NBER WORKING PAPER SERIES

ADDRESSING THE COVID-19 PANDEMIC:
COMPARING ALTERNATIVE VALUE FRAMEWORKS

Maddalena Ferranna
JP Sevilla
David E. Bloom

Working Paper 28601
<http://www.nber.org/papers/w28601>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2021

This paper was finalized in December 2020, and reflects information about COVID-19, its impacts, and the steps taken to respond to the pandemic, as of that date. The paper was prepared for inclusion as a chapter in a forthcoming Cambridge University Press volume entitled “Prioritarianism in Practice” edited by Matthew Adler and Ole Norheim. We thank the editors, James Hammitt, Nicolas Treich and participants at the “Prioritarianism in Practice” meeting on May 21–22, 2020 for extremely thoughtful comments. This work was supported by the Value of Vaccination Research Network (VoVRN) through a grant from the Bill & Melinda Gates Foundation (Grant OPP1158136). The content is solely the responsibility of the authors and does not necessarily reflect the views of the VoVRN or the Bill & Melinda Gates Foundation. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Maddalena Ferranna, JP Sevilla, and David E. Bloom. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Addressing the COVID-19 Pandemic: Comparing Alternative Value Frameworks
Maddalena Ferranna, JP Sevilla, and David E. Bloom
NBER Working Paper No. 28601
March 2021
JEL No. D6,I1,I3

ABSTRACT

The COVID-19 pandemic has forced countries to make difficult ethical choices, e.g., how to balance public health and socioeconomic activity and whom to prioritize in allocating vaccines or other scarce medical resources. We discuss the implications of benefit-cost analysis, utilitarianism, and prioritarianism in evaluating COVID-19-related policies. The relative regressivity of COVID-19 burdens and control policy costs determines whether increased sensitivity to distribution supports more or less aggressive control policies. Utilitarianism and prioritarianism, in that order, increasingly favor income redistribution mechanisms compared with benefit-cost analysis. The concern for the worse-off implies that prioritarianism is more likely than utilitarianism or benefit-cost analysis to target young and socioeconomically disadvantaged individuals in the allocation of scarce vaccine doses.

Maddalena Ferranna
Harvard T.H. Chan School of Public Health
Department of Global Health and Population
665 Huntington Ave.
Building 1, 1202-B
Boston, MA 02115
mferranna@hsph.harvard.edu

JP Sevilla
Department of Global Health and Population
Harvard T.H. Chan School of Public Health
665 Huntington Avenue
Boston, Mass 02115
jsevilla@hsph.harvard.edu

David E. Bloom
Harvard T.H. Chan School of Public Health
Department of Global Health and Population
665 Huntington Ave.
Building 1, Suite 1202
Boston, MA 02115
and IZA
and also NBER
dbloom@hsph.harvard.edu

1. Introduction

Coronavirus disease 2019 (COVID-19), which is caused by severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2), was first reported in Wuhan, China, in December 2019. Since then, COVID-19 has spread explosively around the world, with catastrophic health, social, and economic effects. Through mid-December 2020, roughly 72 million cases had been detected worldwide, including nearly 1.7 million deaths.¹ The socioeconomic costs of the COVID-19 pandemic (and of the nonpharmaceutical interventions adopted to control its spread) are staggering.² The International Monetary Fund estimates that COVID-19 is associated with an expected 4.4% contraction of the world economy in 2020 (IMF 2020). By comparison, the 2009 global financial crisis caused a 0.1% contraction of the world economy (Gopinath 2020).

Several factors underlie the dramatic consequences of COVID-19. SARS-CoV-2 is a novel virus with many features that remain substantially unknown. However, it appears to be very infectious, lethal, and geographically unconstrained. In comparison, the World Health Organization reports that the seasonal flu has an infection fatality rate (i.e., the proportion of deaths among all infected individuals) of 0.1% (WHO 2020), whereas the infection fatality rate of COVID-19 is estimated to be in the range of 0.5%–1% (Levin et al. 2020; Perez-Saez et al. 2020; Streek et al. 2020; Yang et al. 2020), although much uncertainty remains (Ioannidis 2020). Lack of adequate health sector responses (i.e., testing, treatment, and vaccination) has undermined effective efforts to contain the pandemic. However, the pandemic has also spurred an unprecedented innovations race (Bloom et al. 2021). For example, more than 270 vaccines against COVID-19 are under development around the world, with more than 60 already in the clinical testing phase, and eight already approved for full or emergency use in several countries.³ Absent effective treatments and a vaccine, countries have implemented social and physical distancing policies to curb the infections (e.g., quarantine, masks, staying 6 feet apart, and economic shutdowns), which have proven to be socially, culturally, economically, and politically severe and burdensome and which have required a drastic and perhaps permanent reorganization of large segments of society (e.g., social norms about handshakes, hugging, and kissing; patterns of working from home; virtual instead of in-person meetings; air travel; and building planning).

¹ <https://coronavirus.jhu.edu/map.html>

² The World Health Organization declared COVID-19 a pandemic on March 11, 2020.

³ <https://www.nytimes.com/interactive/2020/science/coronavirus-vaccine-tracker.html>

The COVID-19 pandemic has exacerbated existing inequalities through patterns of infections, sickness, and deaths that disproportionately affect disadvantaged populations; socioeconomic differences in healthcare access and in ability to invest resources to tackle the pandemic; and inadequate redistribution mechanisms to equitably share the economic burden of the pandemic and nonpharmaceutical interventions to control it (Avent 2020). Lack of global leadership, coupled with rising nationalism and populism, has hampered a timely, effective, and globally fair solution to the COVID-19 crisis (Hafner et al. 2020).

Management of the pandemic raises many policy questions, the answers to which depend on the ethical framework that is used to evaluate reductions in deaths and morbidity against socioeconomic costs. For example, to what extent and for how long should we restrict social interactions and limit economic activities? When is it safe to reopen schools and businesses? How much should countries spend on research and development (R&D) for a vaccine against COVID-19? How should the financing burden be divided among countries and stakeholders? How should scarce healthcare resources such as ventilators, intensive care unit beds, vaccine doses, or drugs be allocated among the population?

This chapter explores how utilitarian and prioritarian social welfare function (SWF) frameworks can help answer some of these questions. In particular, we will focus on two issues: disease control and vaccine allocation. We discuss these value frameworks in the context of simplified and stylized examples, which aim to demonstrate how pandemic response requires addressing distributional issues of the kind that SWF frameworks can inform rather than to generate real-world policy conclusions. Section 2 describes some of the core pieces of the analytical frameworks and discusses the main issues at stake in the COVID-19 crisis. We will focus on the value of policies to control the spread of the pandemic (e.g., economic lockdowns) and of prioritizing different groups in allocating scarce vaccine supplies. Section 3 qualitatively discusses the main differences among benefit-cost analysis (BCA), utilitarianism, and prioritarianism. Section 4 presents a stylized example of control interventions, and Section 5 a stylized example of vaccine prioritization strategies supported by the three value frameworks. Section 6 concludes. Although we focus on the COVID-19 pandemic, the lessons we draw generalize to other epidemics and to any intervention with health and non-health impacts.

The analysis draws two main sets of conclusions. When deciding the strength of the control policy, what counts is the relative regressivity of COVID-19 burdens and policy costs. The more regressive the burdens of COVID-19 are, the more the distribution sensitivity of the value framework raises the value of strict control policies (i.e., increased regressivity increases the value of control for prioritarianism more than it does for utilitarianism, and for utilitarianism more than it does for BCA). The more regressive the costs of strict control policies are, the more the distribution sensitivity of the value framework will lower the value of strict control policies. Thus, the net impact of the value framework's distribution sensitivity on the value of control depends on the relative strength of the two patterns of regressivity. The sensitivity of utilitarianism and prioritarianism to the distribution of policy costs implies that they will strongly support income redistribution mechanisms (e.g., unemployment benefits).

The second set of conclusions concerns vaccine allocation. None among BCA, utilitarianism, and prioritarianism prioritizes specific groups based solely on the size of the risk of severe health outcomes from COVID-19 that they face. BCA tends to favor the wealthy because they have the larger ability to pay and thus the larger willingness to pay (although, in practice, benefit-costs analysts often use unit values that are invariant to income). Utilitarianism and prioritarianism skew the allocation toward the poor because they bear a disproportionate burden of COVID-19. Because the young and the poor are among the worse-off in the population, prioritarianism is more likely to target young, socioeconomically disadvantaged populations in allocating COVID-19 vaccines than are utilitarianism and BCA. However, because the difference in COVID-19 mortality risks between the old and the young is so high, even prioritarianism will rank socioeconomically disadvantaged elderly at the top of the vaccine priority list (even though wealthy elderly may be ranked lower than the socioeconomically disadvantaged young).

A note of caution before we proceed. The current chapter was completed in December 2020. Although the evaluation tools discussed herein have enduring relevance, their application to COVID-19 is based on the information available as of December 2020 and should be read and understood in that light.

2. Background

2.1 Concept of well-being

The main goal of the chapter is to discuss the evaluation of COVID-19 interventions through the lens of SWFs, specifically utilitarian and prioritarian ones. Under the SWF approach, assessing the value of an intervention involves two steps: (i) estimating the intervention's impacts on all the attributes of individual well-being and (ii) aggregating individuals' well-being gains and losses to assess the intervention's impact on overall social welfare. The specific aggregation rule depends on the adopted SWF (Adler 2019).

We start by specifying a measure of individual well-being. We assume individual well-being depends on attributes including consumption/income, longevity, and health status. Specifying a well-being measure and its dependence on its attributes can be done in several ways. Adler and Decancq (forthcoming) review the most common methods. We rely on utility functions that represent individuals' preferences and risk attitudes regarding alternative probability distributions of attributes over a lifetime. We interpret the period and lifetime utility functions of microeconomic theory as interpersonally comparable ratio-scale indices of individual well-being (with meaningful zeros and defined up to a common multiplicative constant). We therefore treat "well-being" and "utility" interchangeably.

We assess policy impact on lifetime rather than sub-lifetime utility, which is standard in the literature and in this volume. We assume that lifetime utility is additive in period (e.g., annual) utility, there is no time discounting, and the marginal utility of income (or consumption) is diminishing so that a dollar raises the utility of the poor more than it does that of the rich. The time additivity and zero discounting assumptions jointly imply that marginal utility is linear in (expected or realized) longevity. Therefore, all else equal, longevity increments have the same positive effect on lifetime utility independently of the age of the individual.⁴ The time additivity

⁴ Suppose that a 20-year-old and a 70-year-old are going to die soon, but they can both live five more years if they take an experimental treatment. The quality of life of those additional five years is the same for both individuals. Time additivity and zero discounting imply that the two individuals attach the same value to the life-extension treatment. If marginal utility were decreasing in longevity, then the value of the treatment for the 20-year-old would be larger than the value of the treatment for the 70-year-old on the grounds that the life of the former is shorter.

assumption, although standard in the literature, bears on the comparison between utilitarianism and prioritarianism and, in particular, on the weight given to the young in policy evaluation.⁵

2.2 Impact of COVID-19 on well-being

The COVID-19 pandemic has health and non-health negative impacts on individuals' well-being. The health burden includes the probability and severity of infections, which can lead to death, temporary and long-term physical disability, and temporary or long-term mental health burdens in patients and in friends and family. Health-system-related burdens include treatment and other costs; congestion of intensive care units and other resources leading to rationing, delay, and disruption of medical and public health services (including immunization programs critical to the health and the cognitive and physical development of children around the world, especially in developing countries); and out-of-pocket expenditures on health. The non-health impacts of COVID-19 include the potential loss of income due to work stoppage, unemployment, furlough, or sickness.

The health impacts of COVID-19 are strongly associated with pre-existing health and socioeconomic status. Some demographic groups seem to be disproportionately exposed to risks of COVID-19 complications and deaths, particularly the elderly, those with underlying health conditions, and males (e.g., Guan et al. 2020; Zhou et al. 2020; Chow et al. 2020). Risk of infection seems higher for low-socioeconomic (SES) individuals and communities due to their overrepresentation in essential jobs and service jobs requiring physical presence (and disallowing working from home), more crowded living arrangements, higher probability of using public transportation, and lower financial ability to stockpile food (Blundell et al. 2020; Brown and Ravallion 2020). Risk of infection also seems to be higher in polluted neighborhoods (Wu et al. 2020) and in areas with less access to healthcare resources. In the United States, COVID-19 deaths are overrepresented among African Americans, Hispanics, and Native Americans (Subbaraman 2020).

⁵ As discussed later, if we retain the time additivity and zero discounting assumptions, under utilitarianism a given increase in longevity for the young has the same value as an equal increase in longevity for the old, while that is not true for prioritarianism. If we remove the time additivity and zero discounting assumptions, both utilitarianism and prioritarianism would give greater priority to increases in longevity for the young, although more for the latter than for the former.

COVID-19 can have direct negative economic impacts on infected individuals and their families, e.g., through medical bills, salary losses due to sick days or to taking care of sick family members, and decreased productivity due to the onset of a long-term disability. Social distancing to control infections or reduce exposure, whether mandated or voluntary, creates profound macroeconomic costs. People who can afford to do so will spend more work and leisure time at home, which can shift the pattern of demand (e.g., from restaurant meals to online groceries, or from office space to home office supplies). Overall levels of consumption and demand may also fall if income is reduced (e.g., among the unemployed) or if people are saving for precautionary reasons or in anticipation of spending more after the pandemic. In turn, such reduced demand can lead to unemployment or furloughs, particularly in some contact-intensive sectors (e.g., restaurants); bankruptcy and business closures; and increased exposure to socioeconomic strains and the risk of poverty among the most vulnerable groups. Moreover, outbreaks in the workplace will prevent normal functioning of the locale, possibly threatening its own existence and the health and well-being of its employees. In addition, outbreak-induced interruptions in global supply chains can have worldwide negative effects, even in regions not directly affected by the outbreak (UNIDO 2020).

2.3. Control policies

2.3.1. Flattening the curve, suppression, and mitigation: disease control basics

Since the worldwide outbreaks of COVID-19 began in winter 2019/2020, epidemiologists have warned that the best immediate course of action to reduce the burden of COVID-19 is to “flatten the infection curve” through social and physical distancing (Roberts 2020). Figure 1 depicts the idea. The steep curve represents the projected number of infections over time without any measure to control the spread of the disease (e.g., handwashing or economic lockdown), while the flat curve represents the projected number of infections over time when social and physical distancing measures are enforced. Absent any intervention, existing epidemiological models predict a steep surge of infections well beyond countries’ healthcare system capacity for treatment, represented by the thick horizontal line (e.g., Ferguson et al. 2020, Giordano et al. 2020, Kissler et al. 2020). Measures to control the spread of the disease aim to reduce this surge and thereby contain the

demand for scarce healthcare resources. Implementing these measures results in a curve with a flatter slope, denoting a slower rate of infection extended over a longer time.

Figure 1. Number of COVID-19 cases over time with and without nonpharmaceutical interventions and projected rationing

Note: The steep curve represents the projected number of cases over time without nonpharmaceutical intervention to control the spread of the disease (such as handwashing or economic lockdown); the flat curve represents the projected number of cases over time with those measures. The horizontal lines represent the healthcare system capacity constraints, and the arrows represent the peak rationing of healthcare resources (i.e., the number of cases that cannot receive treatment). Implementing nonpharmaceutical interventions and relaxing the healthcare system capacity constraint are both useful strategies to reduce the required rationing.⁶

Flattening the curve achieves three main goals. First, it reduces the daily number of positive cases, thereby preventing the healthcare system from being overwhelmed with cases; allowing more people to get treatment; and, hopefully, reducing the number of deaths. Second, flattening the curve diminishes the total number of cases and deaths throughout the pandemic by allowing more people to get treated daily and by reducing the probability of spreading the disease. Therefore, not

⁶ We thank Daniel Cadarette for creating the picture.

only will fewer cases and deaths occur per day, but the total number of cases and deaths is also expected to reduce, despite the longer time that the epidemic takes to be resolved. Finally, flattening the curve buys some time to develop and test a vaccine or new therapeutic drugs. Expanding the supply of scarce healthcare resources (i.e., moving from the thick horizontal line to the dotted one) is a complementary strategy to flattening the curve. Interventions to relax the healthcare system capacity constraint include, for example, developing faster and more accurate tests for COVID-19 and for COVID-19 antibodies, postponing elective procedures, and easing supply chain bottlenecks and financial constraints to improve the distribution of resources.

Strategies to control the course of the pandemic can be divided into two broad categories: suppression and mitigation (Yglesias 2020). Suppression aims at keeping the number of infections as close to zero as possible. Mitigation allows a controlled outbreak to occur, with the aim of keeping the number of cases within the healthcare system capacity. Stay-at-home orders and economic lockdowns suppress the infection as long as they are in place. Contact tracing, quarantine of positive cases, and social distancing of the elderly are examples of measures to mitigate the infection without completely suppressing it. A useful reference point for delineating the control strategy has been the so-called reproduction number, labeled R , which measures the expected number of infections that one case directly generates (Adam 2020). When $R < 1$, the outbreak is shrinking. Mitigation strategies that aim at slowing down transmission without necessarily ending it may have $R > 1$; suppression strategies always result in $R < 1$.

Suppression measures could delay the outbreak, spare many lives, and buy time until a vaccine or treatment is developed. However, suppression strategies are also very costly, potentially triggering massive unemployment, business closures, educational setbacks, and worsening mental health. In addition, if suitable mitigation policies do not follow once they are lifted, a new surge of cases is likely to occur as herd protection has typically not yet been achieved.⁷ Mitigation strategies, which may be used with or instead of suppression, are less invasive as they do not rely on complete economic shutdown. However, they require extensive testing and contact tracing to control the

⁷ Herd protection occurs when the number of immune individuals in a population is sufficiently large that nonimmune individuals are also protected from infection thanks to reduced likelihood of transmission. For COVID-19, the herd protection threshold is estimated to be at least 60% in the absence of any intervention (Fontanet and Cauchemez 2020).

pandemic successfully and the political and technological capacity to quickly revert to more aggressive suppression strategies if the outbreak gets out of control.

After the initial outbreak of COVID-19, most countries around the world implemented economic shutdowns and placed strict limitations on individuals' mobility to suppress the pandemic and ease overloaded healthcare systems.⁸ Once the outbreak was considered under control, countries started lifting stay-at-home orders and slowly reopening their economies, with the underlying goal of moving from a suppression to a mitigation phase. Sweden is the famous outlier, with its decision to not suppress the outbreak through school and business closures, instead relying on social distancing and non-mandatory guidance to reach herd protection in a controlled way (Vogel 2020). As we write this chapter, the relative success of these alternative strategies is unclear.

2.3.2. Health-income trade-offs

Some argue that a critical issue in COVID-19 control is how to trade off health and income, while others argue that effectively no such trade-off exists (Mahoney 2020). Under the trade-off view, societies must choose between a “low-health, high-income” state achieved with less aggressive control and a “high-health, low-income” state achieved with more aggressive control. The preferred state depends on the distribution of gains and losses in health and income produced by the control policies required to get to those states.

In contrast, under the no-trade-off view, an uncontrolled pandemic will ultimately lead to economic unravelling, so the surest way to protect income is to protect health first. In such a case, the only two states to choose from are a “high-health, high-income” state brought about by aggressive control and a “low-health, low-income state” from less aggressive control. The plausibility of the no-trade-off view rises if the economic costs of inaction are high, if individuals stringently self-protect, and if the policy horizon is long.

The experience in the first few months of the pandemic showed that control measures can meaningfully reduce the number of infections and deaths due to COVID-19, often at the expense

⁸ <https://www.imf.org/en/Topics/imf-and-covid19/Policy-Responses-to-COVID-19>

of staggering social and economic costs. The more aggressive the control is, the higher the costs seem to be, suggesting a trade-off.

However, this appearance of a trade-off can be misleading for several reasons. First, it fails to recognize that the “do nothing” scenario (i.e., a policy of uncontrolled spread of the virus) can also have huge negative economic impacts beyond the dreadful burden of many fatalities. As described earlier, these costs relate to the endogenous social distancing implemented by individuals for fear of infection and the disruptions in global supply chain. If we account for the economic costs of inaction, the stark trade-off between health and income (or between “saving lives” and “saving livelihoods”) is mitigated, if not completely gone.

The length of the policy horizon matters as well. A stringent control policy enacted at the beginning of the pandemic (e.g., with school and business closures and shelter-at-home orders) allows countries to gain quick control over the spread of the virus and to transition to more lenient mitigation policies in the medium run (based, e.g., on extensive testing, contact tracing and quarantine, physical distancing, and adoption of protective behaviors like wearing masks and washing hands) while keeping the economy open. Such a policy scenario entails staggering economic losses in the short run when the economy shuts down, but much lower costs in the medium to long term when the economy is open and the pandemic is under control. Specifically, the overall costs are likely to be lower than in a no-action scenario with a lot of fatalities and a recession triggered by the panic and the burden of disease. The overall costs are also likely to be lower than in a scenario where virus spread is initially left uncontrolled, panic spreads, and aggressive policies are later introduced to stop the virus and to control the economic downturn. A short-sighted policy maker will focus only on the short-run stark trade-off between lives and livelihoods. However, for any reasonably long policy horizon, saving lives is also good for the economy and for preserving individuals’ livelihoods.

The effectiveness of such a narrative depends on many factors. The virus can be kept under control without resorting to lockdowns (or limiting economic lockdowns to short periods of time) only if a country has an efficient testing, contact tracing, and quarantine system; constraints on healthcare system capacity have been relaxed (e.g., the number of hospital beds); and access to healthcare resources is constant and equitably shared (e.g., easy and affordable access to testing). The compliant behavior of the population matters as well, particularly the spread of precautionary

measures such as handwashing and wearing a mask. Additionally, complementary redistributive measures such as higher unemployment benefits and recovery grants for businesses can relieve the economic burden of control strategies on individuals and businesses and ensure that no effective trade-off between saving lives and saving livelihoods occurs.

In particular, as the burden of COVID-19 is likely to exacerbate pre-existing inequities in health, income, and education, so do the policies to control its spread. Younger workers, women, and low-SES individuals have primarily felt the costs of economic lockdowns and social distancing because they are more likely to work in the high-contact sectors that the policies mainly affect (Dingel and Neiman 2020). Women tend also to bear the brunt of childcare and home schooling. In many places, school closures are likely to widen the educational gap between high- and low-SES pupils, due to the difficulties in effective online learning that the latter face (Andrew et al. 2020). Retired people may have been protected from the economic hardships of control policies, but they suffer from social isolation and deteriorating physical and mental health (Armitage and Nellums 2020). These inequities are going to be even more apparent in low-resource settings.

Thus, context also matters. How best to handle the pandemic and its health and economic costs has no single recipe. The best compromise between saving lives and saving livelihoods depends, among other factors, on the amount of healthcare resources available to a country, its capacity to provide economic relief, overall access to sanitation, and the ability to practice social and physical distancing effectively (e.g., the state of living arrangements).⁹ Aggressive lockdown policies coupled with a poorly equipped social security system and overcrowded living arrangements may have moderate effects on the number of infections, shrink the economy, cause extreme food insecurity, and prevent people from accessing basic healthcare services, threatening not only the livelihoods but also the lives of the poorest populations. In this case, saving livelihoods seems to be a precondition for saving lives (e.g., Brown et al. 2020; Cash and Patel 2020).

2.4 The value of vaccines

The staggering health and non-health burdens of COVID-19 and of the interventions to control its spread highlight the broad benefits of having a vaccine for COVID-19 (Bloom et al. forthcoming).

⁹ As well as the evaluation framework, of course.

The vaccine will reduce the morbidity and mortality impacts of the disease. It will also avert the implementation of control measures and their associated costs, e.g., massive unemployment, business closures, disruptions in global supply chains, and increased learning gaps between rich and poor children.

As of December 2020, more than 270 R&D efforts have been undertaken around the world to find a vaccine against SARS-CoV-2, with more than 60 vaccines in human trials and eight candidate vaccines approved for full or limited use in several countries, including Canada, China, Russia, Singapore, the United Kingdom, and the United States.¹⁰ Such a race raises several issues, from the trade-off between safety and speed to the most efficient and fair way of organizing the development, manufacture, and distribution of vaccines (Bloom et al. 2021).

An important topic concerns the allocation of vaccine doses among the population, given that manufacturing capacity constraints will make it a scarce resource, at least at the beginning of the distribution process. The COVID-19 crisis has seen growing “vaccine nationalism,” in which countries take a “my nation first” approach to developing, manufacturing, and distributing a vaccine and other treatments against COVID-19. Against that has been a surge of appeals to overcome nationalistic forces and to make the vaccine available everywhere in the world and to all socioeconomic groups, independently of countries’ and individuals’ ability to pay for it (Emanuel et al. 2020a).¹¹

Several guidelines for COVID-19 vaccine prioritization strategies within countries have been proposed (e.g., JCVI 2020; NASEM 2020; Toner et al. 2020).¹² These guidelines share the ethical objective of maximizing some kind of overall good (e.g., individuals’ aggregate well-being, or cumulative health and non-health benefits), while accounting for the specific epidemiological

¹⁰ <https://www.nytimes.com/interactive/2020/science/coronavirus-vaccine-tracker.html>

Two of these vaccines (the Pfizer-BioNTech one and the Moderna one) use novel technologies that have never been approved for use in humans before (messenger RNA). Several of the other most promising vaccines are also using methods that are at the technological innovation frontier (e.g., viral vector vaccines).

¹¹ For instance, more than 140 world leaders and experts called for a “people’s vaccine” on May 14, 2020, demanding that COVID-19 vaccines, diagnostics, tests, and treatments be provided free of charge to everyone and everywhere in the world:
https://www.unaids.org/en/resources/presscentre/featurestories/2020/may/20200514_covid19-vaccine-open-letter

¹² Debates about the allocation of other scarce medical resources within countries have occurred as well, especially at the beginning of the pandemic, when the explosive pace of infections overwhelmed hospitals in some locales, forcing hard choices of whom to treat and how to allocate resources such as ICU beds and ventilators (e.g., Emanuel et al. 2020b; Bloom et al. 2020).

characteristics of COVID-19 (e.g., age differences in severity and risk). Most guidelines seek also to be sensitive to equity concerns, in particular to pre-existing health and economic inequities and to the apparent correlation between burdens of COVID-19 and structural inequalities (Schmidt et al. 2020).

One hotly debated issue is whether essential workers within the healthcare sector and service industries should be prioritized over older adults and individuals with serious medical conditions, despite the fact that the latter face a considerably higher risk of severe health consequences (Goodnough and Hoffman 2020). Epidemiological, economic, and ethical reasons have been suggested to support this argument (Ferranna et al. 2021). Essential workers are at highest risk of exposure and transmission of the virus because their work often involves unavoidable high-frequency indoor proximity or interaction with others. If the vaccine can block transmission from vaccinated individuals (an uncertain feature of the approved vaccines as we write this chapter, see, e.g., Peiris and Leung 2020), then vaccinating essential workers first may be more beneficial. Economic reasons for prioritizing essential workers refer to their instrumental value in keeping essential services in the economy open (in addition to the instrumental value in saving other people's lives). Ethical reasons appeal to reciprocity and compensation for the work done by these workers in keeping the economy open while overexposing themselves to the risk of infection. Social justice arguments have also been raised highlighting the fact that many essential workers belong to disadvantaged racial and ethnic minorities and socioeconomically vulnerable groups.

3. Value frameworks to evaluate control and vaccine allocation strategies

An emergent and rapidly growing economic literature assesses the value of various measures to control the spread of COVID-19 and to determine the best course of action. Some of the questions this literature addresses are the optimal extent and length of economic shutdowns (how many sectors of the economy to close and for how long), the timing and process of reopening the economy, and the opportunity for age-targeted interventions.¹³ In terms of evaluation methods, two main approaches have emerged: BCA, using value-of-statistical-life measures (see, e.g., Acemoglu et al. 2020, Alvarez et al. 2020, Favero et al. 2020, Gollier 2020), and utilitarian SWFs

¹³ See, for instance, Bloom, Kuhn and Prettnner (forthcoming) for a review of the literature on the macroeconomic costs of COVID-19 and policies to control its spread.

(see, e.g., Eichenbaum et al. 2020, Glover et al. 2020, Jones et al. 2020, Quaas et al. 2020). Only a couple of papers adopt a prioritarian perspective in evaluating COVID-19 control policies (Adler et al. 2020 and Adler 2020).¹⁴

The literature on vaccine allocation, instead, is based mainly on health metrics. While ethical guidelines for vaccine prioritization recommend also considering non-health outcomes and being sensitive to the burden of COVID-19 impacts among socioeconomically vulnerable groups, most modeling exercises focus on strategies that maximize some health benefits, e.g., deaths averted or years of life lost averted (e.g., Bubar et al. 2021; Buckner et al. 2020; Rodríguez et al. 2020). Note that maximizing (quality-adjusted) years of life is tantamount to a utilitarian approach where well-being is assumed to depend only on (quality-adjusted) longevity (see also Cookson, Norheim, and Skarda forthcoming).

In broad terms, the main difference among BCA, utilitarianism, and prioritarianism concerns the relative weight attached to the policy impacts experienced by different individuals.¹⁵ BCA relies on the principle that “a dollar is a dollar is a dollar.” BCA evaluates a given policy by converting all its health and economic impacts into monetary equivalents based on individuals’ willingness to pay (WTP) to avoid these impacts and then summing up these equivalents. Consequently, the method is unconcerned with the distribution of willingness to pay in the population: health and economic impacts experienced by low-income individuals have the same moral importance as impacts experienced by wealthy individuals, where impacts are measured by individuals’ willingness to pay.

The “dollar is a dollar is a dollar” principle has two main consequences. First, because willingness-to-pay measures are inextricably linked to individuals’ ability to pay, the interests of the better-off count more than the interests of the less well-off. Wealthy individuals are likely to offer more to reduce the risk of COVID-19 infection and death than their less-wealthy counterparts not because they value life more, but because they have more resources and are thereby more inclined to trade wealth for a given change in health. BCA might value saving 100 rich lives more than saving 100 poor lives because the total willingness to pay of the former is considerably larger than the total

¹⁴ Adler et al. (2020) is based on the COVID-19 simulator created by Marc Fleurbaey and co-authors (<https://sites.google.com/site/marcfleurbaey/Home/covid>).

¹⁵ See also Adler (2019), chapter 5, and Hammitt and Treich (forthcoming) for a comparison of BCA, utilitarianism, and prioritarianism.

willingness to pay of the latter. As an extreme case, BCA might condone a control policy that saves 100 rich lives from COVID-19 while condemning 100 poor lives to starvation and death. Additionally, BCA would allocate scarce medical resources (including vaccines) to wealthy individuals first, on the grounds that they have a larger willingness to pay than poorer individuals for the same expected improvement in health through immunization.

Note that, to avoid these ethically objectionable results, benefit-cost analysts in practice often use a constant willingness to pay throughout the population, independent of individuals' characteristics such as income or age (Robinson 2007). The analysis in the chapter focuses on “textbook” BCA, which allows for heterogeneous willingness to pay across the population.

Second, because BCA attaches the same value to dollars paid by the rich and dollars paid by the poor (according to the “a dollar is a dollar is a dollar” rule), the distribution of the (net) policy costs plays no role in the evaluation process. Willingness-to-pay measures express how much individuals are ready to pay to reduce their own risk of contracting the virus and suffering the pandemic costs. But, if the policy is implemented, some individuals might end up paying more than their willingness-to-pay threshold, while supporting the risk reduction of other individuals. In particular, absent adequate support policies (e.g., unemployment benefits, paid sick leave, and support for virtual learning), disadvantaged populations, like ethnic minorities and low socioeconomic groups, will mainly feel the costs both of COVID-19 and of the nonpharmaceutical interventions to control it. It is reasonable to assume that a control policy that saves lives while condemning thousands to poverty should rank lower than another control policy that saves the same number of lives, costs the same amount of money in aggregate, but whose burden is felt mainly by the wealthy. However, BCA would be indifferent between the two policies. The same principle holds for the financing of alternative vaccine allocation strategies. BCA is indifferent about who pays for the vaccine doses, whether they are privately paid by single individuals or paid through tax revenues (i.e., sharing the costs collectively).

In contrast, utilitarianism relies on the principle that “a util is a util is a util,” where “util” is a quantum of well-being. The utilitarian SWF evaluates policies based on their impacts on individuals' total expected well-being. A known criticism of utilitarianism is its indifference to the distribution of well-being across the population. Consequently, the utilitarian approach is

indifferent to whether policy benefits go mainly to the well-off or the worse-off, as long as it leads to an increase in the sum of individuals' well-being.

Although the utilitarian approach is insensitive to inequalities in individual well-being levels, it is sensitive to inequalities in individual income because of the diminishing marginal utility of money assumption. Thus, policies that reduce the income of the poor are valued less than similar policies that induce an identical reduction in the income of the rich. Applied to COVID-19, control policies whose costs are paid mainly by the high-income quintiles are considered better than identical policies (in terms of infection and fatality reductions) whose costs are paid mainly by the low-income quintiles. In contrast, BCA is indifferent to the distribution of the policy costs. Concerns about the incidence and distribution of policy costs have arisen, for instance, in some developing countries, where lack of suitable redistribution mechanisms (e.g., unemployment benefits) make policy costs fall harder on the economically worse-off.¹⁶

Additionally, because utilitarianism values policy changes in terms of their impact on individual well-being rather than in terms of the individuals' willingness to pay for those changes, ability to pay is implicitly accounted for (in other words, utilitarianism can be thought of as a form of WTP-driven assessment where willingness to pay is adjusted to reflect differential abilities to pay). As a result, the approach does not inflate the policy benefits accruing to the well-off relative to the worse-off merely because money has relatively lower marginal value for the well-off.

A further consequence is that the utilitarian optimal allocation of vaccines is only sensitive to differences in well-being impacts of the vaccine; individual differences in the marginal utility of money (which reflect differences in their ability to pay) do not, as such, change the optimal allocation. In other words, the framework prioritizes individuals that benefit more in well-being terms from being vaccinated. This will typically include those that face the higher burden of COVID-19 (e.g., those with a high COVID-19 fatality risk). However, the framework is sensitive also to life expectancy (the larger the number of years one is expected to live if vaccinated, the larger the increase in well-being) and to the quality of life preserved.

Finally, prioritarianism relies on the principle that “a priority-weighted util is a priority-weighted util is a priority-weighted util”; i.e., unlike utilitarianism, the method is concerned with the

¹⁶ See, for instance, Cash and Patel (2020).

distribution of the individual well-being impacts of the pandemic, attaching higher value to well-being increments that accrue to the worse-off than to identical well-being impacts that accrue to the better-off. Applied to the evaluation of COVID-19 control and vaccine allocation strategies, the prioritarian approach favors interventions with the overall effect of benefiting mainly the worse-off in the population.

Two main forms of prioritarianism can be considered, which differ in the definition of the “worse-off” (Adler forthcoming). *Ex-ante prioritarianism* is concerned with the distribution of *expected* lifetime well-being across the population. From an ex-ante point of view, the worse-off individuals are those who in expectation have poor lifetime prospects. For example, individuals with low expected longevity, individuals who are expected to have low income, and individuals with high chances of suffering from illness. In contrast, *ex-post prioritarianism* is concerned with the expected distribution of *realized* lifetime well-being. From an ex-post point of view, the worse-off individuals are those who have experienced some hardship during their life, expressed in a low lifetime well-being. For example, individuals who die prematurely, individuals who have low lifetime income, and disabled individuals who suffer from poor health and lack access to many lifetime opportunities.

Compared with utilitarianism, the concern for the worse-off has two consequences. First, prioritarianism is even less likely than utilitarianism to support situations where low-income groups bear the brunt of COVID-19 (net) control policy costs. Equivalently, a prioritarian approach to control policy evaluation is likely to support more redistributive transfers favoring unemployed and essential workers because they bear the greatest economic and health burdens, respectively, of the COVID-19 pandemic and thus are among the worse-off in the population. Second, while utilitarianism is indifferent between saving the life of a well-off individual or the life of a worse-off individual, provided they enjoy the same increase in well-being if saved, prioritarianism prefers the latter. Because young individuals are among the worse-off in lifetime terms (because they have not yet had the opportunity to live a full life), prioritarian SWFs attach higher value than utilitarianism to interventions that benefit mostly the young. One consequence of this is that prioritarianism is more likely than utilitarianism to prioritize the young in the allocation of vaccine doses, despite the fact that they might have lower overall benefits from being vaccinated.

Note that all three methods eventually support the introduction of control policies over an uncontrolled spread of the virus because of the death toll imposed by the pandemic and the positive economic feedback from controlling the virus (e.g., the poor can only go back to work when the pandemic is under control). However, the three methods will differ in the specific characteristics of the recommended control policy (e.g., when to start the lockdown and when to start reopening, whether to have age-specific policies, how to allocate tests and healthcare resources, and so on). In particular, if the trade-off between saving lives and saving livelihoods exists, BCA, utilitarianism, and prioritarianism are, in that order, increasingly sensitive to distribution and so could differ in the choice between more or less aggressive control. If the net costs of control are regressively distributed, for example, aggressive control may be more optimal under BCA than under utilitarianism or prioritarianism. In contrast, if the trade-off between health and the economy does not exist, aggressive control may be optimal under all three views, and distributional sensitivity may largely manifest in the amount of income support to provide the worse-off during the control period.

4. Example 1: The value of COVID-19 control policies in the United States

This section considers a simple example to illustrate the differences among BCA, utilitarianism, and prioritarianism when it comes to evaluating COVID-19 control policies. The example applies to the U.S. economy. Here, we discuss the main features of the example, while Appendix A describes the model in full. The example is not intended to provide policy recommendations, but to clarify the implications of different evaluation methods.

Suppose that, at the beginning of the pandemic, the policy maker must choose between a no intervention strategy (i.e., letting the virus spread in an uncontrolled way) and a strict control strategy that prevents all deaths and is maintained until a vaccine or an effective treatment is found. Further suppose that the vaccine and/or treatment is expected to be available in one year. What is the maximum socially acceptable GDP loss from the intervention to consider the control strategy better than no intervention at all? If the actual cost of the policy is lower than the maximum socially acceptable GDP loss, then the policy maker will implement the policy. In contrast, if the actual cost of the policy is larger than the maximum acceptable GDP loss, then no intervention is better than the control strategy. What is regarded as “socially acceptable” depends on the chosen

evaluation method. Obviously, no real-world policy maker faces a stark choice between strict control and doing nothing. But the lessons from this simple example generalize to more realistic settings in which the optimal degree of strictness of the control can be set in a manner more or less sensitive to the distributional implications of such control.

In sum, this example shows that, all else equal: i) the more regressive are the burdens of COVID-19 (i.e., the more concentrated these burdens are on the worse-off), the higher is the value of control to distribution-sensitive value frameworks (i.e., as regressivity increases, the value of control increases more for prioritarianism than it does for utilitarianism, and more for utilitarianism than it does for BCA); ii) the more regressive are the economic costs of control (net of any redistribution of such costs), the lower is the value of control to distribution-sensitive value frameworks; iii) thus, whether the value of control rises with the distribution-sensitivity of the value framework depends on the relative magnitude of these two regressive patterns.

4.1. Setup

We assume that the population is divided into five income quintiles and two age groups, the “young” (<65 years old) and the “old” (65+). The young constitute 84% of the population. Individuals remain in the same quintile all their life. We use the U.S. lifetable to construct the survival function of the two age groups in a non-COVID-19 situation (see Appendix A for more details). Ample evidence indicates that life expectancies are lower for individuals in low-income groups. To account for the socioeconomic gradient of health, we assume that the survival curve derived from the U.S. lifetable applies to individuals with median income, and we scale up and down the survival chances of the other income quintiles.

We introduce heterogeneity across ages and income quintiles in COVID-19 mortality. To simplify, we assume that, absent any intervention, all age and income groups have the same probability of being infected and of spreading the infection.¹⁷ However, conditional on being infected, older people have a higher risk of being hospitalized and of dying. Based on the estimates of infection fatality risk by age in Verity et al. (2020) and on the 2019 U.S. population structure (U.S. Census Bureau 2020), the average infection fatality rate is 0.95%. However, the infection fatality rate of

¹⁷ In reality, members of low-income groups tend to face higher probabilities of infection because they are often essential workers or otherwise more likely to be exposed to the virus. To simplify, we abstract from this.

the young is 0.28%, while the infection fatality rate of the old is 4.36%. Assuming that, absent any intervention, 70% of the population will be infected, 0.67% of the population will die of COVID-19 on average. From this it follows that the COVID-19 mortality rate of the young is 0.2%, while the mortality rate of the old is 3.05% (i.e., 0.2% of the young and 3.05% of the old die of COVID-19).

Additionally, we assume that the fatality rate may be a decreasing function of income to capture the socioeconomic gradient of health (Seligman et al. 2021). Low socioeconomic groups often face constraints in accessing healthcare (e.g., because they have access to lower-quality healthcare, or because they wait to see a doctor even if ill to avoid losing their job), thereby exposing them to a higher fatality risk. The presence of comorbidities such as hypertension and diabetes has also been found to be correlated with the risk of death from COVID-19, and low socioeconomic groups disproportionately suffer from these conditions. In the benchmark scenario, we assume that low-income quintiles bear a disproportionate number of deaths, and we compare it with the case in which all income quintiles suffer the same number of deaths (i.e., independently of income, 0.2% of the young die of COVID-19 vs. 3.05% of the old). Table 1 summarizes the COVID-19 mortality rates by age and income group if low-income quintiles bear a disproportionate number of deaths.

Table 1. COVID-19 mortality rates by age and income group.

	1 st quintile	2 nd quintile	3 rd quintile	4 th quintile	5 th quintile
Young	0.32%	0.23%	0.18%	0.15%	0.10%
Old	4.96%	3.51%	2.87%	2.31%	1.62%

Assumptions: 70% of the population is infected, the infection fatality rate for the young is 0.28%, the infection fatality rate for the old is 4.36%. COVID-19 deaths occur disproportionately in low income groups, with an income elasticity of the mortality rate equal to -0.5 (see Appendix A). The 1st income quintile is the poorest, and the 5th is the richest. The numbers are rounded to the second significant digit.

We introduce heterogeneities across age and income quintile in the economic costs of control. COVID-19 control policies reduce economic activity and, as a result, induce a GDP contraction. The costs of this contraction are heterogeneously distributed across the population. Lower

socioeconomic groups are more likely to bear the brunt of the costs, but adopting relief programs (such as extending unemployment benefits) can substantially redistribute the costs away from low-income groups. As noted earlier, the pandemic itself can induce an economic contraction because of spontaneous reductions in economic activity brought about by fear of infection (e.g., working from home, avoiding crowded places, etc.), job and income loss, and asset depreciation, but in the example we ignore that. In other words, we assume a trade-off between saving lives and saving livelihoods, i.e., avoiding COVID-19 deaths induces a GDP contraction relative to the no intervention scenario. Stated differently, we assume that GDP with an uncontrolled COVID-19 pandemic and no intervention is the same as pre-pandemic GDP. We consider two cases: one in which the monetary costs of the policy are borne by individuals in proportion to their income, and one in which they are disproportionately borne by lower income individuals. Additionally, young individuals may be more exposed to the negative economic effects of control policies, and so we account for the age distribution of policy costs as well.

We assume that individual well-being depends on two attributes: income and longevity. COVID-19 reduces expected longevity (and possibly income). COVID-19 control policies shrink the longevity reduction at the expense of a larger income contraction. To simplify, we assume that other individual attributes (e.g., physical and mental health status, the ability to interact socially, human capital, etc.) do not affect well-being. We also assume that the income contraction brought about by the control policy has only short-term effects (i.e., once the pandemic is over, the economy goes back to pre-COVID-19 levels), and we do not consider increased mortality due, for instance, to overburdened healthcare systems. The short-term effects assumption implies that the policy-induced GDP contraction lasts only one year, i.e., the time until the vaccine or treatment is available.

4.2. WTP calculations

Before determining the maximum GDP loss triggered by the intervention that the policy maker would find socially acceptable (based on the chosen evaluation method) such that the intervention is better than no intervention, let us compute the maximum income loss that each individual would find acceptable for the same intervention. This is defined as the maximum percentage of current *own* income the individual is willing to sacrifice to eliminate the COVID-19 mortality risk she

faces. In other words, the individual is indifferent between, on the one hand, suffering the COVID-19 mortality risk and, on the other hand, losing income equal to the willingness to pay while suffering no COVID-19 mortality risk (the individual is indifferent in the sense that her well-being is constant in the two situations). This willingness to pay is the backbone of BCA. Because the intervention we are considering eliminates all COVID-19-related fatalities, such a willingness to pay can also be interpreted as the individual burden of the pandemic, absent any intervention (and no economic loss triggered by the pandemic itself).

Table 2 summarizes individuals' willingness to pay for the intervention assuming that low-income quintiles bear a disproportionate number of deaths, i.e., adopting the mortality rates displayed in Table 1. We derive this willingness to pay through a modeling exercise that posits a reasonable functional form for individual expected well-being, and that relies on the distribution of income and longevity across the population and on existing information about individuals' risk attitudes and individuals' incentives to trade-off income and longevity. Appendix A more formally describes our computations.

Because of the higher COVID-19 fatality risk, the old age group attaches a much higher value to policies that eliminate the risk of COVID-19 than the young age group. Moreover, even though the low-income quintiles face a larger fatality risk, they are willing to sacrifice much less than the top quintiles to eliminate the risk of COVID-19. This is a direct result of the lower ability to pay of individuals at the bottom of the income distribution. For example, although an old individual in the 1st income quintile has three times the mortality risk of an old individual in the 5th income quintile, the difference in income implies that the former is willing to pay only a third of her income to avoid the risk of a COVID-19 death, while the latter is willing to pay more than two thirds of her (much higher) income to avoid a relatively lower risk of death.¹⁸ Table A.1 in Appendix A shows individuals' willingness to pay when the distribution of fatalities is independent of income. The same pattern of income-related differences applies, although high-income (low-income) individuals have higher (lower) willingness to pay than the values in Table 2 due to the higher (lower) risk they face.

¹⁸ In absolute terms, and assuming that the 1st quintile income is \$16,250, while the 5th quintile income is \$152,750, the old individual in the 1st quintile is willing to pay \$6,200 and the old individual in the 5th income quintile \$118,978.

Table 2. Individuals' willingness to pay as a percentage of own income for a policy that eliminates the risk of death from COVID-19, by age and income quintile.

	1 st quintile	2 nd quintile	3 rd quintile	4 th quintile	5 th quintile
Young	9.1%	17.7%	23%	28.6%	38.3%
Old	38.2%	56.4%	63.7%	69.9%	77.9%

Assumptions: COVID-19 fatalities occur disproportionately in low-income groups (COVID-19 mortality rates by age and income quintile are taken from Table 1); there is no income loss associated to an uncontrolled pandemic.

4.3. Comparison of BCA, utilitarianism, and prioritarianism: increasing sensitivity to the distribution of the costs of control

BCA evaluates an intervention by summing the amounts individuals are willing to sacrifice to implement the intervention. Using the values of Table 2 and the share of each age-income group in the population, BCA implies that the maximum socially acceptable cost of an intervention that avoids all COVID-19 deaths equals 36.8% of annual aggregate GDP. Therefore, the policy maker adopting a BCA approach is ready to sacrifice up to 36.8% of GDP to eliminate the pandemic, independently of who bears the brunt of those policy costs. In this case, if the policy costs are proportional to income, all individuals pay 36.8% of their own income to finance the intervention. But if the policy costs are regressive (i.e., the low-income quintiles pay a disproportionate amount of them), the individuals in the lowest income quintile would end up sacrificing much more than 36.8% of their income, while the top-income quintile would sacrifice much less than that.¹⁹ However, BCA would find the two situations equally valuable, and would recommend any control policy to eliminate the pandemic as long as, on average, it costs less than 36.8% of GDP. Such a result occurs because BCA is insensitive to the distribution of costs, thereby overemphasizing the interests of the wealthy individuals.

¹⁹ For example, if the income elasticity of policy costs is equal to 0.5 (distribution of policy costs is mildly regressive), the 1st income quintile pays 11% of the total costs, the 2nd 15%, the 3rd 19%, the 4th 23%, and the 5th 33% (see Appendix A). As a percentage of their own income, individuals in the 1st income group have to sacrifice 79% of their own income, while the other income groups sacrifice, respectively, 56%, 45%, 37%, and 26% of their own income.

Table 3 shows the maximum percentage reduction in aggregate GDP that utilitarian and prioritarian policy makers find acceptable to finance an intervention that eliminates the COVID-19 mortality risk. In other words, the maximum acceptable GDP loss is defined such that the policy maker finds a situation where no COVID-19 deaths occur but individuals bear that GDP loss socially equivalent to a situation where individuals die of COVID-19 but bear no policy-related income loss (i.e., the two situations yield the same social welfare). Because the intervention under consideration eliminates all COVID-19 deaths, the maximum acceptable GDP loss can be interpreted as the societal burden of the pandemic, absent any intervention. We compute the societal burden of the pandemic for the three SWFs discussed: utilitarianism, ex-ante prioritarianism, and ex-post prioritarianism. We assume an Atkinson prioritarian SWF of the form $g(w) = \frac{w^{1-\gamma}}{1-\gamma}$, where the parameter γ represents the degree of priority to the worse-off in (expected or realized) lifetime well-being w .²⁰ We consider two different degrees of priority $\gamma = \{1, 2\}$.²¹

To determine the sensitivity of the utilitarian and prioritarian societal burden of the pandemic to the distribution of policy costs and COVID-19 fatalities, we consider five scenarios: (1) The distribution of COVID-19 deaths is regressive (i.e., a disproportionate number of COVID-19 fatalities occur among low-income groups), the distribution of policy costs is also regressive (i.e., low-income groups bear the brunt of the policy costs), and only the young population suffers an income loss due to the intervention. (2) The distribution of COVID-19 deaths is regressive, the distribution of policy costs is proportional to income, and only the young population suffers an income loss due to the intervention. (3) The distribution of COVID-19 deaths is regressive, the distribution of policy costs is progressive (i.e., high-income groups bear the brunt of the policy

²⁰ See Adler (forthcoming) for a description of Atkinson SWFs. Note that if the priority parameter $\gamma = 1$, then the Atkinson SWF is of the form: $g(w) = \ln w$.

²¹ Suppose that individual utility has the form $u(c) = \frac{c^{1-\lambda}}{1-\lambda} - k$, with $\lambda = 2$ the coefficient of relative risk aversion, and $k = -0.000125288$ the utility at subsistence consumption level (see Appendix A for the calibration details). Take two individuals, one in the 1st income quintile (income=\$16,250) and one in the 5th income quintile (income=\$152,750). The former has about 10% the income of the latter. Based on the example (in particular, coefficient of relative risk aversion equal to 2), a utilitarian decision maker (for which the priority parameter $\gamma = 0$) would be indifferent between donating \$100 to the low-income person and \$1.13 to the high-income person, with $\frac{u'(c_5)}{u'(c_1)} = \left(\frac{16,250}{152,750}\right)^2 = \1.13 , and $u'(c) = c^{-2}$ the marginal utility of consumption. A prioritarian decision maker with priority parameter equal to 1 is indifferent between donating \$100 to the low-income person and \$0.6 to the high-income person, with $\frac{g'(c_5)}{g'(c_1)} = \left(\frac{u_1}{u_5}\right)^\gamma \left(\frac{16,250}{152,750}\right)^2 = \0.6 . A prioritarian decision maker with priority parameter equal to 2 is indifferent between donating \$100 to the low-income person and \$0.3 to the high-income person.

costs), and only the young population suffers an income loss due to the intervention. (4) The distribution of COVID-19 deaths is regressive, the distribution of policy costs is also regressive, and all age groups suffer an income loss due to the intervention. (5) The distribution of COVID-19 deaths is independent of income (i.e., COVID-19 mortality risk depends only on age), the distribution of policy costs is regressive, and only the young population suffers an income loss due to the intervention.

We have already commented on the fact that the value determined through BCA is independent of the distribution of policy costs, based on the principle that “a dollar is a dollar is a dollar.” Note that if the COVID-19 mortality risk is independent of income (Scenario 5), the maximum socially acceptable cost is larger, at 44.7% of GDP, due to the relatively higher burden of deaths suffered by high-income quintiles in this scenario (see Table A.1 in Appendix A). Because BCA values the mortality risk reductions of wealthy individuals more than those of low-income ones, the value of an intervention that saves wealthy lives is larger than the value of an intervention that saves the same number of lives, but mainly poor lives. This is because willingness to pay depends on ability to pay (thus wealthy people value a 1 percentage point reduction in mortality risk considerably more than low-income ones), and BCA takes the unweighted sum of individuals’ willingness to pay, i.e., without adjusting for differences in abilities to pay.

A striking result from Table 3 is that, for most scenarios, BCA condones more expensive policies to eliminate the risk of COVID-19 than does utilitarianism or prioritarianism. Only when the distribution of policy costs is progressive (Scenario 3), i.e., when high-income groups pay proportionately more of the policy costs than low-income groups, do both utilitarianism and prioritarianism attach higher values to COVID-19 risk suppression policies than BCA. For example, if the intervention to suppress COVID-19 triggers a 30% reduction in yearly GDP, BCA would find such an intervention better than no intervention at all; utilitarianism and prioritarianism, in contrast, would recommend against implementing such a control strategy, unless the distribution of policy costs is highly progressive.

Table 3. Maximum percentage GDP loss that is considered socially acceptable to pay for an intervention that eliminates the COVID-19 mortality risk.

Scenario	BCA (%)	Utilitarianism (%)	Ex-ante prioritarianism (%)		Ex-post prioritarianism (%)	
			$\gamma = 1$	$\gamma = 2$	$\gamma = 1$	$\gamma = 2$
<u>Scenario 1:</u> Regressive distribution of deaths Regressive distribution of costs Only the young pay the costs	36.8	15.3	12.7	10.6	13.1	11.1
<u>Scenario 2:</u> Regressive distribution of deaths Distribution of costs proportional to income Only the young pay the costs	36.8	26.4	23.2	20.3	24	21.3
<u>Scenario 3:</u> Regressive distribution of deaths Progressive distribution of costs Only the young pay the costs	36.8	40.9	39.7	38.2	41	40
<u>Scenario 4:</u> Regressive distribution of deaths Regressive distribution of costs All age groups pay the costs	36.8	16.2	13.5	11.2	14	11.9
<u>Scenario 5:</u> Distribution of deaths independent of income Regressive distribution of costs Only the young pay the costs	44.7	16.1	12.7	9.9	13.2	10.5

Assumptions: there is no income loss associated with an uncontrolled pandemic; scenarios 1 through 4 adopt the age and income specific mortality rates due to COVID-19 from Table 1, while scenario 5 adopts the mortality rates from Table A.1.

Unlike BCA, utilitarianism and prioritarianism are sensitive to the distribution of policy costs. Moving from Scenario 1 to Scenario 3, as the distribution of costs becomes more progressive (i.e., as the economic burden of the control policy shifts from low- to high-income groups), the maximum socially acceptable GDP loss associated with the intervention increases. In other words, the more likely the utilitarian or prioritarian policy maker will rank the control policy over no intervention. Consequently, if income-support policies are in place (e.g., unemployment benefits) to redistribute policy costs from low- to high-income quintiles, the societal value of eliminating COVID-19 risk increases.

The sensitivity of utilitarianism and prioritarianism to the distribution of policy costs reflects a shared assumption about the decreasing value of money, and in the prioritarian case an extra concern for the worse off. The decreasing value of money assumption implies that dollars paid by low-income quintiles decrease their well-being and total utilitarian welfare more than an equal number of dollars paid mainly by high-income quintiles. Therefore, for a given total reduction in COVID-19 deaths, progressive distributions of costs lead to a lower decrease in welfare. In the prioritarian case, the sensitivity to policy costs reflects the additional prioritarian concern for preserving the (expected or realized) well-being of the worse-off. Because low-income individuals are among the worse-off, less regressive distributions of costs are even more valuable under prioritarianism than under utilitarianism. For instance, the maximum utilitarian socially acceptable GDP loss is 2.7 times larger with a progressive rather than regressive distribution of policy costs (40.9% in Scenario 3 vs. 15.3% in Scenario 1). In the ex-ante and ex-post prioritarian cases the ratio between the two scenarios is even higher: The maximum socially acceptable prioritarian GDP loss is more than three times larger with a progressive rather than regressive loss, and the higher the priority parameter, the larger the ratio.

As a consequence, the more regressive the costs of controlling COVID-19, the more the distribution sensitivity of the value framework lowers the value attached to suppressing the pandemic. Increasing the regressivity of the distribution of policy costs reduces the maximum socially acceptable GDP loss, and such a reduction is larger under prioritarianism than under utilitarianism, and is in turn larger under utilitarianism than under BCA (where there is effectively no reduction at all). Analogous results would hold if the pandemic itself caused financial losses. The more regressive the financial burdens of COVID-19, the larger is the increment in the socially acceptable GDP loss from suppression for both utilitarianism and prioritarianism. And this

increment is larger under prioritarianism than it is under utilitarianism, which in turn is larger than the increment under BCA, this latter equaling zero given BCA's distribution insensitivity. What ultimately matters is the distribution of policy costs net of the COVID-19 costs: the more regressive the net burden, the more the distribution sensitivity of the value framework lowers the value attached to suppression.

Consider again Scenario 1, where the distribution of costs is regressive. Compared with utilitarianism, ex-ante prioritarianism places a lower value on policies aimed at eliminating the risk of COVID-19 (i.e., the ex-ante prioritarian maximum acceptable GDP loss is lower than the utilitarian one). Additionally, the higher the priority conferred to the worse-off in expected well-being (i.e., the higher the parameter γ), the lower the maximum reduction in GDP that an ex-ante prioritarian policy maker would find acceptable to fund an intervention that completely suppresses the pandemic and its death toll. In Scenario 1, older individuals are most at risk of COVID-19 death, while only young and mainly less wealthy individuals bear the costs of the policy. From an ex-ante point of view, the poor and the young are among the worse-off in the population: the poor because they have less access to resources and opportunities and have lower living standards overall, and the young because they have lived for fewer years and have not yet had a chance to live a full life. Thus, ex-ante prioritarianism is less likely to support eliminating COVID-19 risk because doing so hurts the prospects of those who will certainly pay the costs of the policy but who are unlikely to die from COVID-19. Note that the difference between ex-ante prioritarianism and utilitarianism declines once the distribution of costs becomes more progressive (moving from Scenario 1 to Scenario 2 to Scenario 3) because of the reduced burden on low-income quintiles and the associated reduced need to protect the worse-off.

The ex-post prioritarian maximum acceptable GDP loss for eliminating COVID-19 also tends to be smaller than the utilitarian one, and decreases as the degree of priority to the worse-off increases. Moreover, in the example the ex-post prioritarian acceptable loss is larger than the ex-ante prioritarian one. Undoubtedly, from an ex-post point of view, those who die poor and prematurely are the worst-off in the population, especially if they die young. Young, low-income individuals are also the worst-off ex-ante because they face a larger lifetime mortality risk and have lower expected well-being. But they are the worst-off only in expectation, i.e., the situation might get rosier. That is likely one reason why ex-post prioritarianism values COVID-19 elimination more than ex-ante prioritarianism does. Everyone bears the costs of the policy (both

those who survive and those who die), and many individuals face a risk of dying of COVID-19 ex-ante, but only a few will be dead because of COVID-19 ex-post. Protecting their interests is thus a priority.

Note that as the degree of priority increases, the maximum acceptable GDP loss associated with the policy reduces. The reason is that, from an ex-post point of view, the worse-off includes not only those who die prematurely of COVID-19, but also those who die prematurely of other causes (not necessarily related to COVID-19 or indirectly triggered by COVID-19). In other words, the ex-post prioritarian policy maker must trade off the interests of those who die of COVID-19 (who are better off with the policy), those who die prematurely from other causes but still bear the costs of a COVID-19 policy (who are worse off with the policy), and those who do not die but pay the costs (who are worse off with the policy). Because the second group represents the majority among those who die prematurely (after all, they are more likely to die of other causes than of COVID-19), protecting their interests means not investing in the control policy. The larger the priority parameter, the larger the concern for protecting their livelihood, because preventing their death is not possible. However, the ex-post prioritarian concern for protecting the interests of those who die of COVID-19 is sufficient for accepting larger policy-induced income losses than utilitarianism does.

BCA highly values policies that reduce the fatality risk of high-income groups, as shown by a comparison of Scenario 5 with Scenario 1. In contrast, the distribution of fatalities by income group has only a marginal effect on the evaluation under the utilitarian and prioritarian approaches. Utilitarianism attaches a slightly higher value to preventing COVID-19 if fatalities were independent of income rather than disproportionately falling on low-income quintiles. This is because well-being increases with income in the model, that is, high-income groups have larger expected well-being than low-income groups. Consequently, for a utilitarian policy maker saving a wealthy life is more valuable than saving a less well-off life because differences in income are associated with differences in quality of life.

Ex-ante prioritarianism finds COVID-19 more burdensome if fatalities disproportionately affect low-income quintiles because that increases well-being inequality among individuals.²²

²² With $\gamma = 1$, the two values are pretty close: 12.697% when the distribution of deaths is regressive, 12.685% when the distribution of deaths is independent of income.

Individuals in low-income quintiles are the worst-off in expectation because they have lower income and lower life expectancy; the risk of COVID-19 further erodes their lifetime well-being, particularly if they face a higher fatality rate than the top income quintiles. Ex-post prioritarianism also values interventions against COVID-19 more if the fatalities occur disproportionately in the low-income quintiles as long as the priority parameter is sufficiently high.²³ Dying young and poor is the worst-off situation from an ex-post point of view, thereby increasing the value of policies aimed at preventing that.

Note that distributing the policy costs among all age groups rather than only the young (Scenario 4) does not seem to have a sizable impact on the value of the control policy relative to Scenario 1. If the costs are divided among all age groups, the societal burden of COVID-19 increases slightly under both utilitarianism and prioritarianism as compared with Scenario 1. A more equal distribution of costs across the age groups explains the increase in the maximum socially acceptable GDP loss.

This example shows that: i) the more regressive are the burdens of COVID-19, the more the distribution sensitivity of the value framework raises the value of control²⁴; ii) regardless of the SWF approach, the value of controlling COVID-19 decreases if doing so hurts the poor; prioritarianism reinforces this result with respect to utilitarianism; iii) the net impact of distribution sensitivity on the value of control depends on the relative strength of the two patterns of regressivity.

Note that, no matter the evaluation method, the simplified model analyzed here suggests that the societal burden of uncontrolled COVID-19 is likely to be substantial; “do nothing” is unlikely to be the best policy option compared with a control policy, i.e., the actual costs of the control policy are likely to be lower than the maximum acceptable economic losses computed in Table 3.

The example also illustrated that ex-ante prioritarianism and ex-post prioritarianism evaluate policies to eliminate the risk of COVID-19 quite differently, with the latter attaching more value

²³ When the priority parameter is close to zero, the ex-post prioritarian socially acceptable GDP loss is smaller with a regressive distribution of deaths than with a distribution independent of income. This resembles the utilitarian case.

²⁴ The more regressive are the health burdens of COVID-19, the more prioritarianism values control policies, while that does not necessarily hold for utilitarianism and BCA. However, if we add also non-health burdens of COVID-19, then, by parallel with the effects of regressive policy costs, we get that the more regressive are the total burdens of COVID-19, the more the value of control increases going from BCA, to utilitarianism, to prioritarianism.

to such policies than the former. The difference in perspective explains the last result. Even if some individuals face a substantial risk of infection and death from COVID-19, that is just one of many possible outcomes—after all, they might survive COVID-19. Ex-post, some individuals have died of COVID-19 and were definitely victims of the pandemic. Their fate is worse than the fate of someone who might die because of the pandemic, but who has not yet died. Thus, the ex-post prioritarianism societal burden of COVID-19 is higher than the ex-ante prioritarian one. Adler (forthcoming) discusses the pros and cons of the two approaches. One reason why the ex-post perspective may be more appropriate for evaluating COVID-19 policies is that, at the population level, there is no risk. Some people will definitely die of COVID-19, we just do not know their identity. In contrast, at the individual level death from COVID-19 occurs only with some (small) probability. Focusing on the relatively small individual risk discounts the negative fate of those who die due to COVID-19.

The model we have considered so far has several limitations. Addressing them changes the societal burden of COVID-19 and possibly the comparison between BCA, utilitarianism and the various forms of prioritarianism. For instance, we have neglected the morbidity effects of COVID-19, for example, pain and suffering when ill, or the long-term health effects of the disease; adding them would make COVID-19 even more costly. Even absent governmental interventions, people might change their behavior for fear of infection, for example, deciding to work from home if possible or avoiding crowded places like restaurants and bars. Accounting for this endogenous social distancing reduces the value of policies to control the spread of COVID-19 (see, e.g., Farboodi et al. 2020, Toxvaerd 2020).

Moreover, the example assumes that all financial impacts caused by the control policy are borne in the current period (i.e., while the pandemic ravages and we wait for a vaccine or better treatment). This is a strong assumption, because governments often increase their debt in the face of a crisis, thereby shifting some of the burden to future periods or future generations. The impact of such smoothing on the societal burden of COVID-19 and on the differences across evaluation methods is likely to depend on the regressivity of the resulting distribution of costs across individuals and across time.

Additionally, because we do not account for behavioral changes in the absence of intervention, we underestimate the economic losses of the pandemic. But the “do nothing” strategy might end up

being far more costly than eliminating the COVID-19 risk, thereby leading to no effective trade-off between saving the economy and saving lives. In other words, the actual costs of the control policy may be lower than the economic loss triggered by the pandemic itself. If the pandemic causes deaths *and* threatens individuals' livelihoods, the maximum reduction in pre-pandemic GDP triggered by the control intervention that the policy maker would find socially acceptable will be even larger than the values computed so far. This result holds independently of the evaluation method and is because the pandemic is bad both for the economy and for population health. Because the threshold for acceptability has increased, the control policy is more likely to be the optimal strategy.

Policy evaluations based on BCA require a methodology to determine individuals' willingness to pay. The standard methodology is to employ value-of-statistical-life (VSL) measures to monetize the value of avoiding a death (Hammitt 2000; Kniesner and Viscusi 2019). VSL represents the monetary equivalent of saving one (unidentified) life among a group of identical people and derives from the rate at which individuals are willing to substitute a small change in their income (or wealth) for a small change in their survival probability. This rate is estimated either through revealed preference data from observing individuals' choices that affect both income and mortality risk (e.g., purchases of protective equipment or decisions to select a safer job) or through stated preference methods such as surveying people about their hypothetical choices regarding income and mortality risk. For example, if individuals are willing to pay \$10 to reduce their risk of premature death by 1 in a million, then the monetary value of saving one statistical life is \$10,000,000. In a group of 10 million identical people, one death is avoided for sure, and if everyone is willing to pay \$10, the total value of saving one life is exactly \$10 million.

There is an extensive literature providing VSL estimates for the population of interest (Viscusi 2018). In the example, we used those estimates to infer individuals' preferences between income and longevity, and to calibrate the functional form of individual well-being. We then calculate individuals' willingness to pay through a modeling exercise.

Although the calibration relies on VSL estimates from the literature, we did not directly adopt those estimates to measure individuals' WTP for two reasons. First, although VSL is not constant across the population, reflecting the preferences and circumstances of different individuals (e.g., their life expectancy or their wealth), in practice it is often assumed to be uniform across the

population. This avoids the ethically objectionable result that saving the life of a wealthy individual is more valuable than saving the life of a less well-off individual because of differences in abilities to pay. However, assuming a constant VSL also implies that life expectancy differences do not matter, i.e., saving the life of a young or of an old person is equally valuable, even though the young will live more years if saved. To account for differences in life expectancy, value of statistical life years (VSLY) measures are often employed (Hammit 2007). VSLY is computed by dividing the constant VSL by the average life expectancy in the population. The value of extending the life of an individual is then equal to the product of VSLY and the number of life years gained if the person is saved, thereby accounting for age and life expectancy differences.

The second issue with VSL (and VSLY) is the marginality assumption. VSL is the *marginal* rate of substitution between wealth and survival probability, but COVID-19 may represent a *nonmarginal* risk of death. Standard theory implies that the individual's willingness to pay for a nonmarginal increase in survival probability is an increasing, concave function of the change in survival probability (Hammit 2020). In other words, the higher the increment in survival probability, the larger the individual's willingness to pay, but as the increment grows larger, the rate at which the individual is willing to trade wealth for incremental increases in survival probability decreases.²⁵ For example, Adler (2020) shows that, for fatality rates of 1% of the population (a nonmarginal risk and close to the estimated mortality rate of COVID-19), the individual's willingness to pay to eliminate such a risk is only half the amount associated with VSL.²⁶ Therefore, for large probabilities of preventing death, the VSL is likely to overestimate the value that individuals place on saving lives.

²⁵ For example, the individual might be willing to pay \$1,000 for increasing her chances of survival by 0.1 percentage points, but she might be ready to pay at most \$1,800 for increasing her chances of survival by 0.2 percentage points. The concavity of individual willingness to pay with respect to the size of the increment in survival probability is due to the constraints on the individual's ability to spend and the fact that the opportunity cost of spending increases the more the individual has already reduced her mortality risk.

²⁶ Adler (2020) assumes a VSL of more than \$9 million, which, in a population of 100 people, corresponds to an individual willingness to pay to marginally reduce the fatality risk of about \$90,000. In contrast, the individual willingness to pay to eliminate the 1% fatality risk (a nonmarginal reduction) is only \$46,000, half the value implied by VSL. Thus, VSL tend to overestimate the value individuals place on reducing the risk. See also Cutler and Summers (2020), Greenstone and Nigam (2020), and Thunström et al. (2020) for analyses using the VSL framework that suggest that policies to reduce fatalities from COVID-19 are worth trillions of dollars.

5. Example 2: Vaccine allocation

Manufacturing and delivery constraints make vaccine against SARS-CoV-2 a scarce and valuable resource, at least initially. The questions in this initial phase include how to distribute the scarce vaccine doses and who should get them first. Different approaches have been proposed to allocate vaccines among the population. For instance, they could be allocated randomly through a lottery, based on the principle that everyone should be treated equally—although not necessarily fairly, because individuals may have different needs. Or a vaccine could be allocated based on a first-come, first-served principle, thereby potentially favoring those with easier access to healthcare resources or with better information. Or they could be allocated based on need by giving them first to those who are at higher risk, e.g., the elderly in the case of COVID-19 or essential workers. Alternatively, they could be allocated based on the size of expected health benefits, thereby favoring not only high-risk individuals, but also those who will have a long and healthy life if saved (thus, according to this view, a younger person may receive the vaccine before an older person, even if the latter is at greater risk).²⁷ Or they could be allocated based on the size of both health and non-health benefits, e.g., prioritizing those who cannot earn labor income unless they prove they have been vaccinated. Or they could be allocated based on the instrumental value created by the immunized, e.g., prioritizing essential workers within the healthcare sector and service industries given that keeping them healthy is instrumental in saving other people's lives or keeping essential services open. Another approach is to allocate resources based on ability to pay, thus favoring the wealthy. Or, in contrast, priority could be given to the worse-off to compensate for pre-existing injustices.

Adapting the example of section 4, we comment now on the best approach to allocate vaccines through the lenses of utilitarian and prioritarian SWFs and BCA. Although the example is better suited to illustrate vaccine allocation within a country, similar considerations can also guide vaccine allocation among countries.

Individuals are divided into five income quintiles and two age groups, the “young” (<65) and the “old” (65+). Individuals' well-being depends on income and longevity; the risk of COVID-19 threatens both. Suppose that a limited number of vaccine doses is administered. Who should

²⁷ We say “may” but not “will.” If fatality risk is much higher for the older adults, and the vaccine is effective at reducing their risk, expected health gains are maximized by prioritizing older adults even though they have lower healthy life expectancy.

receive the first dose of vaccine? Should we allocate it to a young person or an old one, a person in the bottom of the income distribution or someone in the high-income quintiles? Appendix B presents the model more formally.

Vaccines can provide both direct protection to the vaccinated individual and indirect protection to others who are at reduced risk of infection from the vaccinated individual.²⁸ The example focuses on the direct social value of protecting the vaccinated person from COVID-19, leaving aside the indirect social value created by reduced community transmission (e.g., reduced risk of death for unvaccinated persons, increased social and economic activity by nonvaccinated members of society, increased nonmarket productive activities such as childcare or volunteering). Our goal is to determine who in the population has the largest direct social value of being vaccinated. This individual will be the first to get the vaccine.

The general conclusions are that BCA allocates vaccines based on individuals' WTP for vaccination, while the utilitarian and prioritarian allocations are based, respectively, on well-being gains and priority-adjusted well-being gains. In our example, the groups that tend to have larger WTP for vaccination are the old (because of higher fatality risk from COVID-19) and the wealthy (because of higher ability to pay). Larger well-being gains from vaccination are, instead, experienced by the old (because of higher fatality risk from COVID-19) and by the less wealthy (because COVID-19 disproportionately affects low-income individuals), although not the poorest (because in the model quality of life, and thus well-being, increases in income). Because the groups that tend to be worse-off are the young and the poor, the prioritarian allocation will account for the size of the well-being gains but skew the prioritization strategy toward the young and poor. Depending on the relative strength of empirical facts and priority concerns, young poor populations might be vaccinated before some of the older adults, especially older wealthy adults.

Interestingly, a prioritarian approach seems to substantiate the argument that essential workers should be among the first to be vaccinated for social justice reasons. Many essential workers belong to disadvantaged racial and ethnic minorities and socioeconomically vulnerable groups. Therefore, a prioritarian approach to vaccine allocation suggests that essential workers should be

²⁸ Many vaccines block the symptoms of a disease, but do not prevent infection or onward transmission of the pathogen (Warfel et al. 2014; Hodgson et al. 2020). As we write this chapter (December 2020), whether the COVID-19 candidate vaccines will be effective at reducing the risk of transmission is unclear (Peiris and Leung 2020).

prioritized not only because they face a high risk of exposure to COVID-19 and of transmission of the virus, but also because of their social and economic vulnerability.

5.1. Epidemiological considerations in vaccine allocation

Given our simplified model focused only on longevity and income impacts, the direct social value of vaccinating an individual in a given group depends on (i) the clinical effectiveness and safety of the vaccine for the individual; (ii) the individual's pre-vaccination probability of suffering from COVID-19; (iii) the individual's infection fatality rate; (iv) the potential non-health benefits of being vaccinated, e.g., the income gains if proof of vaccination is required to work; and (v) the social value of a given reduction in mortality risk and a given increase in income for the individual. This last value depends on the ethical framework adopted for the analysis, i.e., BCA, utilitarianism, or prioritarianism (or any other approach not considered in this chapter).²⁹

Individuals' infection fatality rates (iii) for COVID-19 vary by age and socioeconomic status, with older individuals and individuals in low socioeconomic groups more likely to suffer severe health consequences, e.g., due to the higher incidence of comorbidities, more frequent use of public transit, and reduced access to high-quality healthcare (O'Driscoll et al. 2021; McLaren 2020). This heterogeneity in fatality risks may call for vaccinating older people first, especially older people in low-SES groups. However, the safety and effectiveness of a vaccine (i) usually vary across individuals. Vaccines (e.g., flu vaccines) are frequently less effective in older adults because the immune system weakens with age (Bridle and Sharif 2020).³⁰ Additionally, early clinical trials rarely include older populations to test the safety and efficacy of vaccines. Pressure to release a vaccine as soon as possible implies that the first vaccines may not be proven effective and safe for older people, thereby raising the question of how to balance effectiveness and mortality risk. If effectiveness reduces with age, vaccinating the younger population first may be better, even if they face lower risks of dying from COVID-19.

²⁹ An often debated issue is whether vaccine allocation strategies should depend only on number of deaths or also on life years gained, and potentially on other metrics (Emanuel and Wertheimer 2006). These questions pertain to the ethical framework that is adopted in the analysis, and the extent to which the ethical framework values those outcomes.

³⁰ As of December 2020, there is some evidence suggesting that leading COVID-19 vaccines work well also among the elderly (Mandavilli 2020).

Moreover, infection fatality rates may depend on the existence of an economically viable and effective treatment and on the effectiveness of this treatment by age, comorbid status, etc. Therefore, what matters for vaccine allocation is not just the individuals' infection fatality risk, but rather the infection fatality risk net of the potential protection conferred by the treatment. For example, if a treatment is forthcoming soon that is more effective for older people than the vaccine, vaccinating the young generations and letting the older generation wait for the treatment may be better (even if that wait will cause some deaths among the elderly).

The probability of being infected with COVID-19 (ii) depends on individuals' characteristics; for example, older people may be more susceptible to infection (i.e., conditional on a contact with an infected person, they are more likely to catch the virus), and individuals with many social contacts are more likely to be infected (Ferguson et al. 2020; Viner et al. 2021). However, this probability depends also on the social and economic context wherein the vaccine is introduced, e.g., the number and characteristics of the people still susceptible (assuming that the once infected have acquired permanent or temporary immunity), the presence of nonpharmaceutical interventions (those interventions will likely remain in place in the initial phases if the number of vaccine doses is limited), and individuals' behavior and compliance with social-distancing norms. For example, if we are in a hypothetical society where populations at higher fatality risk can be safely and perfectly isolated from the rest of the population, and where such an isolation is not burdensome for them and their families, vaccinating first the individuals with high contacts, but low fatality risk, rather the individuals with low contacts, but high fatality risk, may be better. Likewise, the non-health benefits of being vaccinated (iv) depend on the economic burden of the counterfactual scenario and how this burden is distributed across the population. For instance, if months of economic lockdown have strained young generations economically and mentally, vaccinating them first rather than the elderly may be more valuable if employability requires proof of vaccination, even if the latter are more at risk.

Another factor to consider is the quantity of vaccine supplied and the timing of its availability. If individuals are not vaccinated today, will they be vaccinated before a new surge of cases and deaths? For example, if the next stock of vaccine doses will be distributed only in six months, the argument for prioritizing the high-fatality-risk individuals gains strength, because those individuals are likely to die while waiting for the next stock.

The discussion and the example focus on the direct value of vaccination. However, note that if the vaccine could stop the transmission of the virus, vaccinating the super-spreaders first (i.e., those individuals with many contacts), rather than the high-fatality-risk individuals, might be better (Medlock and Galvani 2009).

To abstract from many of these considerations, while keeping the main trade-offs, we will make the following assumptions. A limited stock of vaccine doses is distributed at the beginning of the pandemic; no more doses will be distributed in the foreseeable future (or at least in the relevant time horizon of the analysis). Vaccination reduces mortality risk, but has no impact on income; except for the vaccine, no other intervention is implemented, in particular, there is no social or physical distancing policy and no treatment is discovered. Vaccination reduces the mortality rate of the vaccinated, but has no impact on the mortality rate of the unvaccinated (i.e., we neglect the disease transmission channels). Therefore, we assume that the unvaccinated individuals face the same mortality risk as in the benchmark uncontrolled outbreak scenario analyzed in Section 4. Table 1 summarized the probability of dying from COVID-19 for each age and income group, under the assumption that deaths over-proportionately occur in low-income groups.

5.2. Prioritization groups under BCA, utilitarianism and prioritarianism

Benefit-cost analysis. Consider first the allocation rule that a benefit-cost approach would support. BCA evaluates policies (i.e., ways of allocating vaccine doses) by estimating individuals' willingness to pay and summing up those monetary amounts. A safe and 100% effective vaccine eliminates the risk of COVID-19-related death for the individual. Table 2, which measured individuals' willingness to pay to eliminate the risk of COVID-19, is then informative of the monetary value that individuals place on being vaccinated. The table shows that individuals' willingness to pay is larger for the old than for the young because the former is expected to benefit more from being vaccinated (as they face a much higher fatality risk). Additionally, the monetary values are larger for top income quintiles than for bottom income quintiles because the former have a greater ability to pay. Therefore, the allocation rule supported by a benefit-cost approach is sensitive to the size of the expected benefits and to differences in abilities to pay, and thus, the vaccine dose will be given to an old individual in the top income quintile. We have already

commented on the ethically objectionable discrimination of BCA toward the interests of the wealthy.

Utilitarian SWF. Consider now the utilitarian framework. The utilitarian SWF evaluates policies based on the change in the sum of individuals' expected well-being compared with the status quo. The optimal allocation rule of vaccine doses is that which confers the greatest increase in the sum of individuals' well-being. In other words, the first vaccine dose will be given to the individual who is expected to gain most in terms of increase in well-being. The first consequence is that the vaccine is likely to be given to individuals who will experience a larger reduction in their COVID-19-related fatality risk thanks to vaccination, i.e., individuals with high baseline COVID-19 fatality risk and/or individuals with high vaccine effectiveness. The second consequence is that the size of the mortality risk reduction is insufficient to determine who gets the vaccine first, because changes in well-being are more than just changes in health.

To determine which demographic group gains the most well-being from COVID-19 vaccination, let us first introduce the concept of social value of mortality risk reduction (SVRR).³¹ The $SVRR_{ia}$ is defined as the change in social welfare induced by a small reduction (e.g., a 1 in 100,000 probability reduction) in the mortality risk of an individual in age group a and income group i . If $SVRR_{ia} \geq SVRR_{jb}$, then reducing the mortality risk of an individual in age group a and income group i generates more social welfare than reducing the mortality risk of an individual in age group b and income group j by the exact same amount. The SVRR is the moral weight associated with a given reduction in fatality risk.

The *utilitarian* $SVRR_{ia}^U$ is the change in *utilitarian* social welfare (i.e., the change in individual lifetime expected well-being) induced by a small reduction in the mortality risk of an individual in age group a and income group i . Suppose that vaccination reduces the COVID-19 fatality risk of an individual in age group a and income group i by x_{ia} , while reducing the fatality risk of an individual in age group b and income group j by x_{jb} . The utilitarian vaccine allocation rule prioritizes the individual in age group a and income group i over the individual in age group b and income group j if and only if $SVRR_{ia}^U x_{ia} \geq SVRR_{jb}^U x_{jb}$. Suppose that $x_{ia} > x_{jb}$, i.e., the fatality risk reduction an individual in age group a and income group i experiences thanks to the COVID-

³¹ See also Hammitt and Treich (forthcoming) and Adler et al. (2021) for a definition and some properties of social value of mortality risk reduction.

19 vaccine is larger than the fatality risk reduction an individual in age group b and income group j experiences. If reducing the fatality risk of the former is considered more valuable from a utilitarian point of view than reducing the fatality risk of the latter (i.e., $SVRR_{ia}^U \geq SVRR_{jb}^U$), the individual experiencing the larger fatality risk reduction obtains the vaccine first. However, if reducing the fatality risk of an individual in age group a and income group i is considered less valuable from a utilitarian point of view than reducing the fatality risk of an individual in age group b and income group j (i.e., $SVRR_{jb}^U > SVRR_{ia}^U$), then the vaccine is allocated to the former first only if the associated COVID-19 fatality risk reduction is large enough, $x_{ia} \geq \frac{SVRR_{jb}^U}{SVRR_{ia}^U} x_{jb}$.

Table 1 indicates that older individuals in the lowest income quintile face the largest COVID-19 mortality risk. Additionally, for each income quintile, the COVID-19 mortality rate of the older population is about 15 times larger than the corresponding mortality rate of the younger population. Unless the age difference in vaccine effectiveness is very large (i.e., the vaccine is at least 15 times more effective for the young than for the old), the older people in low-income groups are likely to benefit the most from a vaccine with any average vaccine effectiveness in terms of mortality risk reduction (i.e., they have the larger x). However, they will be allocated the vaccine first only if the SVRR attached to a given reduction in their mortality risk is large enough.

Table 4 shows the ratio of the utilitarian SVRR for each age and income group to the corresponding utilitarian SVRR of an old individual in the 1st income quintile. For each age group a and income group i , we divided $SVRR_{ia}^U$ by the corresponding SVRR attached to mortality risk reductions an old individual in the 1st income quintile enjoys ($SVRR_{1o}^U$). Values greater than 1 in age group a and income group i mean that, from a utilitarian perspective, reducing the mortality risk of an individual in age group a and income group i by a given amount is more valuable than reducing the mortality risk of an old individual in the 1st income quintile by the same amount.

Table 4. Utilitarian SVRR compared with the corresponding utilitarian SVRR attached to an older person in the 1st income quintile, by age and income group.

	1 st quintile	2 nd quintile	3 rd quintile	4 th quintile	5 th quintile
Young	2.51	3.85	4.36	4.71	5.11
Old	1	1.48	1.64	1.76	1.86

Although older people in the 1st income quintile are expected to experience the largest risk-reduction benefit from the vaccine (Table 1), they have the lowest utilitarian SVRR, while young wealthy individuals have the highest value. For example, a small reduction in mortality risk is valued almost five times more for a young individual in the 4th income quintile than an old individual in the 1st income quintile. More generally, the utilitarian SVRR decreases in age and increases in income. This is because younger individuals have longer remaining life expectancies if saved (i.e., a higher number of life years gained). Also, in the simplified model used in this chapter, income buys happiness, i.e., utility increases in income. Thus, reducing the mortality risk of an individual with better quality of life is more valuable from a utilitarian perspective.

Let $\frac{SVRR_{ia}}{SVRR_{1o}} x_{ia}$ be the “utilitarian-adjusted” COVID-19 mortality risk reduction that an individual in age group a and income group i experiences if vaccinated, and x_{1o} the vaccine-induced COVID-19 mortality risk reduction experienced by an older individual in the 1st income quintile. Even though older individuals in the 1st income quintile have the lowest utilitarian SVRR, they should still receive the vaccine first if their mortality risk reduction is larger than the “utilitarian-adjusted” mortality risk reduction any other group would experience, i.e., $x_{1o} > \frac{SVRR_{ia}}{SVRR_{1o}} x_{ia}$, for any $ia \neq 1o$.

For example, suppose the vaccine is 100% effective for all age groups. Then, Table 1 implies that such a vaccine reduces the COVID-19 mortality risk of an old individual in the 1st income quintile by 4.96 percentage points. The same vaccine reduces the mortality risk of a young individual in the 4th income quintile by only 0.15 percentage points. However, any given reduction in mortality risk is valued 4.71 times more if it occurs to a young individual in the 4th income quintile rather than an old one in the 1st income quintile. Therefore, the 0.15 percentage point reduction accruing

to the young is as valuable as a 0.7 percentage point reduction (4.71×0.15) accruing to the old. In other words, a vaccine that saves 0.15% of the young people in the 4th income quintile is as valuable, from a utilitarian point of view, as a vaccine that saves 0.7% of the old people in the 1st income quintile (if the two groups have the same number of members).

Table 5 shows the utilitarian-adjusted reductions in COVID-19 mortality risk induced by a vaccine that is 100% effective for all age groups (where the numbers in Table 5 are derived by multiplying the SVRRs in Table 4 by the COVID-19 mortality rates in Table 1). Elderly in the 2nd income quintile have the largest utilitarian-adjusted mortality risk reduction if vaccinated, i.e., vaccinating them creates the largest increase in utilitarian social welfare. Indeed, vaccinating an elderly individual in the 2nd income quintile--which reduces their mortality risk by 3.51 percentage points--raises utilitarian social welfare as much as reducing the mortality risk of an older individual in the 1st income quintile by 5.2 percentage points. Vaccinating elderly in the 2nd income quintile therefore yields more utilitarian welfare than vaccinating those in the 1st income quintile, which would only reduce the latter's mortality risk by 4.96 percentage points.

Table 5. Utilitarian-adjusted reduction in COVID-19 mortality risk if the vaccine is 100% effective for all age groups, by age and income group.

	1 st quintile	2 nd quintile	3 rd quintile	4 th quintile	5 th quintile
Young	0.8%	0.87 %	0.8%	0.7%	0.53%
Old	4.96%	5.2%	4.71%	4.06%	3.01%

Assumption: COVID-19 deaths occur disproportionately in low income groups, with COVID-19 mortality rates displayed in Table 1. If the vaccine is less than 100% effective, the values in the table should be multiplied by the reduced effectiveness (e.g., if effectiveness is 60%, the utilitarian-adjusted reduction for a young individual in the 1st income quintile is $0.6 \times 0.8\%$).

Therefore, the utilitarian allocation rule will prioritize elderly in low-income quintiles, though not necessarily the poorest ones. Differences in length and quality of life by income explain this result: individuals in the 2nd quintile face a relatively high COVID-19 mortality rate and, if saved from death can expect a longer life and higher income than those in the 1st quintile. More generally, our

assumptions imply that for any age-invariant vaccine effectiveness, the mortality risks faced by poor elderly are large enough to ensure that the utilitarian allocation rule will prioritize them (although not the poorest ones), despite the lower relative social value attached to extending their life.

If vaccine effectiveness decreases with age, the likelihood of prioritizing the young individual increases. Earlier, we pointed out that vaccinating the young in our example could produce larger mortality risk reductions than vaccinating the elderly if vaccine effectiveness in the young were at least 15 times that for the old. However, the lower utilitarian value of mortality risk reductions in the elderly implies that smaller age differences in vaccine efficacy are required to prioritize the young. Indeed, vaccine efficacy in the young that is six times higher than that in the old suffices to prioritize the young (we divide the utilitarian-adjusted risk reduction of an old person in the 2nd income quintile by the utilitarian-adjusted risk reduction of a young person in the 2nd income quintile³²). For example, if the vaccine is 50% effective for the young and only 8% effective for the old, the young get priority.

Ex-ante prioritarian SWF. Ex-ante prioritarianism evaluates policies based on changes in the sum of strictly increasing and strictly concave transformations of individuals' expected lifetime well-being compared with the status quo. The concave transformation guarantees that, for a given increase in individuals' well-being, the worst-off in expected lifetime well-being are prioritized. The ex-ante prioritarian allocation rule then depends on the change in an individual's expected well-being if vaccinated (i.e., the rule should consider who gains more from the vaccine in terms of well-being, as in the utilitarian rule) and on the identity of the vaccinated person, with a preference for the most disadvantaged in expected well-being. Individuals with low expected well-being are typically people in low socioeconomic groups, people with short life expectancies (i.e., the sick and disabled), and the young because they have lived for fewer years and face a higher lifetime risk of dying prematurely than the individuals who have already reached an old age.

To determine the ex-ante prioritarian allocation strategy, define the ex-ante prioritarian social value of mortality risk reduction $SVRR_{ia}^{EA}$ as the change in ex-ante prioritarian social welfare induced by a small mortality risk reduction of an individual in age group a and income group i . If

³² We chose the 2nd income quintile because it represents the largest value of vaccination, from Table 5.

$SVRR_{ia}^{EA} \geq SVRR_{jb}^{EA}$, then reducing the mortality risk of an individual in age group a and income group i is more valuable from an ex-ante prioritarian point of view than reducing the mortality risk of an individual in age group b and income group j by the same amount. Let x_{ia} and x_{jb} be the COVID-19 mortality risk reductions induced by vaccination; individuals in age group a and income group i get priority over individuals in age group b and income group j if and only if $SVRR_{ia}^{EA} x_{ia} \geq SVRR_{jb}^{EA} x_{jb}$, or equivalently, $x_{ia} \geq \frac{SVRR_{jb}^{EA}}{SVRR_{ia}^{EA}} x_{jb}$.

Table 6 shows the ratio of the ex-ante prioritarian SVRR for each age and income group to the corresponding ex-ante prioritarian SVRR of elderly in the 1st income quintile. It is the ex-ante prioritarian counterpart of Table 4 for the utilitarian case. To compute the values in Table 6, we assume an Atkinson SWF of the form $g(w) = \frac{w^{1-\gamma}}{1-\gamma}$, with priority parameter γ equal to 2. Table 6 shows that given our assumptions, from an ex-ante prioritarian perspective, a 1 percentage point reduction in the mortality risk of a young individual in the 1st income quintile is valued 2.87 times more than a 1 percentage point reduction in the mortality risk of an old individual in the 1st income quintile.

Table 6. Ex-ante prioritarian SVRR compared with the corresponding ex-ante prioritarian SVRR attached to an older person belonging to the 1st income quintile, by age and income group.

	1 st quintile	2 nd quintile	3 rd quintile	4 th quintile	5 th quintile
Young	2.87	1.91	1.69	1.58	1.47
Old	1	0.66	0.58	0.54	0.5

Assumption: Atkinson SWF with priority parameter equal to 2.

As in the utilitarian case, reducing the mortality risk of young people is valued more than reducing the mortality risk of older people by the same amount. This priority to the young is even stronger than in the utilitarian case: the young are prioritized because they are expected to live longer if

saved (as in the utilitarian case) and because they are among the worse-off in lifetime terms (they have not yet lived a full life). For instance, if we take the first income quintile, the ex-ante prioritarian SVRR of a young individual is 2.87 times the value of an older one, while the utilitarian value is 2.51 times the value of an older one: reducing the mortality risk of the young rather than the old is more valuable under ex-ante prioritarianism than utilitarianism. A similar comparison holds for the other income quintiles.³³

Unlike in the utilitarian case, our assumptions imply that the ex-ante prioritarian SVRR decreases in income. For a given age, a 1 percentage point reduction in the mortality risk of a low-income individual is more valuable than the equivalent reduction in that of a wealthier one. For example, the ex-ante prioritarian SVRR of an old individual in the 5th quintile is only half that of an old individual in the 1st quintile. This is because low-income individuals are among the worse-off, both because of their poorer economic conditions and because of their shorter life expectancy due to the income gradient in health and the larger COVID-19 mortality risk if not vaccinated. Note that this result of SVRR declining in income is specific to our model assumptions and not generalizable. Table B.1 in Appendix B reports the ex-ante prioritarian SVRR when the priority parameter is equal to 1 (i.e., giving less priority to the worse-off than when the parameter equals 2). While reducing the mortality risk of the young remains more valuable than reducing that of the old, reducing the risk of a rich young person is now more valuable than reducing that of a poor one. Although low-income individuals are among the worse-off, they also expect to gain less from having their life saved because of their shorter life expectancies and lower incomes.

Consider again Table 6. Poor young individuals have the largest ex-ante prioritarian social value of mortality risk reduction. Will they receive the vaccine first? As before, we can compute the “ex-ante-prioritarian-adjusted” mortality risk reductions to determine which age and income group has the largest ex-ante prioritarian value of vaccination. The ex-ante prioritarian-adjusted mortality risk reduction of an individual in age group a and income group i is defined as $\frac{SVRR_{ia}^{EA}}{SVRR_{10}^{EA}} \chi_{ia}$. Table 7 displays the results for a vaccine that is 100% effective for all age groups.

³³ For each income quintile, we can divide the value for the young by the value for the old to determine how much more valuable it is to prioritize the young. In the utilitarian case, we find 2.51 (1st), 2.59 (2nd), 2.65 (3rd), 2.68 (4th), and 2.74 (5th). In the ex-ante prioritarian case, we find 2.87 (1st), 2.91 (2nd), 2.92 (3rd), 2.93 (4th), and 2.96 (5th). For each income quintile, the value is higher in the ex-ante prioritarian case than in the utilitarian one.

Table 7. Ex-ante-prioritarian-adjusted reduction in COVID-19 mortality risk if the vaccine is 100% effective for all age groups, by age and income group.

	1 st quintile	2 nd quintile	3 rd quintile	4 th quintile	5 th quintile
young	0.92%	0.43%	0.31%	0.23%	0.15%
old	4.96%	2.3%	1.66%	1.24%	0.8%

Assumption: Atkinson SWF with priority parameter equal to 2. If the vaccine is less than 100% effective, the values in the table should be multiplied by the reduced effectiveness (e.g., if effectiveness is 60%, the ex-ante-prioritarian-adjusted reduction for a young individual in the 1st income quintile is $0.6 \times 0.92\%$).

Table 7 indicates that if the vaccine is 100% effective for everyone, it will be allocated first to elderly in the 1st quintile because of the high expected reduction in mortality risk despite the low ex-ante prioritarian SVRR of a given reduction in such risk (see Table 6). For example, the vaccine reduces the mortality risk of the old in the 1st quintile by 4.96 percentage points, while the ex-ante-prioritarian-adjusted reduction in the mortality risk of the young in the 1st quintile is only 0.92 percentage points. In other words, vaccinating a young individual in the 1st quintile is equivalent, from an ex-ante prioritarian perspective, to reducing the mortality risk of an old individual in the 1st quintile by 0.92 percentage points, a lesser reduction than that from vaccinating an old person in the 1st quintile (who would experience a 4.96 percentage point reduction in mortality risk). These numbers account for the lower ex-ante prioritarian social value associated with reductions in mortality risk experienced by the young compared with those experienced by old.

These results differ from those of the utilitarian case in two respects. First, the ex-ante prioritarian rule prioritizes individuals in the 1st income quintile over people in higher income quintiles, which was not always the case with a utilitarian SWF. Second, the rule prioritizes poor young individuals over rich older ones. A strong priority for the worse-off implies that young low-income individuals get priority over older wealthy people because of their young age and despite the lower expected death rate from COVID-19. For example, vaccinating a young person in the 1st income quintile is more valuable than vaccinating an old one in the 5th income quintile: vaccinating the former is equivalent to achieving a 0.92 percentage point reduction in mortality risk of an old individual in

the 1st income quintile, while vaccinating the latter is equivalent to achieving only a 0.8 percentage point reduction. Note that a policy maker with less concerns for the worse-off (priority parameter equal to 1 instead of 2) will also prioritize the old poor individuals in allocating the vaccine, but not the poor young over the wealthy old (Table B.2 in Appendix B).

If the vaccine effectiveness reduces with age, will a young low-income individual get the vaccine before an old low-income one? The answer is positive if the effectiveness for the young is at least 5.4 times the effectiveness for the old (where the figure has been computed by dividing the ex-ante-prioritarian-adjusted mortality risk reduction of the old in the 1st income quintile from Table 7 by the corresponding value of a young person). Therefore, young individuals get priority despite the lower gain in mortality if the vaccine is sufficiently more effective for them than it is for old people. Note that the “excess” effectiveness for prioritizing the young is lower than in the utilitarian case. Once again, that is because, compared with utilitarianism, ex-ante prioritarianism attaches higher moral importance to preserving the life of a young person than the life of an older one.

Ex-post prioritarian SWF. Ex-post prioritarianism evaluates policies based on the expected sum of concave transformations of individuals’ realized well-being. The concave transformation implies that, for a given increase in realized well-being, priority is given to those with lower realized lifetime well-being. From an ex-post point of view, the worse-off are those who die prematurely, the individuals in low-income groups, and the young (because dying prematurely while young is considered worse than dying prematurely while old).

Let $SVRR_{ia}^{EP}$ represent the ex-post prioritarian social value of reducing the mortality risk of an individual in age group a and income group i . Such a value is defined as the change in ex-post prioritarian social welfare induced by a small mortality risk reduction of an individual in age group a and income group i . If $SVRR_{ia}^{EP} \geq SVRR_{jb}^{EP}$, then reducing the mortality risk of an individual in age group a and income group i is considered more valuable from an ex-post prioritarian point of view than reducing the mortality risk of an individual in age group b and income group j by the exact same amount. Moreover, let x_{ia} and x_{jb} be the COVID-19 mortality risk reductions induced by vaccination in these two individuals. The ex-post prioritarian policy maker allocates the vaccine

to the individual in age group a and income group i before the individual in age group b and income group j if and only if $SVRR_{ia}^{EP} x_{ia} \geq SVRR_{jb}^{EP} x_{jb}$, or equivalently, $x_{ia} \geq \frac{SVRR_{jb}^{EP}}{SVRR_{ia}^{EP}} x_{jb}$.

Table 8 shows the ratio of the ex-post prioritarian SVRR for each age and income group to the corresponding ex-post prioritarian SVRR for an old individual in the 1st quintile, assuming an Atkinson SWF with priority parameter equal to 2. For example, from an ex-post prioritarian perspective, a 1 percentage point reduction in the mortality risk of a young individual in the 1st income quintile is valued 5.58 times more than a 1 percentage point reduction in the mortality risk of an old individual in the 1st income quintile.

Table 8. Ex-post prioritarian SVRR compared with the corresponding ex-post prioritarian SVRR attached to an older person belonging to the 1st income quintile, by age and income group.

	1 st quintile	2 nd quintile	3 rd quintile	4 th quintile	5 th quintile
young	5.58	3.89	3.59	3.39	3.27
old	1	0.67	0.61	0.57	0.54

Assumption: Atkinson SWF with priority parameter equal to 2.

As in the ex-ante prioritarian case, the value of reducing mortality risk by a fixed amount decreases in income. Additionally, as in the utilitarian and ex-ante prioritarian cases, reducing the mortality risk of a young individual is more valuable than reducing the mortality risk of an older individual by the same amount. Note that the priority conferred to the young is considerably higher than in both the utilitarian and ex-ante prioritarian cases. Consider, for example, the 1st income quintile. The ex-post prioritarian SWF values risk reductions accruing to the young 5.58 times more than comparable risk reductions accruing to the old; the corresponding values in the utilitarian and ex-ante prioritarian cases are, respectively, 2.51 and 2.87. Comparable differences hold for the other income quintiles.

All three SWFs prioritize the young because they are expected to experience a bigger gain in lifetime well-being if saved from COVID-19 (i.e., a larger number of life years gained). On top of that, ex-ante and ex-post prioritarianism prioritize the young because they rank among the worse-off in well-being terms since they have lived for fewer years. However, this “fair innings” argument³⁴ is stronger under ex-post prioritarianism than under ex-ante prioritarianism. This difference depends on the fact that the ex-post prioritarian favors the young because they are definitely the worst-off if they die prematurely; the ex-ante prioritarian favors them because they face a higher risk, that is, they are the worst-off in expectation. Table 8 assumes a relatively high degree of priority to the worse-off. Table B.3 in Appendix B shows that, with a lower degree of priority, differences in income do not matter any longer, i.e., for a given age, a 1 percentage point reduction in mortality risk has approximately a constant value independently of whether it occurs to a poor or a rich individual. In contrast, differences in age still matter considerably, with higher values attached to interventions that benefit primarily the young.

Multiplying the COVID-19 mortality risk each demographic group faces (Table 1) and the ex-post prioritarian SVRRs of Table 8, we can determine to whom the ex-post prioritarian policy maker will allocate a 100% uniformly effective vaccine first. Table 9 presents the “ex-post-prioritarian-adjusted” mortality risk reductions achieved by such a vaccine, where the ex-post prioritarian-adjusted mortality risk reduction of an individual in age group a and income group i is defined as $\frac{SVRR_{ia}^{EP}}{SVRR_{10}^{EP}} x_{ia}$. The age and income group with the largest ex-post-prioritarian-adjusted mortality risk reduction receives the vaccine first.

³⁴ The “fair innings” argument refers to the idea that mortality risk reduction policies benefiting younger people should be counted more because the young have not yet had a fair innings in their life, i.e., they have not had a chance to live for a normal lifespan (Hammitt and Treich forthcoming; Adler et al. 2021).

Table 9. Ex-post-prioritarian-adjusted reduction in COVID-19 mortality risk if the vaccine is 100% effective for all age groups, by age and income group.

	1 st quintile	2 nd quintile	3 rd quintile	4 th quintile	5 th quintile
young	1.78%	0.88%	0.66%	0.5%	0.34%
old	4.96%	2.37%	1.74%	1.32%	0.87%

Assumption: Atkinson SWF with priority parameter equal to 2. If the vaccine is less than 100% effective, the values in the table should be multiplied by the reduced effectiveness (e.g., if effectiveness is 60%, the ex-post-prioritarian-adjusted reduction for a young individual in the 1st income quintile is $0.6 \times 1.78\%$).

Once again, an old individual in the 1st income quintile will receive the vaccine first because the achievable reduction in COVID-19 mortality risk outweighs the lower ex-post prioritarian value attached to interventions that benefit the old. For example, vaccinating an old person in the 1st income quintile reduces her probability of death by 4.96 percentage points. Vaccinating a young person in the 1st income quintile is equivalent, from an ex-post prioritarian perspective, to reducing the probability of death of an old person in the 1st income quintile by only 1.78 percentage points. Clearly, vaccinating the old produces more value.

Additionally, as with the ex-ante prioritarian SWF, the priority conferred to the young is so high that a young, low-income individual will get the vaccine before an old, high-income one. Specifically, vaccinating a young individual in the 1st income quintile produces more value than vaccinating an old individual in the 3rd (or higher) income quintile. This result hinges on the relatively high priority parameter, and it disappears if the policy maker has a low concern for the worse-off (Table B.4 in Appendix B).

If vaccine effectiveness decreases with age, under which conditions will it be given first to a young individual? Fix, for instance, the 1st income quintile. If the vaccine is at least 2.8 times more effective for a young than for an old person, the ex-post prioritarian rule will allocate it to a young person first, regardless of the lower risk faced by the young (where the 2.8 figure has been computed by dividing the ex-post-prioritarian-adjusted mortality risk reduction of the old in the 1st income quintile from Table 9 by the corresponding value of a young person). For example, if the vaccine is 50% effective for the young, and only 20% effective for the old, the higher ex-post

prioritarian value conferred to saving a young life implies that a young person in the 1st income quintile will get the vaccine first.

One clear conclusion from this discussion is that age matters in allocating vaccine doses and that both ex-ante and ex-post prioritarian SWFs are more likely than utilitarianism to allocate the vaccine to young people first (or at least to some poor young people over wealthier older people), despite their lower COVID-19 mortality risk. Note that this “ageist” feature of prioritarian SWFs is likely to be softened if we recognize that current young generations have higher standards of living than previous young generations (i.e., those who are currently old). Even if the young have not yet had a chance to live a full life, the quality of their life (both the one already lived and the expected one) might be, all things considered, higher than that of the current older generations. Accounting for that is likely to reduce the higher priority attached to interventions that benefit primarily the young.

The discussion so far has completely neglected the costs of different allocation rules. Logistical and political constraints may make it infeasible or expensive to distribute the vaccine in some areas, even if the populations in those areas would benefit most from it. Feasibility constraints are also linked to the properties of the vaccine. For example, widely distributing vaccines that need to be stored at very cold temperatures is more difficult. Vaccine hesitancy also interferes with the optimal vaccine allocation mechanism. If complementary interventions are required to persuade people to get the vaccine, the optimal vaccine allocation may be tilted toward nonhesitant individuals independent of the expected benefits. Additionally, we have so far assumed that only one vaccine will be allocated. Competition among multiple vaccines with different properties may result in a vaccine being allocated first to individuals who will benefit less from it.

6. Discussion and Conclusions

This chapter discusses the use of utilitarian and prioritarian SWFs in assessing COVID-19 control and prevention policies, as compared with more standard BCA. We show that SWF analysis is a rich and flexible evaluation framework; unlike BCA, it is sensitive to the distribution of policy costs and benefits among the population, and it does not have the ethically objectionable result that the interests of the rich count more than the interests of the poor in the evaluation process. One key implication is that all else equal, if the net costs of control are regressively distributed,

aggressive control may be more optimal under BCA than under utilitarianism or prioritarianism. (Though if the burdens of COVID-19 are regressive, then all else equal, both utilitarianism and prioritarianism will recommend more stringent control policies than BCA). BCA prioritizes vaccinating the wealthy based on their higher ability to pay. In contrast, SWF analysis (whether prioritarian or utilitarian), value policies not in terms of people's willingness to pay for their effects but rather in terms of those policies' effects on well-being, thus doing away with any priority the rich may get from their greater ability to pay.

The main difference between utilitarianism and prioritarianism is that the former is only concerned with maximizing the aggregate well-being gains from policy while the latter will accept lower aggregate gains for a better distribution of such gains toward the worse off. One example where such a difference plays out is in the allocation of vaccine doses. Young individuals are among the worse-off in the population because they have not yet had the opportunity to live a full life. So both forms of prioritarianism prioritize vaccinating low-income young individuals before wealthy older adults, even though the latter face a higher COVID-19 mortality risk and could enjoy larger utility benefits from being saved. In the utilitarian case, all older adults should receive priority over younger populations independently of their relative income. However, note that all three SWFs agree in ranking low-income older individuals at the top of the vaccine prioritization strategy.

Both BCA and SWF analysis are likely to support policies that control the pandemic over an uncontrolled spread, although the optimal degree of stringency may vary. Both context and policy horizon matter. In a country where the trade-off between saving lives and saving livelihoods is stark (e.g., if the control policy causes widespread unemployment, depletion of savings, fall into poverty, and associated increase in mortality) and the distribution of policy costs is regressive, BCA might support more aggressive control policies than utilitarianism or prioritarianism if the burdens of COVID-19 are not too regressive. If instead there is no trade-off between health and the economy (e.g., if the economic losses from an uncontrolled pandemic are staggering) and the only way to save the economy is through suppression, then all three methods will support suppression and mainly differ on the magnitude of the income support to the worse-off. The policy horizon plays a role as well. Excessive attention to short-run economic costs of suppression leads to less aggressive policies, independently of the chosen method.

One important topic concerning COVID-19 that we have not discussed is the role of uncertainty. Great uncertainty surrounds the transmission mechanisms of SARS-CoV-2, the severity of its effects, and the timing and effectiveness of treatments and vaccines. Countries must make decisions about the best control strategy based on uncertain and evolving scientific evidence. Prudential considerations suggest employing aggressive control strategies to avoid worst-case scenarios with unknown probability and/or magnitude, regardless of the high policy costs.³⁵

Another issue the COVID-19 pandemic raises is that of aversion to catastrophic outcomes. From a population-level perspective, a catastrophic outcome refers to a situation in which many people die or face a loss together. Thus, not only the eventual magnitude, but also the coincidence or bunching together of deaths and losses might affect policy evaluation. Many COVID-19 deaths occur among old people who would likely have died from other causes soon. Yet, instead of having those deaths spread over a few months or years, they occurred all at the same time. Images of rooms and trucks filled with coffins have been frequently broadcasted during the pandemic. Aversion to such a dreadful situation calls for stringent control policies even if they impose high costs on the economy.³⁶

The chapter focuses on COVID-19 policies. However, the evaluation tools described can be applied to other epidemics, and more generally, to the evaluation of health technologies and health interventions. One area in which SWF analysis can be useful is in determining the value of vaccines. Traditional economic evaluations of vaccines adopt a narrowly defined health-centric perspective, focusing on two types of benefits: (i) the direct health benefits for the immunized (e.g., the reduction in the number of fatalities or the number of quality-adjusted life years gained) and (ii) the savings in the healthcare system due to a reduced need for alternative treatments. This perspective implies that developing, manufacturing, and distributing a new vaccine is valuable

³⁵ However, an argument may exist that learning about the virus occurs only as the pandemic unfolds. The effects of the virus have been found more severe among the elderly because we have observed more deaths in this age group. The presence of endogenous learning implies that, from a purely informational perspective, it is better not to completely suppress the virus. This suggests a trade-off between the value of information and the value of being precautions, although, in practice, the former is likely to be small in the face of an unfolding pandemic.

³⁶ See Rheinberger and Treich (2017) for an exact definition of catastrophe aversion and whether catastrophe aversion is a common sentiment. Hammitt and Treich (forthcoming) also discusses the notion of catastrophe aversion.

only insofar as it avoids deaths and disabilities and saves health sector costs.³⁷ A rapidly growing body of evidence suggests that vaccination yields sizable health, economic, and social benefits far beyond the narrow health-centric benefits traditionally captured in economic evaluations (see, e.g., Bloom et al. 2018; Jit et al. 2015; Nandi et al. 2019). These benefits include, e.g., protection against nosocomial infections (i.e., infections acquired in hospitals or in healthcare-associated locales), and herd protection; increasing labor force participation and labor productivity among the working-age population; protection against catastrophic health spending and the risk of falling into poverty traps; better cognitive functioning, higher school attendance, and higher educational attainment among vaccinated children; increases in tax revenues; and promotion of social equity insofar as the vaccine-preventable disease disproportionately affects low-income groups, minorities, and disadvantaged populations in general.

SWF analysis, with its focus on the overall well-being impact of health technologies and its concern for distributions, can easily accommodate all the relevant dimensions of vaccine evaluation.³⁸ Drawing from the analysis in the chapter, when assessing the value of vaccines through a SWF approach, what counts are not just the distributional aspects of the disease burden and the distributional aspects of disease-control interventions, but also the distributional aspects of financing vaccines. A vaccine paid for by rich countries and distributed free of charge throughout the world contributes more to social welfare than an equally effective vaccine whose costs are shared by all countries, also developing ones.

³⁷ Note, in particular, that traditional economic evaluation of vaccines often does not incorporate the vaccine-induced decrease in transmission of the pathogen and the resulting herd protection.

³⁸ Note that the framework can also encompass the potential adverse effects of vaccination.

References

- Acemoglu, D., Chernozhukov, V., Werning, I., Whinston, M.D. (2020). “Optimal Targeted Lockdowns in a Multi-Group SIR Model.” *NBER Working Paper* n. 27102. Available at <https://www.nber.org/papers/w27102.pdf>
- Adam, D. (2020). “The Limits of R.” *Nature*, 583: 346–348.
<https://www.nature.com/articles/d41586-020-02009-w>
- Adler, M.D. (2019). *Measuring Social Welfare. An Introduction*. Oxford University Press.
- Adler, M.D. (forthcoming). “Chapter 2: Theory of Prioritarianism.” In M.D. Adler and O.F. Norheim, eds., *Prioritarianism in Practice*. Cambridge University Press.
- Adler, M.D. (2020). “What Should We Spend to Save Lives in a Pandemic? A Critique of the Value of Statistical Life.” *Duke Law School Public Law and Legal Theory Series* No. 2020-40. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3636550
- Adler, M.D., Bradley, R., Ferranna, M., Fleurbaey, M., Hammitt, J.K., Voorhoeve, A. (2020). Assessing the Well-Being Impacts of the COVID-19 Pandemic and Three Policy Types: Suppression, Control, and Uncontrolled Spread.” *T20 Policy Brief*. Available at <https://www.g20-insights.org/wp-content/uploads/2020/11/assessing-the-wellbeing-impacts-of-the-covid-19-pandemic-and-three-policy-types-suppression-control-and-uncontrolled-spread-1606214930.pdf>
- Adler, M.D., Decancq K. (forthcoming). “Chapter 3: Measuring Well-Being and Respect for Preferences.” In M.D. Adler and O.F. Norheim, eds., *Prioritarianism in Practice*. Cambridge University Press.
- Adler, M.D., Ferranna, M., Hammitt, J.K., Treich, N. (2021). “Fair Innings? The Utilitarian and Prioritarian Value of Risk Reduction Over a Whole Lifetime.” *Journal of Health Economics*, 75, 102412. <https://doi.org/10.1016/j.jhealeco.2020.102412>.
- Alvarez, F., Argente, D., Lippi, F. (2020). “A Simple Planning Problem for COVID-19 Lockdown.” *NBER Working Paper* n. 26981. Available at <https://www.nber.org/papers/w26981.pdf>
- Andrew, A., Cattan, A., Costa Dias, M., Farquharson, C., Kraftman, L., Krutikova, S., Phimister, A., Sevilla, A. (2020). “Inequalities in Children’s Experiences of Home Learning during the

COVID-19 Lockdown in England.” *Fiscal Studies*, 41(3): 653-683. <https://doi.org/10.1111/1475-5890.12240>

Armitage, R., Nellums, L.B. (2020). “COVID-19 and the Consequences of Isolating the Elderly.” *The Lancet, Public Health*, 5(5): E256. [https://doi.org/10.1016/S2468-2667\(20\)30061-X](https://doi.org/10.1016/S2468-2667(20)30061-X)

Avent, R. (2020). “COVID-19 Leaves a Legacy of Increased Inequality.” *The Economist*, Nov 17, 2020.

Bloom, D.E., Cadarette, D., Ferranna, M. (forthcoming). “The Value of Vaccination in the Age of COVID-19.” *American Journal of Public Health*

Bloom, D.E., Cadarette, D., Ferranna, M., Hyer, R.N., Tortorice, D.L. (2021). “How New Models of Vaccine Development for COVID-19 Have Helped Address an Epic Public Health Crisis.” *Health Affairs*, 40(3): 1–8. <https://doi.org/10.1377/hlthaff.2020.02012>

Bloom, D.E., Cadarette, D., Ferranna, M., Seligman, B. (2020). “A Matter of Life and Death.” *Finance and Development*. Available at <https://www.imf.org/external/pubs/ft/fandd/2020/04/allocating-scarce-medical-resources-during-the-COVID19-pandemic-bloom.htm>

Bloom, D.E., Fan, V.Y., Sevilla, J.P. (2018). “The Broad Socioeconomic Benefits of Vaccination.” *Science Translational Medicine*, 10(441): eaaj2345. <https://doi.org/10.1126/scitranslmed.aaj2345>

Bloom, D.E., Kuhn, M., Prettnner, K. (forthcoming). “Modern Infectious Diseases: Macroeconomic Impacts and Policy Responses.” *Journal of Economic Literature*.

Blundell, R., Costa Dias, M., Joyce, R., Xu, X. (2020). “COVID-19 and Inequalities.” *Fiscal Studies*, 41: 291–319. <https://doi.org/10.1111/1475-5890/12232>.

Bridle, B.W., Sharif, S. (2020). “Why Vaccines Are Less Effective in the Elderly, and What It Means for COVID-19.” *The Conversation*. Available at <https://theconversation.com/why-vaccines-are-less-effective-in-the-elderly-and-what-it-means-for-covid-19-141971>

Brown, C.S., Ravallion, M. (2020). “Inequality and the Coronavirus: Socioeconomic Covariates of Behavioral Responses and Viral Outcomes Across US Counties.” *NBER Working Paper* n. 27549. <https://doi.org/10.3386/w27549>

- Brown, C.S., Ravallion, M., van de Walle, D. (2020). “Can the World’s Poor Protect Themselves from the New Coronavirus?” *NBER Working Paper* n. 27200. Available at <https://www.nber.org/papers/w27200.pdf>
- Bubar, K.M., Kissler, S.M., Lipsitch, M., Cobey, S., Grad, Y.H., Larremore, D.B. (2021). “Model-Informed COVID-19 Vaccine Prioritization Strategies by Age and Serostatus.” *Science*, eabe6959. <https://doi.org/10.1126/science.abe6959>
- Buckner, J.H., Chowell, G., Springborn, M.R. (2020). “Dynamic Prioritization of COVID-19 Vaccines When Social Distancing Is Limited for Essential Workers.” *medRxiv*. <https://doi.org/10.1101/2020.09.22.20199174>
- Cash, R., Patel, V. (2020). “Has COVID-19 Subverted Global Health?” *The Lancet*, 395(10238): 1687–1688. [https://doi.org/10.1016/S0140-6736\(20\)31089-8](https://doi.org/10.1016/S0140-6736(20)31089-8)
- Chetty, R., Stepner, M., Abraham, S., et al. (2016). “The Association between Income and Life Expectancy in the United States, 2001–2014.” *Journal of the American Medical Association*, 315(16): 1750–1766. <https://doi.org/10.1001/jama.2016.4226>
- Chow, N., Fleming-Dutra, K., Gierke, R., Hall, A., Hughes, M., Pilishvili, T., et al. (2020). “Preliminary Estimates of the Prevalence of Selected Underlying Health Conditions among Patients with Coronavirus Disease 2019—United States, February 12–March 28, 2020.” *Morbidity and Mortality Weekly Report, Centers for Disease Control and Prevention* 69(13): 382–386. <https://doi.org/10.15585/mmwr.mm6913e2>
- Cookson, R., Norheim, O.F., Skarda, I. (forthcoming). “Chapter 6: Health Care.” In M.D. Adler and O.F. Norheim, eds., *Prioritarianism in Practice*. Cambridge University Press.
- Cutler, D.M., Summers, L.H. (2020). “The COVID-19 Pandemic and the \$16 Trillion Virus.” *Journal of the American Medical Association*, 324(15): 1495–1496. <https://doi.org/10.1001/jama.2020.19759>
- Dingel, J.I., Neiman, B. (2020). “How Many Jobs Can Be Done at Home?” *NBER Working Paper* n. 26948. Available at <https://www.nber.org/papers/w26948.pdf>
- Eichenbaum, M.S., Rebelo, S., Trabandt, M. (2020). “The Macroeconomics of Epidemics.” *NBER Working Paper* n. 26882. Available at <https://www.nber.org/papers/w26882.pdf>

- Emanuel, E.J., Persad, G., Kern, A., Buchanan, A., Fabre, C., Halliday, D., Heath, J., Herzog, L., Leland, R.J., Lemango, E.T., Luna, F., McCoy, M.S., Norheim, O.F., Ottersen, T., Schaefer, G.O., Tan, K.-C., Wellman, C.H., Wolff, J., Richardson, H.S. (2020a). “An Ethical Framework for Global Vaccine Allocation.” *Science*, 369(6509): 1309–1312. <https://doi.org/10.1126/science.abe2803>
- Emanuel, E.J., Persad, G., Upshur, R., Thome, B., Parker, M., Glickman, A., Zhang, C., Boyle, C., Smith, M., Phillips, J.P. (2020b). “Fair Allocation of Scarce Medical Resources in the Time of COVID-19.” *New England Journal of Medicine*, 382: 2049–2055. <https://doi.org/10.1056/NEJMs2005114>
- Emanuel, E.J., Wertheimer, A. (2006). “Who Should Get Influenza Vaccine When Not All Can?” *Science*, 312(5775): 854-855. <https://doi.org/10.1126/science.1125347>
- Farboodi, M., Jarosch, G., Shimer, R. (2020). “Internal and External Effects of Social Distancing in a Pandemic.” *NBER Working Paper* n. 27059. Available at <https://www.nber.org/papers/w27059.pdf>
- Favero, C.A., Ichino, A., Rustichini, A. (2020). “Restarting the Economy while Saving Lives under COVID-19.” Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3580626
- Ferguson, N.M., Laydon, D., Nedjati-Gilani, G., et al. (2020). “Impact of Non-Pharmaceutical Interventions (NPIs) to Reduce COVID-19 Mortality and Healthcare Demand.” Imperial College, London, Working Paper. <https://doi.org/10.25561/77482>
- Ferranna, M., Cadarette, D., Bloom, D.E. (2021). “COVID-19 Vaccine Allocation: Modeling Health Outcomes and Equity Implications of Alternative Strategies.” Unpublished manuscript.
- Fontanet, A., Cauchemez, S. (2020). “COVID-19 Herd Immunity: Where Are We?” *Nature Reviews Immunology*, 20: 583–584. <https://doi.org/10.1038/s41577-020-00451-5>
- Giordano, G., Blanchini, F., Bruno, R., et al. (2020). “Modelling the COVID-19 Epidemic and Implementation of Population-Wide Interventions in Italy.” *Nature Medicine*, 26: 855–860. <https://doi.org/10.1038/s41591-020-0883-7>
- Glover, A., Heathcote, J., Krueger, D., Ríos-Rull, J.-V. (2020). “Health versus Wealth: On the Distributional Effects of Controlling a Pandemic.” *NBER Working Paper* n. 27046. Available at <https://www.nber.org/papers/w27046.pdf>

- Gollier, C. (2020). “Cost-Benefit Analysis of Age-Specific Deconfinement Strategies.” *COVID Economics*, 24: 1–31.
- Goodnough, A., Hoffman, J. (2020). “The Elderly vs. Essential Workers: Who Should Get the Coronavirus Vaccine First?” *The New York Times*, Dec. 05, 2020. Available at <https://www.nytimes.com/2020/12/05/health/covid-vaccine-first.html>
- Gopinath, G. (2020). “The Great Lockdown: Worst Economic Downturn Since the Great Depression.” Available at <https://blogs.imf.org/2020/04/14/the-great-lockdown-worst-economic-downturn-since-the-great-depression/>
- Greenstone, M., Nigam, V. (2020). “Does Social Distancing Matter?” *COVID Economics*, 7: 1–23. Available at https://bfi.uchicago.edu/wp-content/uploads/BFI_WP_202026.pdf
- Guan, W., Liang, W., Zhao, Y., Liang, H., Chen, Z., Li, Y., et al. (2020). “Comorbidity and Its Impact on 1590 Patients with COVID-19 in China: A Nationwide Analysis.” *European Respiratory Journal*, 55(5): 20000547. <https://doi.org/10.1183/13993003.00547-2020>
- Hafner, M., Yerushalmi, E., Fays, C., Dufresne, E., Van Stolk, C. (2020). “COVID-19 and the Cost of Vaccine Nationalism.” *RAND Europe Report*. <https://doi.org/10.7249/RR769-1>
- Hammitt, J.K. (2000). “Valuing Mortality Risk: Theory and Practice.” *Environmental Science and Technology*, 34(8): 1396-1400. <https://doi.org/10.1021/es990733n>
- Hammitt, J.K. (2007). “Valuing Changes in Mortality Risk: Lives Saved versus Life Years Saved.” *Review of Environmental Economics and Policy*, 1(2): 228-240. <https://doi.org/10.1093/reep/rem015>
- Hammitt, J.K. (2020). “Valuing Mortality Risk in the Time of COVID-19.” *Journal of Risk and Uncertainty*, 11:1-26. <https://doi.org/10.1007/s11166-020-09338-1>
- Hammitt, J.K., Treich, N. (forthcoming). “Chapter 7: Environmental, Health and Safety Policy.” In M.D. Adler and O.F. Norheim, eds., *Prioritarianism in Practice*. Cambridge University Press.
- Hodgson, S.H., Mansatta, K., Mallett, G., Harris, V., Emary, K.R.W., Pollard, A.J. (2020). “What Defines an Efficacious COVID-19 Vaccine? A Review of the Challenges Assessing the Clinical Efficacy of Vaccines against SARS-CoV-2”. *The Lancet, Infectious Diseases*. Published online. [https://doi.org/10.1016/S1473-3099\(20\)30773-8](https://doi.org/10.1016/S1473-3099(20)30773-8)

- IMF (International Monetary Fund). (2020). *World Economic Outlook, October 2020: A Long and Difficult Ascent*. Available at <https://www.imf.org/en/Publications/WEO/Issues/2020/09/30/world-economic-outlook-october-2020>
- Ioannidis, J. (2020). “The Infection Fatality Rate of COVID-19 Inferred from Seroprevalence Data.” *Bulletin of the World Health Organization* (published online October 14, 2020). Available at https://www.who.int/bulletin/online_first/BLT.20.265892.pdf
- Jit, M., Hutubessy, R., Png, M.E., et al. (2015). “The Broader Economic Impact of Vaccination: Reviewing and Appraising the Strength of Evidence.” *BMC Medicine*, 13: 209. <https://doi.org/10.1186/s12916-015-0446-9>
- JCVI (Joint Committee on Vaccination and Immunisation). (2020). “Priority Groups for Coronavirus (COVID-19) Vaccination: Advice from the JCVI. 2 December 2020.” Available at <https://www.gov.uk/government/publications/priority-groups-for-coronavirus-covid-19-vaccination-advice-from-the-jcvi-2-december-2020/priority-groups-for-coronavirus-covid-19-vaccination-advice-from-the-jcvi-2-december-2020>
- Jones, C.J., Philippon, T., Venkateswaran, V. (2020). “Optimal Mitigation Policies in a Pandemic: Social Distancing and Working from Home.” *NBER Working Paper* n. 26984. Available at <https://www.nber.org/papers/w26984.pdf>
- Kissler, S.M., Tedijanto, C., Goldstein, E., Grad, Y.H., Lipsitch, M. (2020). “Projecting the Transmission Dynamics of SARS-CoV-2 through the Postpandemic Period.” *Science*, 368(6493): 860–868. <https://doi.org/10.1126/science.abb5793>
- Kniesner, T.J., Viscusi, W.K. (2019). “The Value of a Statistical Life.” *Oxford Research Encyclopedia of Economics and Finance*. <https://doi.org/10.1093/acrefore/9780190625979.013.138>
- Levin, A.T., Hanage, W.P., Owusu-Boaitey, N., Cochran, K.B., Walsh, A.P., Meyerowitz-Katz, G. (2020). “Assessing the Age Specificity of Infection Fatality Rates for COVID-19: Systematic Review, Meta-Analysis, and Public Policy Implications.” *European Journal of Epidemiology* (published online December 8, 2020). <https://doi.org/10.1007/s10654-020-00698-1>

- Mahoney, N. (2020). “Don’t Fall for the False Trade-offs of COVID-19 Policy.” *Chicago Booth Review*, (published online March 26, 2020). Available at <https://review.chicagobooth.edu/economics/2020/article/don-t-fall-false-trade-offs-covid-19-policy>
- Mandavilli, A. (2020). “The Next Vaccine Challenge: Reassuring Older Americans.” *The New York Times*, December 14, 2020. Available at <https://www.nytimes.com/2020/12/14/health/coronavirus-vaccine-elderly.html>
- McLaren, J. (2020). “Racial Disparity in COVID-19 Deaths: Seeking Economic Roots with Census data.” *NBER Working Paper n. 27407*. <https://doi.org/10.3386/w27407>
- Medlock, J., Galvani, A.P. (2009). “Optimizing Influenza Vaccine Distribution.” *Science*, 325(5948): 1705-1708. <https://doi.org/10.1126/science.1175570>
- Nandi, A., Shet, A., Behrman, J.R., Black, M.M., Bloom, D.E., Laxminarayan, R. (2019). “Anthropometric, Cognitive, and Schooling Benefits of Measles Vaccination: Longitudinal Cohort Analysis in Ethiopia, India, and Vietnam.” *Vaccine*, 37(31): 4336–4343. <https://doi.org/10.1016/j.vaccine.2019.06.025>
- NASEM (National Academy of Sciences Engineering and Medicine) (2020). *Framework for Equitable Allocation of COVID-19 Vaccine*. Washington, DC: The National Academy Press.
- NCHS (National Center for Health Statistics) (2019). *National Vital Statistics Reports*. Volume 68, No.7. Available at: https://www.cdc.gov/nchs/data/nvsr/nvsr68/nvsr68_07-508.pdf
- O’Driscoll, M., Ribeiro Dos Santos, G., Wang, L., et al. (2021). “Age-Specific Mortality and Immunity Patterns of SARS-CoV-2.” *Nature*, 590: 140-145. <https://doi.org/10.1038/s41586-020-2918-0>
- Peiris, M., Leung, G.M. (2020). “What Can We Expect from First-Generation COVID-19 Vaccines?” *The Lancet*, 396(10261): 1467–1469. [https://doi.org/10.1016/S0140-6736\(20\)31976-0](https://doi.org/10.1016/S0140-6736(20)31976-0)
- Perez-Saez, J., Lauer, S.A., Kaiser, L., Regard, S., Delaporte, E., Guessous, I. et al. (2020). “Serology-Informed Estimates of SARS-CoV-2 Infection Fatality Risk in Geneva, Switzerland.” *The Lancet Infectious Diseases* (published online July 14, 2020). Available at [https://doi.org/10.1016/S1473-3099\(20\)30584-3](https://doi.org/10.1016/S1473-3099(20)30584-3)

- Quaas, M.F., Meya, J.N., Schenk, H., Bos, B., Drupp, M.A., Requate, T. (2020). “The Social Cost of Contacts: Theory and Evidence for the COVID-19 Pandemic in Germany.” *CESifo Working Paper* n. 8347. Available at https://www.cesifo.org/DocDL/cesifo1_wp8347.pdf
- Rheinberger, C.M., Treich, N. (2017). “Attitudes Toward Catastrophe.” *Environmental and Resource Economics*, 67: 609–636. <https://doi.org/10.1007/s10640-016-0033-3>
- Roberts, S. (2020). “Flattening the Coronavirus Curve.” *The New York Times*, March 27, 2020. Available at <https://www.nytimes.com/article/flatten-curve-coronavirus.html>
- Robinson, L. (2007). “How U.S. Government Agencies Value Mortality Risk Reductions.” *Review of Environmental Economics and Policy*, 1(2): 283-299. <https://doi.org/10.1093/reep/rem018>
- Rodríguez, J., Patón, M., Acuña, J.M. (2020). “Prioritisation of Population Groups with the Most Interactions for COVID-19 Vaccination Can Substantially Reduce Total Fatalities.” *medRxiv*. <https://doi.org/10.1101/2020.10.12.20211094>
- Schmidt, H., Gostin, L.O., Williams, M.A. (2020). “Is It Lawful and Ethical to Prioritize Racial Minorities for COVID-19 Vaccines?” *Journal of American Medical Association*, 324(20): 2023-2024. <https://doi.org/10.1001/jama.2020.20571>
- Seligman, B., Ferranna, M., Bloom, D.E. (2021). “Social Determinants of Mortality from COVID-19: A Simulation Study Using NHANES.” *PLOS Medicine*, 18(1): e1003490. <https://doi.org/10.1371/journal.pmed.1003490>
- Streek, H., Schulte, B., Kümmerer, B.M., et al. (2020). “Infection Fatality Rate of SARS-CoV-2 in a Super-Spreading Event in Germany.” *Nature Communications*, 11, 5829. <https://doi.org/10.1038/s41467-020-19509-y>
- Subbaraman, N. (2020). “How to Address the Coronavirus’s Outsized Toll of People of Colour.” *Nature*, 581: 366–367. <https://doi.org/10.1038/d41586-020-01470-x>
- Thunström, L., Newbold, S.C., Finnoff, D., Ashworth, M., Shogren, J.F. (2020). “The Benefits and Costs of Using Social Distancing to Flatten the Curve for COVID-19.” *Journal of Benefit-Cost Analysis*, 11(2): 179–195. <https://doi.org/10.1017/bca.2020.12>
- Toner, E., Barnill, A., Krubiner, C., Bernstein, J. (2020). “Interim Framework for COVID-19 Allocation and Distribution in the United States.” The Johns Hopkins Center for Health Security,

Baltimore, MD. Available at https://www.centerforhealthsecurity.org/our-work/pubs_archive/pubs-pdfs/2020/200819-vaccine-allocation.pdf.

Toxvaerd, F. (2020). “Equilibrium Social Distancing.” *Cambridge Working Papers in Economics*: 2021. Available at <http://www.econ.cam.ac.uk/research-files/repec/cam/pdf/cwpe2021.pdf>

UNIDO (United Nations Industrial Development Organization). (2020). “Managing COVID-19: How the Pandemic Disrupts Global Value Chains.” Available at <https://iap.unido.org/articles/managing-covid-19-how-pandemic-disrupts-global-value-chains>

U.S. Census Bureau, Population Division (2020). “Annual Estimates of the Resident Population for Selected Age Groups by Sex for the United States: April 1, 2010 to July 1, 2019 (NC-EST2019-AGESEX).” Available at: <https://www.census.gov/newsroom/press-kits/2020/population-estimates-detailed.html>

Verity, R., Okell, L.C., Dorigatti, I., Winskill, P., Whittaker, C., Imai, N. et al. (2020). “Estimates of the Severity of Coronavirus Disease 2019: A Model-Based Analysis.” *The Lancet, Infectious Diseases*, 20(6): 669–677. [https://doi.org/10.1016/S1473-3099\(20\)30243-7](https://doi.org/10.1016/S1473-3099(20)30243-7)

Viner, R.M., Mytton, O.T., Bonell, C., et al. (2021). “Susceptibility to SARS-CoV-2 Infection Among Children and Adolescents Compared with Adults: A Systematic Review and Meta-Analysis.” *Journal of American Medical Association, Pediatrics*, 175(2): 143-156. <https://doi.org/10.1001/jamapediatrics.2020.4573>

Viscusi, W.K. (2018). *Pricing Lives: Guideposts for a Safer Society*. Princeton: Princeton University Press.

Vogel, G. (2020). “Sweden’s Gamble.” *Science* 370(6513): 159-163. <https://doi.org/10.1126/science.370.6513.159>

Warfel, J.M., Zimmerman, L.I., Merkel, T.J. (2014). “Acellular Pertussis Vaccines Protect Against Disease but Fail to Prevent Infection and Transmission in a Nonhuman Primate Model.” *Proceedings of the National Academy of Science*, 111(2): 787–792. <https://doi.org/10.1073/pnas.1314688110>

WHO (World Health Organization). (2020). “Coronavirus Disease 2019 (COVID-19): Situation Report, 46.” Available at https://www.who.int/docs/default-source/coronaviruse/situation-reports/20200306-sitrep-46-covid-19.pdf?sfvrsn%3d96b04adf_4

Wu, X., Nethery, R.C., Sabath, B.M., Braun, D., Dominici, F. (2020). “Exposure to Air Pollution and COVID-19 Mortality in the United States: A Nationwide Cross-Sectional Study.” *Science Advances*, 6(45), eabd4049. <https://doi.org/10.1126/sciadv.adb4049>

Yang, W., Kandula, S., Huynh, M., Greene, S.K., Van Wye, G., Li, W., et al. (2020). “Estimating the Infection-Fatality Risk of SARS-CoV-2 in New York City During the Spring 2020 Pandemic Wave: A Model-Based Analysis.” *The Lancet, Infectious Diseases* (published online October 19, 2020). [https://doi.org/10.1016/S1473-3099\(20\)30769-6](https://doi.org/10.1016/S1473-3099(20)30769-6)

Yglesias, M. (2020). “Flattening the Curve isn’t Good Enough.” *Vox* (published online May 6, 2020). Available at <https://www.vox.com/2020/5/6/21241058/coronavirus-mitigation-suppression-flatten-the-curve>

Zhou, F., Yu, T., Du, R., Fan, G., Liu, Y., Liu, Z., et al. (2020). “Clinical Course and Risk Factors for Mortality of Adult Inpatients with COVID-19 in Wuhan, China: A Retrospective Cohort Study.” *Lancet*, 395(10229), 1054–1062. [https://doi.org/10.1016/s0140-6736\(20\)30566-3](https://doi.org/10.1016/s0140-6736(20)30566-3)

APPENDIX A: The value of COVID-19 control policies in the United States

Distribution of costs and fatalities

The population is divided into two age groups, the “young” (age less than 65) and the “old” (65+), and five income groups, corresponding to the quintiles of the income distribution; based on the 2019 U.S. population age structure (U.S. Census Bureau 2020), 84% of the population is young. Let q_i be the proportion of aggregate GDP that individuals in income group i own, with $\sum_{i=1}^5 q_i = 1$. The income distribution is given by $q_i = \{5\%, 10\%, 15\%, 23\%, 47\%\}$. We follow the COVID-19 simulator created by Fleurbaey and co-authors in representing the distribution of fatalities and costs by income quintiles.³⁹

Let σ_{ia} be the proportion of the total income loss caused by the control policy or by the pandemic that an individual belonging to income group $i = \{1, 2, 3, 4, 5\}$ and age group $a = \{y, o\}$ bears:

$$\sigma_{ia} = \frac{5}{N} \frac{q_i^\omega}{\sum_i q_i^\omega} \frac{\sigma_a}{n_a} \quad (1)$$

where ω is the elasticity of costs to income, σ_a the proportion of costs paid by age group a , n_a the proportion of individuals in age group a , N the total population, and $\frac{N}{5}$ the size of each income group.⁴⁰ If $\omega = 1$, individuals pay proportionally to their income; if $\omega < 1$, lower-income quintiles bear a disproportionate amount of the costs. If $\frac{\sigma_a}{n_a} = 1$, all ages suffer an economic loss; if only the young pay the costs of the pandemic, then $\sigma_o = 0$ and $\frac{\sigma_y}{n_y} = \frac{1}{n_y}$, where n_y is the share of young in the population. If the pandemic reduces aggregate income by the amount ΔGDP , each individual in age group a and income group i pays $\sigma_{ia} \Delta GDP$.

Let δ_{ia} be the proportion of individuals in age group a and income group i that die of COVID-19:

$$\delta_{ia} = \frac{5}{N} \frac{q_i^\xi}{\sum_i q_i^\xi} \frac{\delta_a}{n_a} D \quad (2)$$

³⁹ <https://sites.google.com/site/marcfleurbaey/Home/covid>

⁴⁰ There are ten age-income combinations in the example. The proportion of costs paid by each age-income group is $\frac{q_i^\omega}{\sum_i q_i^\omega} \sigma_a$. Since the size of an age-income group is $0.2 n_a N$, each individual in age group a and income group i pays

$$\sigma_{ia} = \frac{q_i^\omega}{\sum_i q_i^\omega} \sigma_a \frac{1}{0.2 n_a N} = \frac{5}{N} \frac{q_i^\omega}{\sum_i q_i^\omega} \frac{\sigma_a}{n_a}.$$

where D is the total number of COVID-19 fatalities, ξ is the elasticity of fatalities to income and δ_a the share of deaths experienced by individuals in age group a . If there are D total COVID-19 fatalities, $D_{ia} = \frac{q_i^\xi}{\sum_i q_i^\xi} \delta_a D$ denotes the number of deaths accruing to individuals in age group a and income group i , and $\delta_{ia} = \frac{D_{ia}}{0.2n_a N}$ the mortality rate of individuals in age group a and income group i (i.e., the share of individuals in age group a and income group i that die of COVID-19). If $\xi = 0$, mortality risk is independent of socioeconomic status (i.e., every income group bears the same number of COVID-19 fatalities). If $\xi < 0$, individuals in low socioeconomic groups bear a disproportionate number of deaths due to COVID-19. If $\frac{\delta_a}{n_a} = 1$, then the mortality rate would be the same in all age groups. Based on the estimates of infection fatality rate by age in Verity et al. (2020) and on the 2019 U.S. population structure (U.S. Census Bureau 2020), the infection fatality rate (i.e., the percentage of fatalities among the infected) of the young is equal to 0.28% and the infection fatality rate of the old is equal to 4.36%. If both age groups have the same probability of getting infected, 25% of COVID-19 fatalities occur among the young. As a consequence, $\delta_y = 0.25$ and $\delta_o = 0.75$.

Per capita annual GDP in the pre-COVID-19 situation is set equal to \$65,000. The per capita income distribution across quintiles is given by \$16,250; \$32,500; \$48,750; \$74,750; and \$152,750.⁴¹ We also assume that, absent any control policy, the pandemic will cause 70% of the population to be infected, resulting in a mortality rate equal to 0.67% of the population. More specifically, 0.2% of the young will die from COVID-19, while the mortality rate of the old is 3.05%. Regarding the distribution of fatalities across income quintiles, we consider two scenarios:

- i) Elasticity of fatalities to income ξ equals -0.5, i.e., low-income quintiles bear a disproportionate number of deaths. The corresponding percentages of fatalities accruing to the 1st, 2nd, 3rd, 4th, and 5th income quintiles are, respectively, 33%, 23%, 19%, 15%, and 11%. Table 1 presents the resulting mortality rates δ_{ia} per income and age group.⁴²

⁴¹ Per-capita income of individuals in income group i is given by: $\frac{q_i N}{0.2 N}$ \$65,000, where N is total population size and $0.2N$ the size of the income group.

⁴² The mortality rates are computed by multiplying the age-specific mortality rates, $\frac{\delta_a D}{n_a N}$ (0.2% for the young and 3.05% for the old), by five times the quintile-specific proportion of deaths, $5 \frac{q_i^\xi}{\sum_i q_i^\xi}$. For example, the mortality rate of young individuals in the 1st income quintile is equal to $0.2 * 5 * 0.33\%$.

- ii) Elasticity of fatalities to income ξ equals 0, i.e., the distribution of fatalities is independent of income. Consequently, all young individuals have a COVID-19 mortality risk of 0.2%, and all old individuals a mortality risk equal to 3.05%.

Likewise, we consider three different cases for the distribution of policy costs:

- i) Elasticity of policy costs to income $\omega = 1$, i.e., the distribution of costs is proportional to income. The percentages of policy costs paid by the 1st, 2nd, 3rd, 4th, and 5th income quintiles are equal to the income distribution q_i .
- ii) Elasticity of policy costs to income $\omega = 0.5$, i.e., the distribution of costs is regressive, or low-income groups bear a higher proportion of economic losses. The percentages of policy costs paid by the 1st, 2nd, 3rd, 4th, and 5th income quintiles are, respectively, 11%, 15%, 19%, 23%, and 33% (while their income share is 5%, 10%, 15%, 23%, 47%).
- iii) Elasticity of policy costs to income $\omega = 1.5$, i.e., the distribution of costs is progressive, or high-income groups bear the brunt of economic losses. The percentages of policy costs paid by the 1st, 2nd, 3rd, 4th, and 5th income quintiles are, respectively, 2%, 6%, 11%, 21%, and 60% (while their income share is 5%, 10%, 15%, 23%, 47%).

Individual well-being

Individual well-being depends on consumption and longevity. Lifetime well-being is time separable, and the period utility is $u(c_{ia}) = \frac{c_{ia}^{1-\lambda}}{1-\lambda} - k$, where λ is the coefficient of relative risk aversion, c_{ia} the consumption of an individual in age group a and income group i , and k the utility at the subsistence consumption level. We assume that in the pre-COVID-19 situation, consumption was constant across ages and equal to the income of the associated quintile: $c_{ia} = c_i = \frac{5}{N} q_i GDP$, where GDP denotes aggregate income. If the pandemic (or the nonpharmaceutical intervention to control it) causes a GDP contraction of X , consumption is equal to $c_{ia} = c_i - \sigma_{ia}X$, where σ_{ia} is the proportion of total loss borne by an individual in age group a and income group i , and its expression is given by (1). To simplify, we assume that all costs occur in one year.

Let μ_{iy} be the pre-COVID-19 fraction of young individuals who die prematurely, that is, before reaching an old age, and μ_{io} the fraction of old individuals who die prematurely in the pre-COVID-19 setting, i.e., before reaching a reasonable length of life, set at 85. Both probabilities depend on the income group, with low-income quintiles facing a higher risk of dying prematurely. Based on

the 2017 U.S. lifetable (NCHS 2019), and averaging across all ages in an age group, we find that 13% of the young die before reaching age 65 and 43% of the old die prematurely, i.e., before reaching age 85. The average age of the young who do not live past 65 is 51, and of the old who do not live past 85 is 76; the life expectancy of the old who do live past age 85 (the assumed desirable longevity) is 88. We assume that the derived fractions of premature deaths apply to individuals in the 3rd income quintile, and we scale them to compute the probabilities of the other income quintiles. Based on Chetty et al. (2016) and following Adler et al. (2021), we assume that the annual probabilities of dying for individuals in the 1st, 2nd, 4th, and 5th quintiles are, respectively, 1.5, 1.3, 0.9, and 0.75 times the annual probability of dying of individuals with median income. As a result, the fractions of young and old that die prematurely for each income quintile are (19%, 57%), (16%, 49%), (13%, 43%), (12%, 40%), and (10%, 34%).

COVID-19 increases the number of individuals who die prematurely. Let μ_{ia}^* be the total fraction of individuals in age group a and income group i who die prematurely, with

$$\mu_{ia}^* = \mu_{ia} + \delta_{ia} \quad (3)$$

and δ_{ia} the fraction of individuals in age group a and income group i who die of COVID-19, whose expression is given in (2).

Let V_{ia} be the expected lifetime utility of individuals in age group a and income group i . For the “old” we have

$$V_{io} = 75u(c_i) + u(c_i - \sigma_{io}X) + (1 - \mu_{io}^*)12u(c_i).$$

The “old” has already lived for 75 years, sustains a consumption loss in the 76th year of age (the average age of the “old”), and is expected to live for an additional 12 years (until age 88) if she does not die prematurely of COVID-19 or of other causes (i.e., $1 - \mu_{io}^*$ of the “old” reach at least the desirable longevity of 85 years).

For the “young” we have

$$V_{iy} = 50u(c_i) + u(c_i - \sigma_{iy}X) + (1 - \mu_{iy}^*)[\mu_{io}25u(c_i) + (1 - \mu_{io})37u(c_i)].$$

The “young” has lived for 50 years and sustains a consumption loss in the 51st year of age (the average age of the young). A proportion μ_{iy}^* dies before the old age (i.e., at age 51) for either COVID-19 or other causes. Among those that survive the old age, a proportion $1 - \mu_{io}$ reaches

the desirable length of life and dies at age 88, thereby living for an additional 37 years; the remaining individuals die at age 76 before reaching the desirable length of life, thereby living for an additional 25 years.

Moreover, let $U_{ia}(t)$ be the realized lifetime utility of an individual in age group a and income group i who dies exactly at $t = \{51, 76, 88\}$:

$$U_{ia}(t) = (t - 1)u(c_i) + u(c_i - \sigma_{ia}X).$$

For example, a young individual who survives the young age, but dies prematurely when old, will have a lifetime realized well-being equal to $U_{iy}(76) = 75u(c_i) + u(c_i - \sigma_{iy}X)$. The second term denotes the fact that, when young, he had to pay the costs of the pandemic for one year.

To calibrate the utility function, we set $\lambda = 2$ (a standard assumption in the literature), and the parameter k is set such that the value of a statistical life of an individual with median income at age 40 is equal to 10 million.⁴³

Individual willingness to pay to eliminate pandemic risks

Given the definition of expected individual lifetime utility, we can compute the amount of own income each individual is willing to sacrifice to get rid of COVID-19, that is, to eliminate the COVID-19-related mortality risk and income losses absent any intervention.

Consider an old individual and let ΔGDP be the aggregate income loss associated with no control over COVID-19. The individual willingness to pay wtp_{io} of an old individual in income group i is defined as

$$\begin{aligned} V_{io}(\mu_{io}, wtp_{io}) &\equiv 75u(c_i) + u(c_i - wtp_{io}) + (1 - \mu_{io})12u(c_i) \\ &= 75u(c_i) + u(c_i - \sigma_{io}\Delta GDP) + (1 - \mu_{io}^*)12u(c_i) \equiv V_{io}(\mu_{io}^*, \Delta GDP) \end{aligned} \quad (4)$$

Similarly, the willingness to pay of a young individual in income group i is defined as

⁴³ Consider a marginal change in the baseline probability of dying young. Assuming consumption is constant through life, the formula for the VSL is $VSL = \frac{u(c)}{u'(c)} LE_{40}$, where LE_{40} is the remaining life expectancy of a 40-year-old individual (set at 40.16) and $u'(c)$ the marginal utility of consumption. We set consumption at \$48,750 (the consumption value of an individual with median income) and solve the expression to derive k . We find that $k = -0.000125288$.

$$\begin{aligned}
V_{iy}(\mu_{iy}, wtp_{iy}) &\equiv 50u(c_i) + u(c_i - wtp_{iy}) + (1 - \mu_{iy})[\mu_{io}25u(c_i) + (1 - \mu_{io})37u(c_i)] \\
&= 50u(c_i) + u(c_i - \sigma_{iy}\Delta GDP) + (1 - \mu_{iy}^*)[\mu_{io}25u(c_i) + (1 - \mu_{io})37u(c_i)] \\
&\equiv V_{iy}(\mu_{iy}^*, \Delta GDP)
\end{aligned} \tag{5}$$

Table 2 summarizes the individuals' willingness to pay to eliminate the pandemic risk by income and age group, under the assumption that an uncontrolled pandemic does not cause any income loss (i.e., $\Delta GDP = 0$).⁴⁴

Evaluation methods

Under BCA, the value of a policy to eliminate the risk of COVID-19 is equal to the unweighted sum of individuals' willingness to pay to eliminate COVID-19. Let wtp_{ia} be the willingness to pay of an individual in age group a and income group i . Under BCA, the total value X^{BCA} of the intervention is given by:

$$X^{BCA} = N \frac{1}{5} \sum_{i=1}^5 (n_y wtp_{iy} + n_o wtp_{io}),$$

where n_y is the proportion of young people in the population, $n_o = 1 - n_y$ the proportion of old people, and N the total (prepandemic) population.

We consider three alternative SWFs:

1. Utilitarianism. The utilitarian SWF is the sum of expected individual lifetime well-being:

$$S^U = N \frac{1}{5} \sum_{i=1}^5 (n_y V_{iy} + (1 - n_y) V_{io}).$$

2. Ex-ante prioritarianism. The ex-ante prioritarian SWF is the sum of a concave transformation of expected individual lifetime well-being:

$$S^{EA} = N \frac{1}{5} \sum_{i=1}^5 (n_y g(V_{iy}) + (1 - n_y) g(V_{io})),$$

⁴⁴ Applying the definitions of willingness to pay in (4) and (5), and given that $u(c_i) = \frac{c_i^{1-\lambda}}{1-\lambda} + k$ and $\lambda = 2$, the willingness to pay of an individual in age group a and income group i solves the following equation:

$$wtp_{ia} = c_i - \left[k - \left(k - \frac{1}{c_i} \right) \Delta L_{ia} \right]^{-1}$$

with $\Delta L_{iy} = 1 - \delta_{iy}(37 - 12\mu_{io})$ and $\Delta L_{io} = 1 - \delta_{io}12$.

where g is a strictly increasing and strictly concave function, denoting the degree of priority to the worse-off in (expected) lifetime well-being.

3. Ex-post prioritarianism. The ex-post prioritarian SWF is the expected sum of a strictly increasing and strictly concave transformation of realized individual lifetime well-being:

$$\begin{aligned} S^{EP} = N \frac{1}{5} n_y \sum_{i=1}^5 & \left[\mu_{iy}^* g(U_{iy}(51)) + (1 - \mu_{iy}^*) \mu_{io} g(U_{iy}(76)) \right. \\ & \left. + (1 - \mu_{iy}^*)(1 - \mu_{io}) g(U_{iy}(88)) \right] \\ & + N \frac{1}{5} (1 - n_y) \sum_{i=1}^5 \left[\mu_{io}^* g(U_{io}(76)) + (1 - \mu_{io}^*) g(U_{io}(88)) \right]. \end{aligned}$$

We assume that the function g belongs to the Atkinson family, with γ the priority parameter:

$$g(w) = \frac{w^{1-\gamma}}{1-\gamma}.$$

The maximum GDP loss X associated with the policy that the policy maker finds acceptable is defined as the reduction in current aggregate GDP that society is willing to suffer to eliminate the pandemic. Let $S^U(\mu_{iy}^*, \mu_{io}^*, \Delta GDP)$, $S^{EA}(\mu_{iy}^*, \mu_{io}^*, \Delta GDP)$, and $S^{EP}(\mu_{iy}^*, \mu_{io}^*, \Delta GDP)$ be the utilitarian, ex-ante prioritarian, and ex-post prioritarian social welfare levels with COVID-19 and no control policy, where ΔGDP is the potential income loss associated to an uncontrolled policy (note that in the examples considered in the text ΔGDP is set to zero). The utilitarian societal burden X^U is defined as

$$S^U(\mu_{iy}, \mu_{io}, X^U) \equiv N \frac{1}{5} \sum_{i=1}^5 \left(n_y V_{iy}(\mu_{iy}, X^U) + (1 - n_y) V_{io}(\mu_{io}, X^U) \right) = S^U(\mu_{iy}^*, \mu_{io}^*, \Delta GDP),$$

with $V_{iy}(\mu_{iy}, X^U) = 50u(c_i) + u(c_i - \sigma_{iy}X^U) + (1 - \mu_{iy})[25u(c_i) + (1 - \mu_{io})12u(c_i)]$ and $V_{io}(\mu_{io}, X^U) = 75u(c_i) + u(c_i - \sigma_{io}X^U) + (1 - \mu_{io})12u(c_i)$.

The ex-ante prioritarian societal burden X^{EA} is defined as

$$\begin{aligned} S^{EA}(\mu_{iy}, \mu_{io}, X^{EA}) & \equiv N \frac{1}{5} \sum_{i=1}^5 \left(n_y g(V_{iy}(\mu_{iy}, X^{EA})) + (1 - n_y) g(V_{io}(\mu_{io}, X^{EA})) \right) \\ & = S^{EA}(\mu_{iy}^*, \mu_{io}^*, \Delta GDP). \end{aligned}$$

Finally, the ex-post prioritarian societal burden X^{EP} is defined as

$$\begin{aligned}
S^{EP}(\mu_{iy}, \mu_{io}, X^{EP}) \\
&\equiv N \frac{1}{5} n_y \sum_{i=1}^5 \left[\mu_{iy} g(U_{iy}(51; X^{EP})) + (1 - \mu_{iy}) \mu_{io} g(U_{iy}(76; X^{EP})) \right. \\
&\quad \left. + (1 - \mu_{iy})(1 - \mu_{io}) g(U_{iy}(88; X^{EP})) \right] \\
&\quad + N \frac{1}{5} (1 - n_y) \sum_{i=1}^5 \left[\mu_{io} g(U_{io}(76; X^{EP})) + (1 - \mu_{io}) g(U_{io}(88; X^{EP})) \right] \\
&= S^{EP}(\mu_{iy}^*, \mu_{io}^*, \Delta GDP),
\end{aligned}$$

with $U_{ia}(t; X^{EP}) = (t - 1)u(c_i) + u(c_i - \sigma_{ia}X^{EP})$.

The results in Table 3 correspond to the value X^U, X^{EA}, X^{EP} computed by solving the previous expressions.

Additional tables

Table A.1. Individuals' willingness to pay as a percentage of own income for a policy that eliminates the risk of death from COVID-19, by age group and income quintile

	1 st quintile	2 nd quintile	3 rd quintile	4 th quintile	5 th quintile
Young	5.8%	15.8%	24.2%	34.6%	54%
Old	27.5%	53%	65.2%	75.4%	86.9%

Assumptions: COVID-19 fatality risk is independent of income; the young die of COVID-19 with probability 0.2%, the old with probability 3.05%; there is no income loss associated to an uncontrolled pandemic.

APPENDIX B: The allocation of vaccine doses

Consider again the example in Section 4. Individuals are divided into five income and two age groups, and life expectancies depend on the income quintile. We denoted δ_{ia} the probability that an individual in income quintile i and age group a dies because of COVID-19. If the individual is vaccinated, her COVID-19 mortality rate decreases to $\delta_{ia}(1 - e_a)$, where e_a is the effectiveness of the vaccine for individuals of age a . We consider the case where low-income quintiles are overburdened with deaths. The COVID-19 mortality rates by income and age group are given in Table 1.

Consider an old person in income group i and suppose for simplicity that COVID-19 does not entail any income loss, only the risk of premature mortality. If the individual is not vaccinated, her expected well-being is $V_{io}^{no\ vac} = 76u(c_i) + (1 - \mu_{io} - \delta_{io})12u(c_i)$. If the individual is vaccinated, her well-being becomes $V_{io}^{vac} = 76u(c_i) + (1 - \mu_{io} - \delta_{io}(1 - e_o))12u(c_i)$. The well-being gain from being vaccinated then equals

$$\Delta V_{io} = 12u(c_i)e_o\delta_{io}.$$

Consider now a young person in income group i . If the individual is not vaccinated, her expected well-being is $V_{iy}^{no\ vac} = 51u(c_i) + (1 - \mu_{iy} - \delta_{iy})[25u(c_i) + (1 - \mu_{io})12u(c_i)]$. If the individual is vaccinated, her well-being becomes $V_{iy}^{vac} = 51u(c_i) + (1 - \mu_{iy} - \delta_{iy}(1 - e_y))[25u(c_i) + (1 - \mu_{io})12u(c_i)]$. The well-being gain from being vaccinated then equals

$$\Delta V_{iy} = [37 - \mu_{io}12]u(c_i)e_y\delta_{iy}.$$

In the utilitarian case, social welfare is the sum of individuals' expected lifetime well-being, $S^U = \frac{N}{5}\sum_i(n_y V_{iy} + (1 - n_y)V_{io})$. Therefore, one dose of vaccine to allocate will be given to the individual who gains more in terms of increase in expected well-being.

Let ΔV_{1o} be the utilitarian value of vaccinating an old individual in the 1st income quintile. We can compute the utilitarian value of vaccinating an individual at age a and income group i as a proportion of the utilitarian value of vaccinating an old individual in the 1st income quintile:

$$\frac{\Delta V_{iy}}{\Delta V_{1o}} = \frac{[37 - \mu_{io}12]u(c_i)}{12u(c_1)} \frac{e_y\delta_{iy}}{e_o\delta_{1o}}$$

$$\frac{\Delta V_{io}}{\Delta V_{1o}} = \frac{12u(c_i) e_o \delta_{io}}{12u(c_1) e_o \delta_{1o}}$$

In both expressions, the first term represents the utilitarian social value of a given reduction in mortality risk $SVRR_{ia}^U$ (e.g., a 1 percentage point reduction) experienced by an individual at age a and in income group i as a proportion of the utilitarian value $SVRR_{1o}^U$ attached to a comparable reduction accruing to an old individual in the 1st income quintile (Table 4). The second term is the health benefit of vaccination, as compared with the health benefit enjoyed by an old person in the 1st income quintile. The terms $\frac{\Delta V_{iy}}{\Delta V_{1o}} e_o \delta_{1o}$ and $\frac{\Delta V_{io}}{\Delta V_{1o}} e_o \delta_{1o}$ are the utilitarian-adjusted mortality risk reductions achieved by vaccinating an individual (Table 5). For example, vaccinating a young individual in the 1st income quintile is equivalent, from a utilitarian point of view, to reducing the mortality risk of an old individual in the 1st income quintile by $\frac{\Delta V_{1y}}{\Delta V_{1o}} e_o \delta_{1o}$.

In the ex-ante prioritarian case, social welfare is the sum of strictly increasing and strictly concave transformations of individual expected well-being: $S^{EA} = \frac{N}{5} \sum_i (n_y g(V_{iy}) + (1 - n_y) g(V_{io}))$, with g strictly increasing and strictly concave. The ex-ante prioritarian welfare gain from vaccinating a young person in income group i is given by

$$\Delta S_{iy}^{EA} = g(V_{iy}^{vax}) - g(V_{iy}^{no\ vax}) \simeq g'(V_{iy}^{no\ vax}) \Delta V_{iy}.$$

The second expression derives from approximating the g function around the no vaccination case. Note that the approximation is good only if the change in expected well-being is sufficiently small (an assumption that, for the sake of simplicity, we make here). Similarly, the ex-ante prioritarian welfare gain from vaccinating an old person in income group i is given by

$$\Delta S_{io}^{EA} = g(V_{io}^{vax}) - g(V_{io}^{no\ vax}) \simeq g'(V_{io}^{no\ vax}) \Delta V_{io}.$$

The ex-ante prioritarian value of vaccinating an individual at age a and in income group i as a proportion of the ex-ante prioritarian value of vaccinating an old individual in the 1st income quintile is

$$\frac{\Delta S_{iy}^{EA}}{\Delta S_{1o}^{EA}} = \frac{g'(V_{iy}^{no\ vax}) [37 - \mu_{io} 12] u(c_i) e_y \delta_{iy}}{g'(V_{1o}^{no\ vax}) 12u(c_1) e_o \delta_{1o}}$$

$$\frac{\Delta S_{io}^{EA}}{\Delta S_{1o}^{EA}} = \frac{g'(V_{io}^{no\ vax}) 12u(c_i) e_o \delta_{io}}{g'(V_{1o}^{no\ vax}) 12u(c_1) e_o \delta_{1o}}$$

where, in both expressions, the first term is the ex-ante prioritarian social value $SVRR_{ia}^{EA}$ attached to a given reduction in mortality risk experienced by an individual at age a and in income group i as a proportion of the ex-ante prioritarian value $SVRR_{1o}^{EA}$ of a comparable risk reduction accruing to an old individual in the 1st income quintile (Tables 6 and B.1). The results in Tables 7 and B.2 are determined by multiplying the previous expressions by $e_o\delta_{1o}$.

In the ex-post prioritarian case, social welfare is the expected sum of concave transformations of individual realized well-being: $S^{EP} = N\frac{1}{5}n_y \sum_{i=1}^5 [(\mu_{iy} + \delta_{iy})g(U_i(51)) + (1 - \mu_{iy} - \delta_{iy})\mu_{io}g(U_i(76)) + (1 - \mu_{iy} - \delta_{iy})(1 - \mu_{io})g(U_i(88))] + N\frac{1}{5}(1 - n_y) \sum_{i=1}^5 [(\mu_{io} + \delta_{io})g(U_i(76)) + (1 - \mu_{io} - \delta_{io})g(U_i(88))]$, with $U_i(a) = u(c_i)a$. The ex-post prioritarian welfare gain from vaccinating an old person in income group i is given by

$$\Delta S_{io}^{EP} = \delta_{io}[g(U_i(88)) - g(U_i(76))] \simeq \delta_{io}g'(U_i(76))12u(c_i) = g'(U_i(76))\Delta V_{io}.$$

Once again, we implicitly assume that the change in realized well-being is small enough such that the approximation is good.

The ex-post prioritarian welfare gain from vaccinating a young person in income group i is given by

$$\begin{aligned} \Delta S_{iy}^{EP} &= \delta_{iy} [g(U_i(88)) - g(U_i(51)) - \mu_{io}(g(U_i(88)) - g(U_i(76)))] \\ &\simeq \delta_{iy} [g'(U_i(51))37u(c_i) - \mu_{io}g'(U_i(76))12u(c_i)] \\ &= g'(U_i(51))\Delta V_{iy} + \delta_{iy}\mu_{io}12u(c_i)[g'(U_i(51)) - g'(U_i(76))]. \end{aligned}$$

The ex-post prioritarian value of vaccinating an individual at age a and in income group i as a proportion of the ex-post prioritarian value of vaccinating an old individual in the 1st income quintile is

$$\begin{aligned} \frac{\Delta S_{iy}^{EP}}{\Delta S_{1o}^{EP}} &= \frac{g'(U_i(51))37u(c_i) - \mu_{io}g'(U_i(76))12u(c_i)}{g'(U_1(76))12u(c_1)} \frac{e_y\delta_{iy}}{e_o\delta_{1o}} \\ \frac{\Delta S_{io}^{EP}}{\Delta S_{1o}^{EP}} &= \frac{g'(U_i(76))12u(c_i)}{g'(U_1(76))12u(c_1)} \frac{e_o\delta_{io}}{e_o\delta_{1o}} \end{aligned}$$

where, in both expressions, the first term is the ex-post prioritarian value $SVRR_{ia}^{EP}$ attached to a given reduction in mortality risk experienced by an individual at age a and in income group i as a proportion of the ex-post prioritarian value $SVRR_{10}^{EP}$ of a comparable risk reduction accruing to an old individual in the 1st income quintile (Tables 8 and B.3). The results in Tables 9 and B.4 are determined by multiplying the previous expressions by $e_o\delta_{10}$.

Additional tables

Table B.1. Ex-ante prioritarian SVRR compared with the ex-ante prioritarian SVRR attached to an older person belonging to the 1st income quintile, by age and income group.

	1 st quintile	2 nd quintile	3 rd quintile	4 th quintile	5 th quintile
Young	2.69	2.71	2.72	2.73	2.74
Old	1	0.99	0.98	0.97	0.96

Assumption: Atkinson SWF with priority parameter equal to 1.

Table B.2. Ex-ante-prioritarian-adjusted reduction in COVID-19 mortality risk if the vaccine is 100% effective for all age groups, by age and income group.

	1 st quintile	2 nd quintile	3 rd quintile	4 th quintile	5 th quintile
Young	0.86%	0.61%	0.5%	0.41%	0.29%
Old	4.96%	3.46%	2.8%	2.25%	1.56%

Assumption: Atkinson SWF with priority parameter equal to 1. If the vaccine is less than 100% effective, the values in the table should be multiplied by the reduced effectiveness (e.g., if effectiveness is 60%, the ex-ante-prioritarian-adjusted reduction for a young individual in the 1st income quintile is $0.6 \times 0.86\%$).

Table B.3. Ex-post prioritarian SVRR compared with the ex-post prioritarian SVRR attached to an older person belonging to the 1st income quintile, by age and income group.

	1 st quintile	2 nd quintile	3 rd quintile	4 th quintile	5 th quintile
Young	3.75	3.87	3.95	4	4.09
Old	1	1	1	1	1

Assumption: Atkinson SWF with priority parameter equal to 1. When the priority parameter is equal to 1, the approximated formula of the ex-post prioritarian SVRR for an old person is $g'(U_i(76))\Delta V_{io} = [76u(c_i)]^{-1}12u(c_i) = \frac{12}{76}$, which is constant across income groups.

Table B.4. Ex-post-prioritarian-adjusted reduction in COVID-19 mortality risk if the vaccine is 100% effective for all age groups, by age and income group.

	1 st quintile	2 nd quintile	3 rd quintile	4 th quintile	5 th quintile
young	1.19%	0.87%	0.73%	0.59%	0.43%
old	4.96%	3.51%	2.87%	2.31%	1.62%

Assumption: Atkinson SWF with priority parameter equal to 1. If the vaccine is less than 100% effective, the values in the table should be multiplied by the reduced effectiveness (e.g., if effectiveness is 60%, the ex-post-prioritarian-adjusted reduction for a young individual in the 1st income quintile is 0.6×1.19%).