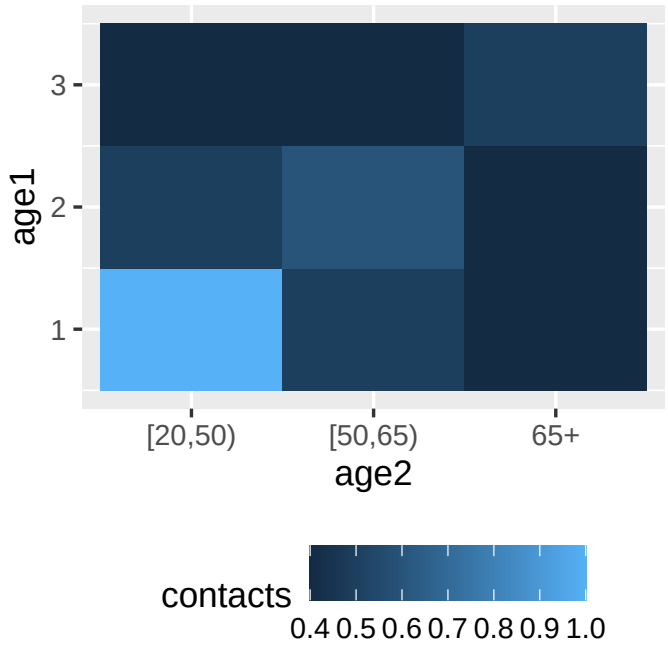


**Figure A.11:** Contact matrix.

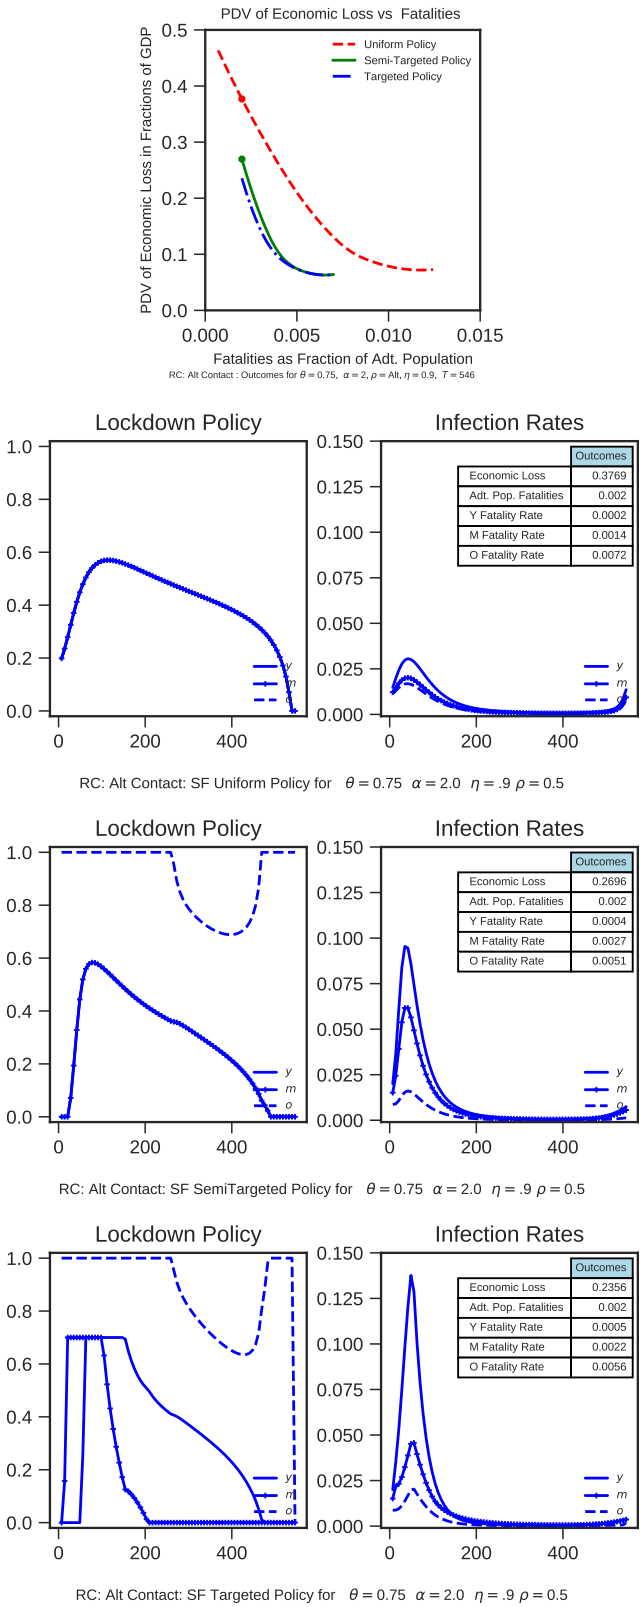


**Figure A.12:** Baseline model with calibrated contact matrix – frontiers and optimal "safety-focused" optimal policies.

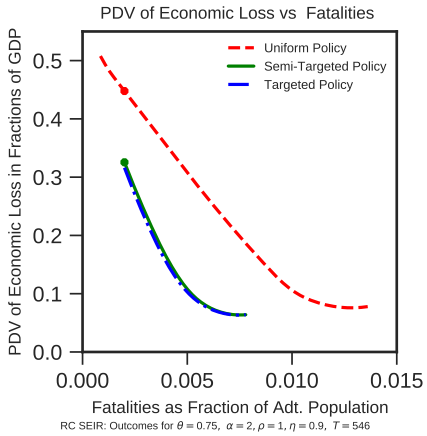


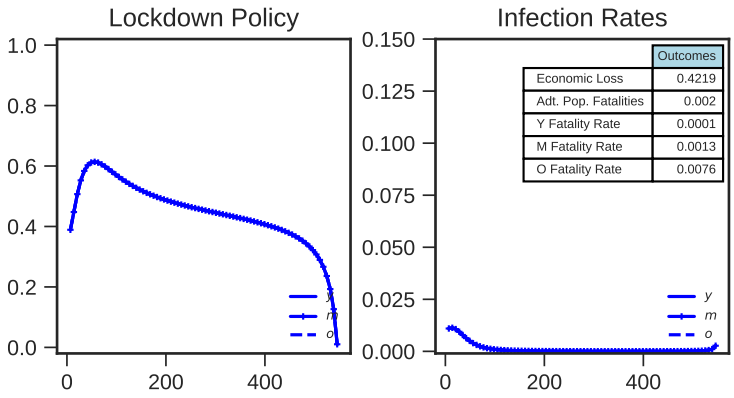


RC: SEIR : SF Uniform Policy for $\theta = 0.75$ $\alpha = 2.0$ $\eta = .9$ $\rho = 1.0$

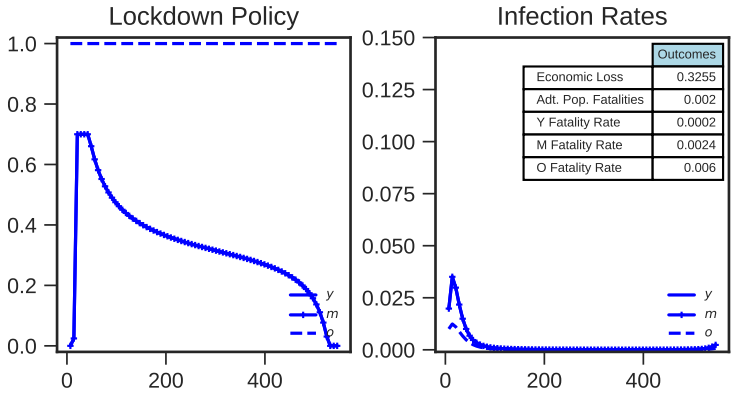


TEST SEIR : SF SemiTargeted Policy for $\theta = 0.75$ $\alpha = 2.0$ $\eta = .9$ $\rho = 1.0$

**Figure A.13:** SEIR model – frontiers and optimal "safety-focused" uniform and semi-targeted policy.

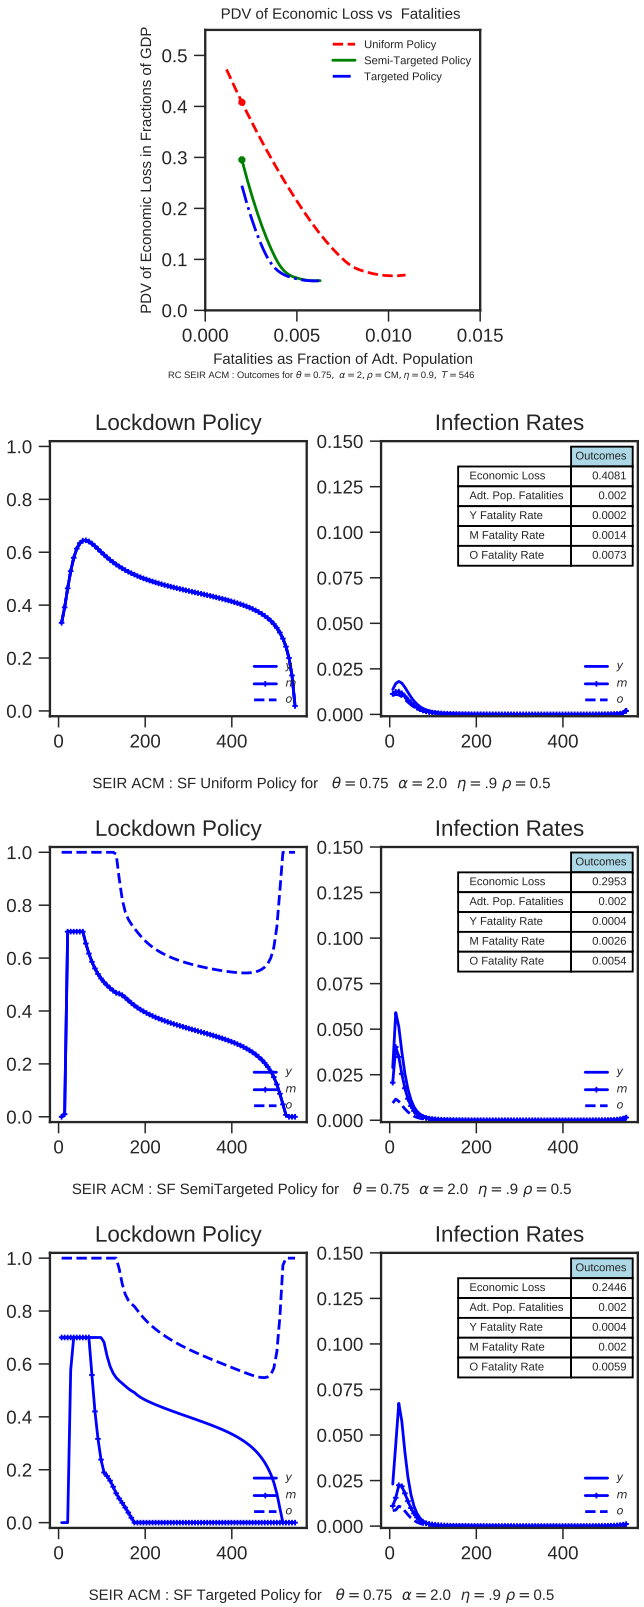


**Figure A.14:** SEIR model with calibrated contact matrix – frontiers and optimal "safety-focused" optimal policies.