ORIGINAL PAPER

Optimal lockdown and social welfare

Pierre Pestieau[1,2] · Gregory Ponthiere[3]

Abstract

This paper examines the robustness of the optimal lockdown strategy to the postulated social welfare criterion. We show that utilitarianism can, under some conditions, imply a COVID-19 variant of Parfit's (1984) Repugnant Conclusion: for any (interior) lockdown with life periods of low quality, there must be a stricter lockdown that is regarded as better, even though this reduces the quality of life periods even more. On the contrary, the ex post egalitarian criterion (giving priority to the worst-off ex post) implies zero lockdown. Varying between its minimal and its maximal levels, the optimal lockdown is not robust to the postulated ethical criterion. We also identify a general ethical dilemma between the goal of saving lives (modeled by the Survivors Number Count axiom) and the goal of giving priority to the worst-off (Hammond Equity).

1 Introduction

In order to battle against the COVID-19 pandemic and to save lives, governments around the world relied massively on lockdown strategies during spring 2020. At

✉ Gregory Ponthiere
gregory.ponthiere@uclouvain.be

Pierre Pestieau
p.pestieau@ulg.ac.be

[1] University of Liege, Liege, Belgium

[2] CORE, Louvain-la-Neuve, Belgium

[3] UCLouvain, Hoover Chair in Economic and Social Ethics, Collège Dupriez, Office D307, 3, Place Montesquieu, 1348, Louvain la Neuve, Belgium

the beginning of April 2020, about 3.9 billion humans—that is, more than half of humanity—were asked or ordered to stay at home to prevent further expansion of the pandemic.[1]

Those (more or less strict) lockdown strategies saved lives, but at the cost of causing a major macroeconomic slowdown and a rise in unemployment. According to Deb et al. (2020), the Great Lockdown reduced the number of fatalities by about 90% in comparison to a baseline with no containment policies. But at the same time, the lockdown had also a major impact on economic performance. According to the OECD (2020), the scale of the decline in the level of output is equivalent to about 2 percentage points of GDP growth lost for each month of strict lockdown. As a consequence, the 3 months of strict lockdown in spring 2020 imply an annual GDP growth that was, for 2020, lowered by about 6 percentage points, and thus unambiguously negative, i.e., a severe contraction of economic activity.[2] The macroeconomic impact of COVID-19 is studied in more details in Atkeson (2020) and Eichenbaum et al. (2020).

The design of optimal lockdown requires weighting the gains and losses from maintaining the economy under more or less strong sanitary constraints. Recent papers by Acemoglu et al. (2020), Alvarez et al. (2020), and Gollier (2020) examined, using various theoretical settings, the optimal lockdown strategy within dynamic economic models including the (potentially amended) susceptible, infected, and recovered (SIR) structure of epidemiological frameworks. Those papers studied not only the optimal extent of the lockdown, but, also, the optimal timing of the lockdown.[3]

The goal of this paper is to focus on another aspect of the optimal lockdown problem: its robustness to the postulated social welfare criterion. The research question that we would like to address in this paper is to know how robust the optimal lockdown strategy is to the underlying ethical criterion.

Our research question is motivated by the fact that the design of an optimal lockdown strategy is an inherently *prescriptive* task. Deriving an optimal lockdown is, by nature, normative, in the sense that this tells us something about what governments or national health agencies *should do*. As emphasized by Moore (1903), a normative proposition—whatever it is—cannot be deduced from positive premises only, but requires necessarily, among the premises, at least one normative proposition. Any policy can only be "optimal" with respect to some pursued goal, which is normative in nature. What is true in general is also true in the times of COVID-19: any "optimal" lockdown can only be "optimal" given some pursued goal. As a consequence, the study of the optimal lockdown cannot avoid a discussion of its normative foundations.

During the COVID-19 crisis, governments gave the impression that the pursued goal was to "save as many lives as possible". Whereas further research would be

---

[1]See Euronews (2020).

[2]Obviously, those figures will have to be updated, in the light of what takes place in 2021, where the COVID-19 epidemic is still active.

[3]Other recent papers on optimal size and timing of the lockdown include Bosi et al. (2021), Garriga et al. (2020), and Piguillem and Shi (2020).

needed to "rationalize" government's strategies during the COVID-19 crisis, one can raise doubts about a social objective that would consist in saving as many lives as possible. Such a normative objective would ignore all other aspects of life, and, as such, would be too restrictive. Ideally, a social objective should incorporate and weight all aspects of life (not only the quantity of life, but, also, the quality of life). This paper explores the implications of several distinct social welfare criteria for the design of optimal lockdown strategies.

In order to address that issue, this study develops a dynamic model of the human life cycle with risky lifetime. In our model, an epidemic, which has its own dynamics, increases mortality and reduces the proportion of old individuals in the population. In that economy, the government has only *one* policy instrument against the epidemic: a lockdown strategy, which decreases the prevalence of the epidemic and the mortality due to the epidemic, but at the cost of reducing consumption possibilities and causing direct disutility.[4] We examine the solution of the optimal lockdown problem under various social welfare criteria.

The model studied in this paper is designed so as to be the simplest framework allowing us to examine how different ethical criteria address the trade-offs between the goals of saving lives and economic prosperity. Our model includes, unlike SIR models, only two states: individuals can be either alive or dead, and the average age at death varies with the prevalence of the pandemic. But despite its simplicity, our model accounts for the fact that individuals of different ages are not equally vulnerable to the pandemic. The prevalence of the pandemic is modelled *as an equilibrium phenomenon*, whose level depends on the strength of lockdown policies. Following Acemoglu et al. (2020), we consider not only the design of optimal (undifferentiated) lockdown policies (as those carried out in 2020), but, also, the design of optimal age-differentiated lockdown policies (with a degree differing between the young and the old).

When examining the design of optimal lockdown policies, this paper differs from the existing literature on optimal timing issues (optimal "stop-and-go" policies). Unlike that literature, this paper focuses on the selection of stationary lockdown policies allowing to reach the equilibrium prevalence of the pandemic, which can be interpreted as a "target" selected by the public authority. Abstracting from transitional issues allows us to pay more attention to the determinants of the selection of the target. Although a simplification, this theoretical shortcut is most relevant for policymakers facing trade-offs between saving lives and prosperity. Governments around the world designed their policies by fixing a *target for the prevalence of the pandemic*, a target with respect to which they adjusted lockdown policies up and down. Our setting allows us to focus on the selection of that target, whose level reveals how the trade-off between saving lives and prosperity is dealt with. This paper aims at examining how different ethical doctrines address that difficult trade-off.

Our main result is that the optimal lockdown strategy is not robust to the postulated social welfare criterion. Utilitarianism leads to a strictly positive lockdown level,

---

[4]This paper considers thus a second-best framework, where other policy instruments, such as vaccination policies, are not available.

and can also, under some conditions, imply a COVID-19 variant of Parfit's (1984) Repugnant Conclusion: for any (interior) lockdown with life periods of low quality, there must be a stricter lockdown that is regarded as better, even though this reduces the quality of life periods even more. On the contrary, the ex post egalitarian criterion, which gives priority to the worst-off ex post (once individual durations of life are known), implies a zero lockdown, a result in sharp contrast with utilitarianism.[5]

Although the sensitivity of an optimal policy to the pursued social objective could be anticipated, the contribution of this paper is to show that, in the case of the design of optimal lockdown, the lack of robustness takes *an extreme form*: the optimal lockdown varies between its *minimal* level (under ex post egalitarianism) and its *maximal* level (prevailing, in general, under utilitarianism). The strong sensitivity of the optimal lockdown policy to the postulated social criterion is robust to extending the set of available policy instruments to allow for age-differentiated lockdown policies. Under utilitarianism, there is still, in general, some support for positive age-differentiated lockdown, whereas the ex post egalitarian optimum involves zero lockdown for all.

Beyond the comparison of optimal lockdown under utilitarianism and ex post egalitarianism, our results point to a general ethical dilemma between two goals: on the one hand, saving lives (modeled by the Survivors Number Count axiom), and, on the other hand, giving priority to the worst-off (Hammond Equity). Preventive policies like the lockdown achieve the first goal, but at the cost of deteriorating the situation of the most disadvantaged in the society.

This paper is organized as follows. Section 2 presents the model. The optimal lockdown is then derived under the utilitarian criterion in Section 3. Section 4 identifies the condition under which the utilitarian criterion leads to a COVID-19 variant of Parfit's Repugnant Conclusion. The ethical dilemma between saving lives and Hammond Equity is studied in Section 5. Section 6 characterizes the optimal lockdown under the ex post egalitarian criterion. In order to study the robustness of our results to the set of available policy instruments Section 7 examines the optimal age-specific lockdown policies under utilitarianism and ex post egalitarianism. Section 8 concludes.

2 The model

Let us consider a simple model of the human life cycle in the presence of an epidemic. Time is discrete, and goes from $t = 0$ to $t = +\infty$. The model is a two-period overlapping generations economy: there are two ages of life—the young age and the old age—and the duration of each period of life is normalized to unity.[6] Young individuals are active, while old individuals are retired. The number of young adults

---

[5]The ex post egalitarian social criterion was previously studied by Fleurbaey and Ponthiere (2013) and Fleurbaey et al. (2014).

[6]For the sake of simplicity, we abstract here from childhood.

is denoted by $N > 0$. For the sake of simplicity, we assume that all cohorts have the same size N (i.e., no population growth).[7]

The survival process An epidemic takes place in the society, and increases the strength of mortality. Let us denote by $0 \leq S_t < 1$ the probability to survive to the old age at time t. Abstracting from infant and childhood mortality, life expectancy at birth in our economy is thus equal to $1 + S_t$.[8]

The probability of survival to the old age is decreasing with the degree of prevalence of the epidemic, denoted by $0 \leq I_t \leq 1$, and increasing with the strength of the lockdown, denoted by Z_t. The degree of lockdown can go from $Z_t = 0$ (no lockdown) to $Z_t = 1$ (full lockdown). We have:[9]

$$S_t = S(I_t, Z_t) \tag{1}$$

In the absence of any epidemic (i.e., $I_t = 0$) and lockdown, we have $S_t = S(0, 0) = \bar{S} < 1$. The first-order derivatives of $S(I, Z)$ satisfy: $S_I < 0$, $S_Z > 0$ when $I_t > 0$ and $S_Z = 0$ when $I_t = 0$.

Using the Law of Large Numbers, the number of old persons at time t is: $N S(I_t, Z_t)$. The strength of the pandemic only affects the number of old individuals, but not the number of young individuals (N). This feature of the model captures an important asymmetry concerning the impact of the COVID-19 epidemic on mortality: the epidemic increased the mortality of the old, while the mortality of the young remained low and only marginally affected.[10]

The prevalence of the epidemic The prevalence of the epidemic at time t depends on prevalence at time $t - 1$, and on the degree of lockdown imposed by the government at time $t - 1$, according to the law:

$$I_t = \frac{M + I_{t-1}}{1 + I_{t-1} + Z_{t-1}} \tag{2}$$

where $0 < M < 1$. That simple dynamic law ensures that the prevalence of the epidemic I_t belongs to the $[0, 1]$ interval, is increasing in the past prevalence I_{t-1}, and decreasing in the strength of the lockdown Z_{t-1}. We have:

$$\Delta I_t \equiv I_t - I_{t-1} = \frac{M - I_{t-1}^2 - I_{t-1} Z_{t-1}}{1 + I_{t-1} + Z_{t-1}} \tag{3}$$

---

[7]That assumption prevents us from considering the effect of unequal cohort sizes on the pandemic. Note, however, that replacement fertility is a standard assumption when studying countries having reached advanced stages in the demographic transition (Lee 2003).

[8]It is assumed that all young individuals live the entire young age (of duration 1), and that only a fraction S of them will reach the old age (also of duration 1).

[9]We assume that individuals cannot invest in preventive behaviors increasing survival chances. The only way to reduce mortality is through lockdown policies. This simplifying assumption is made to be as close as possible to the context of COVID-19, where lockdown policies imposed themselves as major instruments against the pandemic.

[10]Abstracting from infant mortality and young adult mortality has a cost in terms of realism. But introducing mortality at younger ages would complicate the model without bringing new insights for the study of the general trade-off between saving lives and prosperity.

In the absence of lockdown ($Z_{t-1} = 0$), the prevalence of the epidemic grows (i.e., $\Delta I_t > 0$) when $M > I_{t-1}^2$. That condition is quite weak when the prevalence of the epidemic is initially low, as it is always the case with epidemics in the real world. On the contrary, the prevalence declines (i.e., $\Delta I_t < 0$) in the absence of lockdown when $M < I_{t-1}^2$, which can only be achieved when the prevalence is quite high.

Introducing a lockdown $Z_{t-1} > 0$ reduces the growth of the prevalence of the epidemic. The growth of the prevalence is strictly positive when $M > I_{t-1}^2 + I_{t-1}Z_{t-1}$, and declines when $M < I_{t-1}^2 + I_{t-1}Z_{t-1}$. Thus, imposing a lockdown reduces the growth of the prevalence of the epidemic, and can also, if the lockdown is sufficiently strict, decrease the prevalence of the epidemic.

The production process The output is denoted by $Y_t > 0$. For simplicity, production involves only labor. Young individuals supply one unit of labor inelastically. The output obtained from a given quantity of labor depends on the extent of the lockdown policy decided by the government. The stronger the lockdown is, and the lower the output is, for a given quantity of labor.

The production process is described by the following production function:

$$Y_t = F(N, Z_t) \tag{4}$$

where $F_N > 0$, $F_{NN} < 0$ (decreasing marginal productivity of labor), as well as $F_Z < 0$. We assume also that, for any level of labor, the output remains strictly positive even when the lockdown chosen by the government takes its maximal value, that is, that: $F(N, 1) = \hat{Y} > 0$.[11]

The resource constraint At any period of time, the production is consumed by the population. Let us denote by $c_t \geq 0$ the consumption of each young adult at time t. Let us assume that each old adult receives a defined benefit Pay-As-You-Go pension $b_t \geq 0$.[12] Thus, the following resource constraint holds in the economy at any period of time:

$$F(N, Z_t) = Nc_t + NS(I_t, Z_t)b_t \tag{5}$$

The left-hand side (LHS) is equal to the total production at time t, whereas the right-hand side (RHS) is the total consumption, of either young adults (first term) or old adults (second term). The extent of the lockdown policy chosen by the government affects the resource constraint on the two sides. On the one hand, a stronger lockdown reduces the available output (LHS); on the other hand, the degree of the lockdown also increases the proportion of old individuals in the population, which increases aggregate consumption of the old (RHS).

---

[11] For simplicity, we assume that the only channel through which the pandemic affects the production process is through the lockdown policy. Thus, we do not take into account here the direct effects of the pandemic on labor productivity (through the health of workers). This assumption can be justified as a first approximation: it reflects the fact that the pandemic affected much more the elderly than the young adults.

[12] Note that our results would not be affected if we were assuming a PAYGO system with defined contributions instead of defined benefits.

Rewriting the consumption of each young adult as follows:

$$c_t = \frac{F(N, Z_t)}{N} - S(I_t, Z_t) b_t \tag{6}$$

it appears that the lockdown imposes a double pressure on the consumption of the young, by reducing output per worker (first term of RHS) and by increasing the transfer of each worker to the elderly (second term of RHS).

Preferences Individuals are assumed to have preferences over lotteries of life that satisfy the expected utility hypothesis. Expected lifetime well-being for a young adult at time t is given by:[13]

$$EU_t = u(c_t) - vZ_t + S(I_t, Z_t) \left[u(b_{t+1}) - vZ_{t+1} \right] \tag{7}$$

where the utility function $u(\cdot)$ is increasing and concave in consumption.[14]

Expression Eq. 7 represents individual preferences, and, as such, captures how individuals value consumption and lifetime.[15] Following Becker et al. (2005), we assume that there exists a finite consumption level $\bar{b} > 0$ that makes a person indifferent, in the absence of lockdown, between life and death. When consumption exceeds $\bar{b}$, the person would, in the absence of lockdown, prefer being alive to being dead. On the contrary, when consumption is below $\bar{b}$, the person prefers, in the absence of lockdown, being dead to being alive.[16] Throughout this paper, we consider an advanced economy, where consumption lies above the threshold, making lives worth living.

The parameter $v \geq 0$ captures the pure disutility from the lockdown, that is, the welfare loss that arises independently from the fall of consumption. The pure disutility from the lockdown vZ_t has various sources. The lockdown restricts fundamental liberties (freedom of circulation, right of assembly, etc.) and limits social contacts (with family and friends), which cause welfare losses. The term vZ_t allows us to incorporate in the analysis all lockdown-related welfare losses that are not due to variations in consumption.

Note that the postulated utility function treats all premature deaths in a similar way, whatever their cause is (epidemic or not). The underlying intuition is that the damage due to a premature death lies mainly in the opportunity cost of death (i.e., all the things that would have been lived provided death had not taken place), rather than in the circumstances of death.

---

[13] The utility function deliberately abstracts from pure time preferences. Actually, the survival probability $S(I_t, Z_t)$ can be interpreted as a biological discount factor.

[14] The utility of being dead is normalized to 0.

[15] Given that this paper focuses on the trade-off between saving lives and prosperity, we assume here that the effect of the pandemic on welfare goes either through survival conditions, or through consumption possibilities, or through the pure disutility of lockdown. Adding a term for the utility of health would complicate the analysis, without adding key insights for the issue at stake.

[16] The existence of the threshold $\bar{b}$ can be justified by *reductio ad absurdum*. If that finite threshold $\bar{b} > 0$ did not exist, it would be the case either that all lives (even the most miserous ones) are worth living (when $\bar{b} \leq 0$), or that no life (even the richest ones) is worth living (when $\bar{b} \to +\infty$). None of those two cases is attractive. Thus, the existence of a finite threshold $\bar{b} > 0$ makes sense. See Fleurbaey and Ponthiere (2019).

Equilibrium prevalence of the epidemic The prevalence of the epidemic follows the dynamic law:

$$I_t = \frac{M + I_{t-1}}{1 + I_{t-1} + Z_{t-1}} \tag{8}$$

In the absence of any lockdown policy (i.e., $Z_{t-1} = 0 \ \forall t$), the prevalence of the epidemic follows the law:

$$I_t = \frac{M + I_{t-1}}{1 + I_{t-1}} \tag{9}$$

so that the equilibrium prevalence $I_t = I_{t-1} = I$ is given by: $I^* = \sqrt[2]{M} < 1$. Given that $\left|\frac{\partial I_t}{\partial I_{t-1}}\right|_{I=I^*} < 1$, the prevalence of the epidemic will, in the absence of lockdown, converge to a level equal to $\sqrt[2]{M}$. We then have: $S_t = \lim_{I \to \infty} S\left(\sqrt[2]{M}, 0\right) = \check{S}$ with $0 < \check{S} < 1$.

Let us now assume that a constant positive lockdown $Z_{t-1} = Z > 0$ takes place. The equilibrium prevalence of the epidemic is now given by:

$$I^* = \frac{-Z + \sqrt[2]{Z^2 + 4M}}{2} > 0 \tag{10}$$

The equilibrium prevalence of the epidemic I^* under a positive lockdown is smaller than in the absence of lockdown, since I^* is decreasing in the size of the lockdown Z.

Note also that the absolute value of the derivative of the prevalence dynamic equation at the equilibrium prevalence satisfies the condition:

$$\left|\frac{\partial I_t}{\partial I_{t-1}}\right|_{I=I^*} = \frac{1 + Z - M}{\left(1 + Z + \frac{-Z + \sqrt[2]{Z^2 + 4M}}{2}\right)^2} < 1$$

so that the equilibrium prevalence I^* is locally stable. This means that, for any initial prevalence I_0 and any level of lockdown Z, the prevalence of the pandemic will converge towards its equilibrium level I^*.

One can interpret the equilibrium prevalence of the pandemic I^* as a kind of "target" that the government can achieve through the selection of a particular degree of lockdown Z. Obviously, the level of the "target" selected by the government depends on the underlying social welfare criterion. The rest of the paper is dedicated to examining that policy choice.

3 Optimal lockdown under utilitarianism

Let us now consider the design of a socially optimal lockdown policy. For that purpose, this section adopts the standard utilitarian ethical criterion. Introduced by Bentham (1789) and Mill (1863), utilitarianism recommends to implement policies that satisfy the Principle of Utility, or the Principle of the Largest Happiness for the Largest Number. Under utilitarianism, social ethics is reduced to a calculus of pleasures and pains. The question of the optimal lockdown is here reduced to characterizing the degree of lockdown that leads to the largest social welfare, defined as the sum of individual utilities.

Let us consider a utilitarian social planner, who chooses the consumption per young adult c, the PAYGO benefit b, and the degree of lockdown Z in such a way as to maximize the sum of individual utilities prevailing at the steady-state equilibrium, while satisfying the economy's resource constraint, and while assuming that the prevalence of the epidemic takes its equilibrium level.[17]

The problem of the utilitarian social planner can be written as:

$$\max_{c,b,Z} \quad N\left[u(c) - vZ + S(I, Z)\left[u(b) - vZ\right]\right]$$

$$\text{s.t. } c = \frac{F(N, Z)}{N} - S(I, Z)\,b$$

$$I = \frac{-Z + \sqrt[2]{Z^2 + 4M}}{2}$$

$$Z \geq 0 \text{ and } 1 - Z \geq 0$$

The last two conditions insure that $0 \leq Z \leq 1$, that is, the lockdown is between 0 (no lockdown) and 1 (full lockdown).

The above social planning problem can be interpreted as follows. The equilibrium level of the prevalence I can be interpreted as a "target" for the prevalence of the pandemic. The social planner, by selecting a degree of lockdown Z, also selects an equilibrium level for I, and, hence, some "target" in terms of the prevalence of the pandemic. This "target" reveals how the social planner addresses the trade-off between the goals of saving lives and prosperity.

Substituting for the first two constraints, the utilitarian social planner's problem can be rewritten as the following Lagrangian:

$$\max_{b,Z} N\left[\begin{array}{l} u\left(\frac{F(N,Z)}{N} - S\left(\frac{-Z+\sqrt[2]{Z^2+4M}}{2}, Z\right)b\right) - vZ \\ +S\left(\frac{-Z+\sqrt[2]{Z^2+4M}}{2}, Z\right)[u(b) - vZ] \end{array}\right] + \lambda Z + \mu(1 - Z)$$

where λ and μ are two Lagrange multipliers.

First-order conditions (FOCs) are:

$$-S\left(\frac{-Z+\sqrt[2]{Z^2+4M}}{2}, Z\right)u'(c) + S\left(\frac{-Z+\sqrt[2]{Z^2+4M}}{2}, Z\right)u'(b) = 0 \quad (11)$$

and

$$N\left[u'(c)\left(\frac{F_Z(\cdot)}{N} - S_Z(\cdot)\,b\right) - v + S_Z(\cdot)[u(b) - vZ] - S(\cdot)\,v\right] + \lambda - \mu = 0$$
$$(12)$$

---

[17]The social planner faces an informational constraint: the planner cannot identify the persons who will have a long life and those who will die prematurely. This explains why the planner selects a unique level of consumption c at the young age, for all individuals (short-lived or long-lived). This informational constraint can be defended on the grounds of realism. Governments in the real world cannot identify citizens who will turn out to die prematurely in the future. Note, however, that the solution of the utilitarian problem would remain exactly the same if the planner could identify ex ante the persons who will be short-lived or long-lived.

where $S_Z(\cdot)$ is the derivative of $S(I(Z), Z)$ with respect to Z, as well as:[18]

$$\lambda \geq 0, Z \geq 0 \tag{13}$$
$$\mu \geq 0, 1 - Z \geq 0 \tag{14}$$

with complementary slackness.

The first FOC can be simplified to:

$$u'(c) = u'(b) \tag{15}$$

that is, it is socially optimal, from a utilitarian perspective, to smooth consumption along the life cycle. We thus have:

$$c = b = \frac{F(N, Z)}{N(1 + S(I, Z))} \tag{16}$$

Consumption possibilities are thus decreasing with the extent of the lockdown.

Let us now focus on the socially optimal lockdown level under utilitarianism. The second FOC simplifies to:

$$S_Z(\cdot)[u(b) - vZ] = v(1 + S(\cdot)) - u'(c)\left(\frac{F_Z(\cdot)}{N} - S_Z(\cdot)b\right) - \frac{\lambda}{N} + \frac{\mu}{N} \tag{17}$$

The LHS is the marginal gain in social welfare from allowing more elderly people to survive thanks to stricter lockdown, whereas the RHS is the marginal loss in social welfare due to a stronger lockdown, which includes, on the one hand, a direct utility loss from the lockdown (for the young and the old) (first term of RHS), and, on the other hand, a reduction of consumption possibilities due to the fall of output per worker and the additional pressure on resources due to a larger number of elderly persons (second term of RHS).[19]

In a nutshell, the condition for optimal lockdown Z depends on three key elements: first, the marginal impact of a higher degree of lockdown on survival conditions (as valued by individuals); second, the marginal pure disutility from the lockdown (independently from its effect on consumption possibilities); third, the effect of the lockdown on temporal utility via the reduction of consumption possibilities. The relative strengths of those three elements will determine the level of optimal lockdown from a utilitarian perspective.

Let us now examine the condition under which the optimal lockdown is zero. Lagrange multipliers take the values $\lambda > 0$ and $\mu = 0$. This case occurs when, if $Z = 0$, the marginal welfare gain from a rise in Z is lower than the marginal

---

[18] We have: $S_Z(\cdot) = S_1 \frac{dI(Z)}{dZ} + S_2$, where S_1 and S_2 denote, respectively, the partial derivatives of $S(\cdot)$ with respect to its first argument, and with respect to its second argument.

[19] Note that, in the absence of epidemic, we would have $S_Z(\cdot) = 0$, implying that the optimal lockdown would be equal to zero, since in that case the marginal welfare gains from a positive lockdown would be equal to zero, while the associated costs on the RHS would be strictly positive, so that we must have $\lambda > 0$ and $\mu = 0$, that is, $Z = 0$.

welfare loss associated to such a rise. Denoting $\hat{c} \equiv \dfrac{F(N,0)}{N\left(1+S\left(\sqrt[2]{M},0\right)\right)}$, the necessary and sufficient condition for a zero optimal lockdown level is:

$$S_Z\left(\frac{-Z+\sqrt[2]{Z^2+4M}}{2},Z\right)\bigg|_{Z=0}\left[u(\hat{c})-u'(\hat{c})\hat{c}\right] \leq v(1+\check{S})-u'(\hat{c})\left(\frac{F_Z(\cdot)}{N}\bigg|_{Z=0}\right) \quad (18)$$

Since $S_Z(\cdot) > 0$ at $Z = 0$, and assuming $u(\hat{c}) - u'(\hat{c})\hat{c} > 0$, the LHS is strictly positive. The two terms of the RHS are also strictly positive. The higher the pure marginal disutility of the lockdown is, the more likely it is that the optimal lockdown is zero. Another key determinant is the level of the marginal output loss due to a marginal rise in the degree of lockdown when $Z = 0$. If the pure marginal disutility of the lockdown is low (i.e., $v \simeq 0$), and if the output reacts little to a marginal rise in Z (i.e., $F_Z(\cdot) \simeq 0$ at $Z = 0$), the RHS is close to 0, so that the condition is not satisfied, and the optimal lockdown is strictly positive.

Using the above discussion, Proposition 1 presents the necessary and sufficient condition for a strictly positive optimal lockdown.

Proposition 1 *Consider the utilitarian social optimum. The optimal lockdown is strictly positive if and only if*

$$S_Z\left(\frac{-Z+\sqrt[2]{Z^2+4M}}{2},Z\right)\bigg|_{Z=0}\left[u\left(\hat{c}\right)-u'\left(\hat{c}\right)\hat{c}\right] > v(1+\check{S})-u'\left(\hat{c}\right)\left(\frac{F_Z(\cdot)}{N}\bigg|_{Z=0}\right)$$

Proof See above. □

The necessary and sufficient condition for a strictly positive lockdown level $Z > 0$ is relatively weak, especially when the direct disutility of the lockdown is low. The intuition behind that statement is that the survival function is likely to be extremely sensitive to a marginal rise in the extent of the lockdown when the initial lockdown level is zero. Thus, the first factor of the LHS of the condition in Proposition 1 is likely to be quite high, implying that the LHS exceeds the RHS, leading thus to a strictly positive optimal lockdown level.

It is thus possible, under mild conditions, to rationalize, on utilitarian grounds, the existence of a strictly positive lockdown in the presence of an epidemic like COVID-19.[20] Note, however, that the degree of strictness of the optimal lockdown under utilitarianism is more complex to characterize, since this optimal degree (i.e., the level of optimal Z) is the outcome of a complex calculus of pleasures and pains, where various forces are at work, as we discussed above. Having stressed this, we

---

[20]Given that this paper focuses on the selection of optimal stationary policies, utilitarianism would thus recommend here a positive stationary, i.e., permanent, lockdown. While that result may seem counterintuitive at first glance, it should be reminded that we study here a second-best setting where the lockdown is the unique policy instrument against the epidemic.

can nonetheless identify some conditions under which the optimal lockdown would, under utilitarianism, take its maximal level ($Z = 1$).

4 A COVID-19 Repugnant Conclusion

The utilitarian social criterion is a standard benchmark in social ethics and in public economics. However, when applied to an economy with a pandemic, the utilitarian criterion can have quite counterintuitive implications.

A first unattractive implication consists of its tendency to imply a variant of Parfit's 1984 Repugnant Conclusion. In a seminal work in population ethics, Parfit showed that classical utilitarianism suffers, under mild conditions, from the Repugnant Conclusion: for any large population of individuals having a low quality of life, it is always possible to find a larger population where each individual enjoys a lower quality of life, but which is regarded as better.

Within the context of a pandemic, utilitarianism can lead to some form of Repugnant Conclusion. Under some conditions, the society is, under utilitarianism, locked in a kind of COVID-19 variant of the Repugnant Conclusion: *for any (interior) lockdown with life-periods of low quality, there must be a stricter lockdown that is regarded as better, even though this reduces the quality of life-periods even more.* A utilitarian society would thus, under some conditions, be condemned to maximal lockdown leading to life periods of low quality.

In order to understand that result, let us turn back to expression Eq. 17, and let us now consider the condition under which the optimal lockdown is $Z = 1$. In that case, Lagrange multipliers in Eq. 17 take the levels $\lambda = 0$ and $\mu > 0$. Denoting $\bar{S} = S\left(\frac{-1+\sqrt[2]{1+4M}}{2}, 1\right) < 1$, and $\bar{c} \equiv \frac{F(N,1)}{N(1+\bar{S})}$, the necessary and sufficient condition for the maximal lockdown to be socially optimal is:

$$S_Z\left(\frac{-Z + \sqrt[2]{Z^2 + 4M}}{2}, Z\right)\bigg|_{Z=1} \left[u(\bar{c}) - u'(\bar{c})\bar{c} - v\right] > v\left[1 + \bar{S}\right] - u'(\bar{c})\left(\frac{F_Z(\cdot)}{N}\bigg|_{Z=1}\right) \quad (19)$$

That condition depends on how reactive the survival function $S(\cdot)$ is to a variation of Z when the lockdown is at its maximal level. If that reactivity is low (i.e., $S_Z(\cdot) \simeq 0$), the above condition is never satisfied. However, there is no reason to impose that assumption. There is no obvious reason why small deviations from maximal lockdown should have a negligible effect on survival to the old age. It is plausible that the survival function is sensitive to Z even at the maximal lockdown level.[21] If so, the above condition may be satisfied, especially in economies where the pure disutility from lockdown is low.

---

[21] Note that the parallel with the modelling of a production process may be here quite misleading. Indeed, when modelling a production function, it is usual to assume some form of Inada conditions, that is, that the marginal productivity of an input tends to be extremely low when the quantity of that input tends to be extremely large. But the modelling of a survival function is quite different: it is possible that, even at the maximal degree of lockdown, reducing slightly the degree of lockdown may lead to a substantial rise of mortality. Thus, one cannot use Inada conditions to rule out the case where maximal lockdown is optimal.

Under that condition, a level of lockdown less than the maximal lockdown is never socially optimal. There is always a gain in social welfare from shifting from Z to $Z' > Z$. Obviously, such a shift from Z to $Z' > Z$ has a cost in terms of the average welfare level prevailing in the population at a given time period. To see this, note that average welfare at time t is:

$$\frac{N\left(u(c) - vZ\right) + NS\left(\frac{-Z + \sqrt[2]{Z^2 + 4M}}{2}, Z\right)\left(u(c) - vZ\right)}{N\left(1 + S\left(\frac{-Z + \sqrt[2]{Z^2 + 4M}}{2}, Z\right)\right)} = u(c) - vZ \qquad (20)$$

which is *decreasing* in Z. Hence, when shifting from Z to $Z' > Z$, there is always a loss in average welfare at time t, but a gain in aggregate welfare.

Hence, under some conditions, utilitarianism leads to a form of Repugnant Conclusion, since this leads societies to sacrifice the quality of life periods on the grounds of promoting the quantity of life periods. Proposition 2 summarizes our results.

Proposition 2 *Consider the utilitarian social optimum. If*

$$S_Z\left(\frac{-Z + \sqrt[2]{Z^2 + 4M}}{2}, Z\right)\Bigg|_{Z=1} \left[u(\bar{c}) - u'(\bar{c})\bar{c} - v\right] > v\left[1 + \bar{S}\right] - u'(\bar{c})\left(\frac{F_Z(\cdot)}{N}\Bigg|_{Z=1}\right)$$

then for any (interior) lockdown $Z < 1$ with life periods of low quality, there must be a stricter lockdown $Z' > Z$ that is regarded as better, even though it reduces the quality of life periods even more.

Proof See above. □

Proposition 2 states a condition under which any interior lockdown is not optimal from a utilitarian perspective: for any interior lockdown with a low quality of life periods (i.e., a low average welfare at a given period), there exists a higher degree of the lockdown that is regarded as better, even though it reduces the quality of life periods even more. That result can be regarded as a COVID-19 variant of Parfit's Repugnant Conclusion.

Let us say a few words on how the result of Proposition 2 can be related to Parfit's original Repugnant Conclusion.

Parfit's Repugnant Conclusion states that, for any large population of individuals having a low utility level, it is always possible to find another, even larger, population, where each individual enjoys an even lower welfare level, but such that total welfare is larger. The reason why Parfit's result is regarded as a counterintuitive implication of classical utilitarianism lies in the fact that utilitarianism justifies a—somewhat unattractive—*substitution* between the quantity of lives and the quality of lives: the Repugnant Conclusion amounts to replace lives with a high quality by a larger number of born individuals enjoying poor lives (see Blackorby et al. (2005), Arrhenius (2013)).

Proposition 2 identifies the condition under which utilitarianism leads to a COVID-19 variant of the Repugnant Conclusion. In comparison to Parfit's result, the utilitarian criterion justifies here another type of *substitution*: a substitution between average welfare at a given time period and the quantity of life for a given number of born individuals. That form of substitution is somewhat counterintuitive: the quality of each life period is here being sacrificed on the ground of a rise in the total number of life periods. As in Parfit (1984), there is, here again, some form of substitution between quantity and quality. The difference is that Parfit's result involves a substitution between the quantity of lives and the quality of lives, whereas Proposition 2 involves a substitution between the quantity of life periods and the quality of life periods.[22]

Note that the extent to which one regards the result of Proposition 2 as being "repugnant" is a matter of ethical sensibility. As for Parfit's result, some readers might find that this result is in conflict with their basic moral intuitions. Other readers might find this conclusion not "repugnant". We do not want to impose a particular view here. We just want to highlight that utilitarianism can legitimate extreme lockdown policies on the ground of a substitution between the quantity of life periods and the quality of life periods. Whether this kind of substitution is ethically attractive or not goes beyond the scope of this paper.

5 Saving lives versus Hammond Equity

When considering issues of social justice, a major benchmark consists of the property of Hammond Equity (1979). Hammond Equity regards as socially desirable any transfer of well-being from a more advantaged person to a less advantaged person, whatever the sizes of the welfare loss for the former and of the welfare gain for the latter. In some sense, Hammond Equity consists of giving absolute priority to the worst-off in the society.

It can be shown that the utilitarian optimum leads, in the context of selecting an optimal lockdown under a pandemic, to a violation of Hammond Equity. To see this, let us first calculate the welfare levels of all individuals in the society.

The realized or ex post lifetime well-being of the short-lived is, at the utilitarian optimum, given by:

$$U^{SL} = u\left(\frac{F(N, Z)}{N(1 + S(I, Z))}\right) - vZ \tag{21}$$

The realized or ex post lifetime well-being of the long-lived is, at the utilitarian optimum, given by:

$$U^{LL} = u\left(\frac{F(N, Z)}{N(1 + S(I, Z))}\right) - vZ + u\left(\frac{F(N, Z)}{N(1 + S(I, Z))}\right) - vZ \tag{22}$$

---

[22]The policy result derived in Proposition 2 is *ethical* in nature: it does not concern a one-person choice between the quantity of life periods and the quality of life periods, but a social dilemma that involves many persons, everyone making some sacrifices in terms of quality of life periods for the sake of increasing the quantity of life periods of some individuals.

Assuming that a life is worth living, that is, that the utility of being alive at a given time period is higher than the utility of being dead (normalized to 0), that is:

$$u\left(\frac{F(N, Z)}{N(1 + S(I, Z))}\right) - vZ > 0 \tag{23}$$

we can see that, at the utilitarian optimum, persons who have a short life are worse-off than persons who enjoy a long life.

Does the utilitarian criterion satisfy Hammond Equity? The answer to that question is, in general: No. To understand why the utilitarian criterion violates Hammond Equity, let us remind that, under the condition of Proposition 1, the utilitarian optimum involves a strictly positive level of the lockdown $Z > 0$. It should be stressed, however, that the well-being of the short-lived is *decreasing* with the strength of the lockdown:

$$\frac{\partial U^{SL}}{\partial Z} = u'\left(\frac{F(N, Z)}{N(1 + S(\cdot))}\right)\frac{F_Z(N, Z)N(1 + S(\cdot)) - F(N, Z)NS_Z(\cdot)}{(N(1 + S(\cdot)))^2} - v < 0$$

Therefore, if one wants to satisfy Hammond Equity, one needs here to set the lockdown level Z to zero, in such a way as to maximize the lifetime well-being of the short-lived, U^{SL}. This is not the case under utilitarianism in general. Proposition 3 summarizes our results.

Proposition 3 *Hammond Equity requires to set the lockdown level to $Z = 0$. Under the condition of Proposition 1, the utilitarian criterion leads to a strictly positive lockdown, $Z > 0$, thus violating Hammond Equity.*

Proof See above. □

Utilitarianism violates Hammond Equity, since it fails to give priority to the worst-off ex post. Such a failure is problematic: within our economy, being long-lived or short-lived is a pure matter of luck, a pure matter of circumstances. Well-being inequalities due to unequal lifetimes are arbitrary, and the short-lived are victims of circumstances. Hence, if a government wants to be fair with respect to short-lived individuals, it should, in line with Hammond Equity, implement a zero lockdown. Utilitarianism, by recommending a strictly positive lockdown, fails to satisfy Hammond Equity.

At this stage, it is worth explaining the reason why utilitarianism implies a positive lockdown. Actually, as shown above, the average welfare at a given point in time is decreasing in the extent of the lockdown. Moreover, the well-being of the worst-off (the short-lived) is decreasing with the strength of the lockdown. Hence, why does utilitarianism lead to a positive lockdown level? The answer to that question can be found by going back to the FOC characterizing the optimal lockdown under utilitarianism:

$$S_Z(\cdot)[u(b) - vZ] = v(1 + S(\cdot)) - u'(c)\left(\frac{F_Z(\cdot)}{N} - S_Z(\cdot)b\right) - \frac{\lambda}{N} + \frac{\mu}{N}$$

The unique gain, in terms of social welfare, appears on the LHS of the condition: it comes from the fact that the strength of the lockdown strategy increases the survival probability to the old age, and, hence, leads to *saving lives*. If that motive were absent (i.e., $S_Z (\cdot) = 0$), the utilitarian criterion would not legitimate any lockdown.

In the light of all this, it appears that there exists a fundamental dilemma between, on the one hand, Hammond Equity and, on the other hand, saving lives. Under the conditions of Proposition 1, utilitarianism implements a strictly positive lockdown and, hence, violates Hammond Equity. Under the conditions of Proposition 2, utilitarianism implies the maximal lockdown, and, hence, leads to save as many lives as possible. But when the conditions of Proposition 1 are satisfied and not the ones of Proposition 2, utilitarianism implies an interior lockdown, which violates both Hammond Equity and the goal of saving lives.

As formalized by Fleurbaey and Ponthiere (2013), a basic way to present the objective of saving lives takes the form of the Survivors Number Count Axiom. That axiom states that *the certainty to save one life justifies a general rise of the preventive effort in the population, consumption levels being left unchanged.*

Does the utilitarian criterion satisfy the Survivors Number Count Axiom? To answer that question, let us consider the hypothetical choice of a lockdown level *ceteris paribus*, that is, the hypothetical choice of a lockdown level in an economy where consumption possibilities are not affected by the lockdown (i.e., $F_Z (\cdot) = 0$).[23] Under that assumption, the FOC for optimal lockdown becomes:

$$S_Z (\cdot) \left[u(c) - u'(c)c - vZ \right] = v(1 + S (\cdot)) - \frac{\lambda}{N} + \frac{\mu}{N}$$

We can see that, if

$$(S_Z (\cdot)|_{Z=1}) \left[u(\bar{c}) - u'(\bar{c})\bar{c} - v \right] > v(1 + \bar{S})$$

the utilitarian optimum would involve, in that hypothetical case, the maximal lockdown, meaning that the certainty to save one more life would definitely justify as much prevention as possible. That condition is necessary and sufficient for the utilitarian criterion to satisfy the Survivors Number Count Axiom.

Proposition 4 *The Survivors Number Count Axiom requires to set the lockdown to $Z = 1$ when the lockdown does not affect consumption possibilities. The utilitarian criterion satisfies the Survivors Number Count Axiom if and only if:*

$$(S_Z (\cdot)|_{Z=1}) \left[u(\bar{c}) - u'(\bar{c})\bar{c} - v \right] > v(1 + \bar{S})$$

Proof See above.   □

Thus, while the utilitarian criterion fails to satisfy Hammond Equity, it can, in some cases, satisfy the Survivors Number Count Axiom. Obviously, when there is no direct welfare loss due to the lockdown, the condition of Proposition 4 is necessarily

---

[23]That assumption is not realistic, but is made here only for the sake of evaluating whether the utilitarian criterion would satisfy the Survivors Number Count Axiom.

satisfied, and the utilitarian criterion satisfies the Survivors Number Count Axiom. However, in cases where there is a high disutility of lockdown, this may not be the case.

Whereas the Survivors Number Count Axiom may be regarded as a quite mild way of formalizing the goal of "saving lives," that axiom is actually incompatible with Hammond Equity, as shown in Fleurbaey and Ponthiere (2013). This contradiction between those two ethical principles is a general result, which can be expressed also in an economy with a pandemic like the one studied here.

Indeed, Hammond Equity requires to set the lockdown level to $Z = 0$ even in the hypothetical case where the lockdown does not affect consumption possibilities (i.e., when $F_Z (\cdot) = 0$). The reason is that a strictly positive lockdown level would reduce the welfare of the worst-off (the short-lived), because of the pure disutility of lockdown. But at the same time, the Survivors Number Count Axiom requires, in that hypothetical case, to set $Z = 1$ (maximal lockdown), on the ground of the certainty to save one life. Those two implications being contradictory, there exists no ethical criterion that would satisfy both Hammond Equity and the Survivors Number Count Axiom.

Proposition 5 *Assume a strictly positive disutility from lockdown $v > 0$. No ethical criteria can satisfy both Hammond Equity and the Survivors Number Count Axiom.*

Proof See above.   □

Proposition 5 points to a fundamental dilemma between two attractive social goals: on the one hand, giving priority to the worst-off and, on the other hand, saving lives. A choice is to be made between those two objectives.

The utilitarian criterion violates Hammond Equity, but can, under the conditions of Proposition 4, satisfy the Survivors Number Count Axiom. As such, the utilitarian criterion tends to give priority to the goal of saving lives. One may, on the contrary, opt for a social criterion that gives priority to the worst-off. A social objective of that kind is studied in the next section.

6 Optimal lockdown under ex post egalitarianism

Welfare inequalities due to unequal lifetimes are a pure matter of luck, and one may argue that such arbitrary inequalities should be abolished. One way to advocate for this consists of referring to the Principle of Compensation, according to which welfare inequalities due to circumstances should be abolished (see Fleurbaey and Maniquet 2004, Fleurbaey 2009).

Let us now assume an alternative social criterion, which does justice to the Principle of Compensation. Following Fleurbaey et al. (2014), we consider here the ex post egalitarian social criterion, which gives absolute priority to the interests of the worst-off ex post. Note that the social planner cannot identify the persons who will turn out to be short-lived and those who will turn out to be long-lived. As a consequence,

it is difficult for the planner to identify precisely the persons who will be the most disadvantaged ex post (i.e., once durations of life are revealed). But this informational constraint does not prevent the social planner from selecting a triplet $\{c, b, Z\}$ that maximizes the well-being of the worst-off individual ex post. Indeed, the social planner knows that *some* individuals will turn out to be short-lived, and that piece of information is sufficient to select the triplet $\{c, b, Z\}$ that will maximize the lifetime welfare of the persons—not identified—who will turn out to be the worst-off ex post.[24]

Under the ex post egalitarian criterion, the problem of the social planner is:

$$\max_{c,b,Z} \quad \min\{u(c) - vZ, u(c) - vZ + u(b) - vZ\}$$

$$\text{s.t. } c = \frac{F(N, Z)}{N} - S(I, Z) b$$

$$I = \frac{-Z + \sqrt[2]{Z^2 + 4M}}{2}$$

$$Z \geq 0 \text{ and } 1 - Z \geq 0$$

The objective function is not differentiable. However, it is possible to rewrite that planning problem as the maximization of the welfare of the short-lived subject to the egalitarian constraint specifying that the long-lived and the short-lived are equally well-off, and subject to the resource constraint and the equilibrium epidemic prevalence constraint. The problem can be rewritten as:

$$\max_{b,Z} \quad u\left(\frac{F(N, Z)}{N} - S(I, Z) b\right) - vZ$$

$$\text{s.t.} \quad u(c) - vZ = u(c) - vZ + u(b) - vZ$$

$$I = \frac{-Z + \sqrt[2]{Z^2 + 4M}}{2}$$

$$Z \geq 0 \text{ and } 1 - Z \geq 0$$

The first constraint consists of an egalitarian constraint. That constraint specifies that the long-lived persons and the short-lived persons must be equally well-off. That egalitarian condition can only be satisfied provided:

$$u(b) - vZ = 0 \iff b = u^{-1}(vZ)$$

Substituting for this and for the equilibrium prevalence, the planning problem becomes:

$$\max_Z u\left(\frac{F(N, Z)}{N} - S\left(\frac{-Z + \sqrt[2]{Z^2 + 4M}}{2}, Z\right) u^{-1}(vZ)\right) - vZ + \lambda Z + \mu(1 - Z)$$

where λ, μ are two Lagrange multipliers.

---

[24]See Fleurbaey et al. (2014) on this issue. Note that, if the ex post egalitarian planner could identify the persons who will turn out to be short-lived, the planner would differentiate young-age consumptions, and give more resources to the short-lived (unlike what the utilitarian planner would do in the same circumstances). Given the lack of realism of that case, this paper focuses on problems where the planner cannot identify the short-lived ex ante.

The FOCs are:

$$u'(c)\left[\frac{F_Z(\cdot)}{N} - S_Z(\cdot)u^{-1}(vZ) - S(\cdot)\frac{1}{u' \circ u^{-1}(vZ)}\right] - v + \lambda - \mu = 0 \quad (24)$$

as well as:

$$Z \geq 0, \lambda \geq 0 \quad (25)$$

$$1 - Z \geq 0, \mu \geq 0 \quad (26)$$

with complementary slackness.

Note that, since the first, the second, and the third term in brackets of the first FOC are all strictly negative, the above equality cannot prevail, except if $\lambda > 0$ and $\mu = 0$, which implies $Z = 0$. The optimal lockdown under the ex post egalitarian criterion is thus the absence of lockdown, that is, $Z = 0$. That result is in line with what we showed above: Hammond Equity requires a zero lockdown. As a consequence, the ex post egalitarian criterion, which gives absolute priority to the worst-off, implies also a zero lockdown.

Quite interestingly, this policy recommendation is in sharp contradiction with what the utilitarian criterion would recommend. Indeed, under the mild conditions of Proposition 1, utilitarianism implies a strictly positive lockdown level. This lockdown level may also take its maximal level under the conditions of Proposition 2. Those policy recommendations are not compatible with what the ex post egalitarian criterion implies.

This is not the only point of divergence between the two social criteria in the context of a pandemic. There is also a substantial divergence concerning the shape of consumption profiles. As we have seen in Section 3, utilitarianism recommends a flat consumption profile, with young age consumption c being exactly equal to the old-age pension benefit b. But such an equality does not hold at the ex post egalitarian optimum.

Indeed, the egalitarian constraint requires, at a zero lockdown level:

$$u(c) = u(c) + u(b) \iff b = \bar{b}$$

We thus have that the old-age pension benefit is, in the light of the egalitarian constraint, reduced to the level $\bar{b} > 0$.[25]

Consumption at the young age satisfies:

$$c = \frac{F(N,0)}{N} - S\left(\sqrt[2]{M}, 0\right)u^{-1}(0) = \frac{F(N,0)}{N} - \check{S}\bar{b} \quad (27)$$

Under general conditions, we have that $\frac{F(N,0)}{N} - \check{S}\bar{b} > \bar{b}$, which implies that the ex post egalitarian optimum involves a decreasing consumption profile with the age. That policy recommendation is in sharp contrast with what utilitarianism implies.

---

[25]The ex post egalitarian optimum involves thus a strictly positive old-age consumption equal to $\bar{b}$. The intuition is that if old-age consumption were lower than $\bar{b}$, long-lived individuals would be worse-off than short-lived persons. Alternatively, if $b > \bar{b}$, the short-lived would be worse-off than the long-lived.

The intuition behind that policy recommendations goes as follows. The government has here, as an objective, to maximize the well-being of the worst-off ex post. Those worst-off individuals are, under mild conditions (see above), the persons who die prematurely. Those persons cannot be identified ex ante. Moreover, once individual durations of life are revealed, it is generally too late to compensate the disadvantaged (because he is already dead). However, despite those difficulties, it is possible for a government to improve the situation of the worst-off, by transferring resources at young ages of life, in such a way as to increase the consumption of all young individuals. By proceeding in that way, the government will also, indirectly, improve the lives of the persons who will turn out to be short-lived. Hence, consumption profiles decreasing with the age are part of a strategy in order to improve the situation of the worst-off (see Fleurbaey et al. (2014)). Proposition 6 summarizes our results.

Proposition 6 *The ex post egalitarian social optimum involves zero lockdown $Z = 0$, as well as a decreasing consumption profile with the age: $c > b$.*

Proof See above.   □

Proposition 6 points to two major reasons why the Great Lockdown went against the interests of the worst-off.

First, giving priority to the worst-off recommends to impose zero lockdown, contrary to the policies carried out in spring 2020. The Great Lockdown aimed at saving lives, but deteriorated the living standards of young individuals, some of whom will turn out to be short-lived (because of the epidemic or because of other causes).

Second, giving priority to the worst-off requires to impose decreasing consumption profiles with age. However, the large economic crisis that followed the Great Lockdown had, as a major consequence, to reduce incomes and consumption possibilities *especially for young adults*. This goes against the interests of worst-off individuals, since reducing consumption possibilities at the young age also deteriorates the situation of individuals who will turn out to be short-lived.

7 Optimal age-differentiated lockdown

Up to now, we characterized optimal lockdown under various social objectives, while assuming that the lockdown must be global, and cannot be differentiated across age groups. One might criticize this approach, and argue that our analysis assumed a kind of second-best setting. Actually, given that the elderly are most vulnerable to the COVID-19, and given that the cost of the lockdown consists mainly of the economic losses due to the reduction of production, there is a strong support for a *targeted* lockdown that would focus only on the old. That point was made by Acemoglu et al. (2020).

This section aims at examining to what extent our previous results are robust to allowing for age-differentiated lockdown. Throughout this section, we assume that

the government facing the pandemic can choose different degrees of lockdown for the young and for the old. Let us denote by Z_t^y the lockdown for young adults, and by Z_t^o the lockdown for the old.

The survival conditions faced by the elderly depend on the prevalence of the pandemic and on age-specific lockdown levels Z_t^o and Z_t^y:

$$S_t = S\left(I_t, Z_t^y, Z_t^o\right) \tag{28}$$

First-order derivatives of $S\left(I, Z^y, Z^o\right)$ satisfy: $S_I < 0$, $S_{Z^y} > 0$ when $I_t > 0$ and $Z_t^o < 1$ as well as $S_{Z^y} = 0$ when $I_t = 0$ and/or when $Z_t^o = 1$.[26] We also have: $S_{Z^o} > 0$ when $I_t > 0$ as well as $S_{Z^o} = 0$ when $I_t = 0$.

The prevalence of the epidemic follows the dynamic law:

$$I_t = \frac{M + I_{t-1}}{1 + I_{t-1} + \alpha Z_{t-1}^y + (1 - \alpha) Z_{t-1}^o} \tag{29}$$

where $0 < M < 1$ and where $\alpha \in [0, 1]$ measures the strength of the (marginal) impact of the lockdown of the young on the prevalence of the epidemic, while $1 - \alpha \in [0, 1]$ measures the strength of the (marginal) impact of the lockdown of the old on the prevalence of the epidemic. In case of non-differentiated lockdown (i.e., $Z_{t-1}^y = Z_{t-1}^o$), that formula vanishes to the one studied above.

The equilibrium prevalence is:

$$I^* = \frac{-(\alpha Z^y + (1 - \alpha) Z^o) + \sqrt[2]{(\alpha Z^y + (1 - \alpha) Z^o)^2 + 4M}}{2} > 0 \tag{30}$$

Note also that:

$$\left|\frac{\partial I_t}{\partial I_{t-1}}\right|_{I=I^*} = \frac{1 + \alpha Z^y + (1 - \alpha) Z^o - M}{\left(1 + \alpha Z^y + (1 - \alpha) Z^o + \frac{-Z + \sqrt[2]{(\alpha Z^y + (1-\alpha) Z^o)^2 + 4M}}{2}\right)^2} < 1$$

so that the equilibrium prevalence is locally stable.

The production process is only affected by the lockdown faced by young adults (i.e., workers):

$$Y_t = F\left(N, Z_t^y\right) \tag{31}$$

where $F_N > 0$, $F_{NN} < 0$ (decreasing marginal productivity of labor), as well as $F_{Z^y} < 0$.

---

[26]The intuition is that the degree of lockdown for the young does not affect survival conditions of the elderly when there is no pandemic, or when the elderly are under maximal lockdown (i.e., then completely isolated from the remaining of the population).

The utilitarian problem The problem of the utilitarian social planner can be written as:

$$\max_{c,b,Z^y,Z^o} N\left[u(c) - vZ^y + S\left(I, Z^y, Z^o\right)\left[u(b) - vZ^o\right]\right]$$

$$\text{s.t. } c = \frac{F\left(N, Z^y\right)}{N} - S\left(I, Z^y, Z^o\right)b$$

$$I = \frac{-(\alpha Z^y + (1-\alpha)Z^o) + \sqrt[2]{(\alpha Z^y + (1-\alpha)Z^o)^2 + 4M}}{2}$$

$$Z^y \geq 0 \text{ and } 1 - Z^y \geq 0$$

$$Z^o \geq 0 \text{ and } 1 - Z^o \geq 0$$

Substituting for the first two constraints, the utilitarian social planner's problem can be rewritten as the following Lagrangian:

$$\max_{b,Z^y,Z^o} N\left[\begin{array}{l} u\left(\frac{F(N,Z^y)}{N} - S\left(\frac{-(\alpha Z^y+(1-\alpha)Z^o)+\sqrt[2]{(\alpha Z^y+(1-\alpha)Z^o)^2+4M}}{2}, Z^y, Z^o\right)b\right) - vZ^y \\ + S\left(\frac{-(\alpha Z^y+(1-\alpha)Z^o)+\sqrt[2]{(\alpha Z^y+(1-\alpha)Z^o)^2+4M}}{2}, Z^y, Z^o\right)\left[u(b) - vZ^o\right] \end{array}\right]$$

$$+\lambda Z^y + \mu(1 - Z^y) + \varsigma Z^o + \chi(1 - Z^o)$$

where λ, μ, ς, and χ are Lagrange multipliers.
First-order conditions (FOCs) are:

$$u'(c) = u'(b) \tag{32}$$

and

$$N\left[u'(c)\left(\frac{F_{Z^y}(\cdot)}{N} - S_{Z^y}(\cdot)b\right) - v + S_{Z^y}(\cdot)\left[u(b) - vZ^y\right]\right] + \lambda - \mu = 0 \tag{33}$$

and

$$N\left[u'(c)\left(-S_{Z^o}(\cdot)b\right) + S_{Z^o}(\cdot)\left[u(b) - vZ^y\right] - S(\cdot)v\right] + \varsigma - \chi = 0 \tag{34}$$

where $S_{Z^i}(\cdot)$ denotes the derivative of $S(\cdot)$ with respect to Z^i, as well as:

$$\lambda \geq 0, Z^y \geq 0 \tag{35}$$

$$\mu \geq 0, 1 - Z^y \geq 0 \tag{36}$$

$$\varsigma \geq 0, Z^o \geq 0 \tag{37}$$

$$\chi \geq 0, 1 - Z^o \geq 0 \tag{38}$$

with complementary slackness.

The first FOC implies that it is socially optimal, from a utilitarian perspective, to smooth consumption along the life cycle. We thus have: $c = b = \frac{F(N,Z^y)}{N(1+S(I,Z^y,Z^o))}$.

Let us now focus on the socially optimal lockdown levels $\{Z^y, Z^o\}$ under utilitarianism. The second and third FOCs simplify to:

$$S_{Z^y}(\cdot)\left[u(b) - vZ^y\right] + \frac{\lambda}{N} - \frac{\mu}{N} = -u'(c)\left(\frac{F_{Z^y}(\cdot)}{N} - S_{Z^y}(\cdot)b\right) + v \tag{39}$$

and

$$S_{Z^o}(\cdot)\left[u(b) - vZ^y\right] + \frac{\varsigma}{N} - \frac{\chi}{N} = -u'(c)\left(-S_{Z^o}(\cdot)b\right) + S(\cdot)v \qquad (40)$$

Those conditions are studied in the Appendix. Proposition 7 summarizes our results.

Proposition 7 *Assume:*

$$\lim_{Z^o \to 0} S_{Z^o}(\cdot) > S_{Z^y}(\cdot)\,\forall Z^y \in [0, 1] \ \text{ and } \ \lim_{Z^y \to 1} S_{Z^y}(\cdot) < S_{Z^o}(\cdot)\,\forall Z^o \in [0, 1]$$

At the utilitarian optimum under age-differentiated lockdown strategies, only four cases are possible:

- *No lockdown for the young and for the old: $Z^y = 0$, $Z^o = 0$,*
- *No lockdown for the young and interior lockdown for the old: $Z^y = 0, 0 < Z^o < 1$,*
- *No lockdown for the young and maximal lockdown for the old: $Z^y = 0$, $Z^o = 1$,*
- *Interior lockdown for the young and for the old: $0 < Z^y < 1, 0 < Z^o < 1$.*

Proof See the Appendix. □

Proposition 7 states that, under general conditions, the utilitarian optimum excludes maximal lockdown for the young, but may involve maximal lockdown for the old. That result is in sharp contrast with the COVID-19 Repugnant Conclusion of Section 4 and, as such, requires some explanation.

When there is no possibility to apply different lockdown policies to different age groups, it can be optimal, under a large set of conditions, to apply maximal lockdown from a utilitarian perspective (Proposition 2). The underlying intuition is that, in that case, the trade-off between prosperity and saving lives is harsh, and the goal of maximizing social welfare, which requires to save many lives, may, under some conditions, justify to apply maximal lockdown.

However, once the available policies include the possibility to differentiate the degree of lockdown across age groups, the social planner benefits from a larger set of strategies, and the trade-off between prosperity and saving lives takes a different form. It is now possible, by applying age-differentiated lockdown policies, to save many lives by isolating the elderly, without giving up too much prosperity. One can thus now achieve the goal of saving lives without having to pay a too large cost in terms of giving up output. This cheaper way of saving lives makes the maximal lockdown for the young adult socially suboptimal. Given the new form of the trade-off between prosperity and saving lives, it cannot be socially optimal to have maximal lockdown for the young, since this strategy is dominated by the lockdown of the elderly, which saves (marginally) more lives and at a lower cost in terms of prosperity.

The ex post egalitarian optimum Under the ex post egalitarian social objective, the problem of the social planner becomes:

$$\max_{c,b,Z^y,Z^o} \quad \min\left\{u(c) - vZ^y, u(c) - vZ^y + u(b) - vZ^o\right\}$$

$$\text{s.t. } c = \frac{F(N, Z^y)}{N} - S\left(I, Z^y, Z^o\right)b$$

$$I = \frac{-(\alpha Z^y + (1-\alpha)Z^o) + \sqrt[2]{(\alpha Z^y + (1-\alpha)Z^o)^2 + 4M}}{2}$$

$$Z^y \geq 0 \text{ and } 1 - Z^y \geq 0$$

$$Z^o \geq 0 \text{ and } 1 - Z^o \geq 0$$

After some substitutions, the problem can be rewritten as follows:

$$\max_{b,Z^y,Z^o} \quad u\left(\frac{F(N, Z^y)}{N} - S\left(I, Z^y, Z^o\right)b\right) - vZ^y$$

$$\text{s.t.} \quad u(c) - vZ^y = u(c) - vZ^y + u(b) - vZ^o$$

$$I = \frac{-(\alpha Z^y + (1-\alpha)Z^o) + \sqrt[2]{(\alpha Z^y + (1-\alpha)Z^o)^2 + 4M}}{2}$$

$$Z^y \geq 0 \text{ and } 1 - Z^y \geq 0$$

$$Z^o \geq 0 \text{ and } 1 - Z^o \geq 0$$

The first constraint consists of an egalitarian constraint. That egalitarian condition can only be satisfied provided:

$$u(b) - vZ^o = 0 \iff b = u^{-1}\left(vZ^o\right)$$

Substituting for this, the planning problem becomes:

$$\max_{Z^y,Z^o} \quad u\left(\frac{F(N, Z^y)}{N} - S\left(\frac{-(\alpha Z^y + (1-\alpha)Z^o) + \sqrt[2]{(\alpha Z^y + (1-\alpha)Z^o)^2 + 4M}}{2}, Z^y, Z^o\right)u^{-1}\left(vZ^o\right)\right)$$

$$-vZ^y + \lambda Z^y + \mu(1 - Z^y) + \varsigma Z^o + \chi(1 - Z^o)$$

where $\lambda, \mu, \varsigma, \chi$ are Lagrange multipliers.
The FOCs are:

$$u'(c)\left[\frac{F_{Z^y}(\cdot)}{N} - S_{Z^y}(\cdot)u^{-1}\left(vZ^o\right)\right] - v + \lambda - \mu = 0 \qquad (41)$$

$$u'(c)\left[-S_{Z^o}(\cdot)u^{-1}\left(vZ^o\right) - S(\cdot)\frac{1}{u' \circ u^{-1}\left(vZ^o\right)}\right] + \varsigma - \chi = 0 \qquad (42)$$

as well as:

$$Z^y \geq 0, \lambda \geq 0 \tag{43}$$

$$1 - Z^y \geq 0, \mu \geq 0 \tag{44}$$

$$Z^o \geq 0, \varsigma \geq 0 \tag{45}$$

$$1 - Z^o \geq 0, \chi \geq 0 \tag{46}$$

with complementary slackness.

The analysis of the FOCs is provided in the Appendix. Proposition 8 summarizes our results.

Proposition 8 *The ex post egalitarian optimum under age-differentiated lockdown strategies involves zero lockdown for the young and for the old: $Z^y = 0$, $Z^o = 0$.*

Proof See the Appendix. □

Proposition 8 states that the *ex post* egalitarian optimum involves zero lockdown for the young and for the old. Thus, the possibility to differentiate the degree of lockdown with the age does not influence the result that we already had in the baseline model (Proposition 6): giving absolute priority to the worst-off individual *ex post* implies zero lockdown.

The intuition behind that result goes as follows. Take first the lockdown of young adults. That lockdown is a source of direct disutility for the young (through the disutility inherent to restricting liberties), and, also, reduces consumption possibilities for the young through two distinct channels: reduction of output and rise of the burden of the elderly's consumption (because of the associated rise of S). Thus, from the perspective of individuals who will turn out to be short-lived (and thus enjoy only the young age), zero lockdown is optimal for the young. Consider now the lockdown for the old. This reduces consumption possibilities for the young through two channels: first, this increases the burden of the old's consumption, and, second, positive lockdown for the elderly requires to increase the old's consumption (via the egalitarian constraint), which reduces consumption possibilities of the young. Hence, zero lockdown for the old is also optimal from an ex post egalitarian perspective.

Taken together, Propositions 7 and 8 illustrate the robustness of the main result of this paper: the optimal lockdown is strongly sensitive to the postulated social objective. This was true when considering the design of optimal (undifferentiated) lockdown, and this remains true once we consider the choice of age-differentiated lockdown.

8 Concluding remarks

This paper proposed an exploration of the normative foundations for lockdown policies in the context of a pandemic. Our main result consists of the lack of robustness of the optimal lockdown to the postulated ethical criterion. Utilitarianism recommends a strictly positive lockdown level, and may also, under some conditions, lead

to a COVID-19 variant of Parfit's (1984) Repugnant Conclusion: for any (interior) lockdown with life periods of low quality, there must be a stricter lockdown that is regarded as better, even though this reduces the quality of life periods even more. On the contrary, the ex post egalitarian criterion, which gives priority to the worst-off ex post, implies a zero lockdown. The ex post egalitarian optimum involves also decreasing consumption profiles with age. This last result points to a strategy against the epidemic that has not been explored so far: modifying consumption profiles so as to improve the situation of the young, some of whom will turn out to die prematurely.

Quite interestingly, our exploration also reveals that there exists a general dilemma between the goal of giving priority to the worst-off (Hammond Equity) and the objective of saving lives, modeled by the Survivors Number Count Axiom. The Survivors Number Count Axiom would imply, in the hypothetical case where the lockdown would leave the production process unaffected, the maximal lockdown, whereas Hammond Equity would, in that case, recommend zero lockdown. Thus, there exists a fundamental dilemma between saving lives and giving priority to the worst-off.

All in all, this paper highlights that the weight assigned to the worst-off determines the extent to which a lockdown is socially desirable, and, if yes, how strict it should be. This general result holds whatever we consider the choice of (undifferentiated) lockdown, or the selection of age-differentiated degrees of lockdown. This paper, by highlighting the lack of robustness of optimal lockdown to the ethical criterion, underlines that the design of lockdown policies is a matter of social choice. Normative foundations behind those policies are far from obvious, and these should be more discussed within the public debate.

Finally, it should be stressed that, although the paper focuses on a second-best setting where the lockdown is the only available policy against the epidemic, some of our results could be used to study the normative foundations of other policies, such as vaccination. Introducing vaccination as another policy instrument could, under some conditions, crowd-out lockdown policies. However, some of our results would still remain relevant. In particular, the conflict between the goal of saving lives (Survivors Number Count) and the goal of giving priority to the worst-offs (Hammond Equity) is a general conflict, which concerns all preventive policies that save lives at the cost of decreasing the welfare of the worst-offs. Those two goals being attractive (at least for progressive governments), we face here a general ethical dilemma, which concerns all societies, not only in the context of COVID-19, but, also, beyond that pandemic.

Appendix

Proof of Proposition 7 Let us rewrite the conditions for optimal Z^y and Z^o as follows:

$$S_{Z^y}(\cdot)\left[u(b) - u'(b)b - vZ^y\right] + \frac{\lambda}{N} - \frac{\mu}{N} = -u'(c)\left(\frac{F_{Z^y}(\cdot)}{N}\right) + v \qquad (47)$$

and

$$S_{Z^o}(\cdot)\left[u(b) - u'(b)b - vZ^y\right] + \frac{\varsigma}{N} - \frac{\chi}{N} = S(\cdot)v \qquad (48)$$

9 cases are a priori possible:

1. $Z^y = 0, Z^o = 0 : \lambda = \varsigma > 0, \mu = \chi = 0.$
2. $Z^y = 0, 0 < Z^o < 1 : \lambda > 0, \mu = 0, \varsigma = \chi = 0$
3. $Z^y = 0, Z^o = 1 : \lambda > 0, \mu = 0, \varsigma = 0, \chi > 0$
4. $0 < Z^y < 1, Z^o = 0 : \lambda = \mu = 0, \varsigma > 0, \chi = 0.$
5. $0 < Z^y < 1, 0 < Z^o < 1 : \lambda = \mu = \varsigma = \chi = 0.$
6. $0 < Z^y < 1, Z^o = 1 : \lambda = \mu = 0, \varsigma = 0, \chi > 0.$
7. $Z^y = 1, Z^o = 0 : \lambda = 0, \mu > 0, \varsigma > 0, \chi = 0.$
8. $Z^y = 1, 0 < Z^o < 1 : \lambda = 0, \mu > 0, \varsigma = \chi = 0.$
9. $Z^y = 1, Z^o = 1 : \lambda = \varsigma = 0, \mu, \chi > 0.$

The RHS of the FOC for Z^y is larger than the RHS of the FOC for Z^o. Hence, the analysis of the feasibility of the different cases requires to check the conditions under which the LHS of the FOC for Z^y is also larger than the LHS of the FOC for Z^o.

Assume $Z^o = 1$. We know that $S_{Z^y}(\cdot) = 0$ when $Z^o = 1$. Hence, cases 6 and 9 are excluded.

Moreover, under $\lim_{Z^o \to 0} S_{Z^o}(\cdot) > S_{Z^y}(\cdot) \forall Z^y \in [0, 1]$, cases 4 and 7 are excluded.

Finally, under $\lim_{Z^y \to 1} S_{Z^y}(\cdot) < S_{Z^o}(\cdot) \forall Z^o \in [0, 1]$, case 8 is excluded.

As a consequence, the only possible cases are cases 1, 2, 3, and 5. □

Proof of Proposition 8 The FOCs for optimal Z^y and Z^o can be rewritten as:

$$u'(c) \left[\frac{F_{Z^y}(\cdot)}{N} - S_{Z^y}(\cdot) u^{-1}(vZ^o) \right] = -\lambda + \mu + v \tag{49}$$

$$u'(c) \left[-S_{Z^o}(\cdot) u^{-1}(vZ^o) - S(\cdot) \frac{1}{u' \circ u^{-1}(vZ^o)} \right] = -\varsigma + \chi \tag{50}$$

9 cases are a priori possible:

1. $Z^y = 0, Z^o = 0 : \lambda = \varsigma > 0, \mu = \chi = 0.$
2. $Z^y = 0, 0 < Z^o < 1 : \lambda > 0, \mu = 0, \varsigma = \chi = 0$
3. $Z^y = 0, Z^o = 1 : \lambda > 0, \mu = 0, \varsigma = 0, \chi > 0$
4. $0 < Z^y < 1, Z^o = 0 : \lambda = \mu = 0, \varsigma > 0, \chi = 0.$
5. $0 < Z^y < 1, 0 < Z^o < 1 : \lambda = \mu = \varsigma = \chi = 0.$
6. $0 < Z^y < 1, Z^o = 1 : \lambda = \mu = 0, \varsigma = 0, \chi > 0.$
7. $Z^y = 1, Z^o = 0 : \lambda = 0, \mu > 0, \varsigma > 0, \chi = 0.$
8. $Z^y = 1, 0 < Z^o < 1 : \lambda = 0, \mu > 0, \varsigma = \chi = 0.$
9. $Z^y = 1, Z^o = 1 : \lambda = \varsigma = 0, \mu, \chi > 0.$

Take the first condition. Its LHS is unambiguously negative. Hence, the equality can only hold provided the RHS is also negative. This excludes cases 4, 5, 6, 7, 8, and 9.

Take the second condition. Its LHS is negative. Hence, the equality can only hold provided the RHS is also negative. This excludes cases 2 and 3.

Hence, only one case remain: zero lockdown for the young and for the old. □

Declarations

Conflict of interest The authors declare no competing interests.

References

Acemoglu D, Chernozhukov V, Werning I, Whinston M (2020) A multi-risk SIR model with optimally targeted lockdown. NBER Working Paper 27102

Alvarez FE, Argente D, Lippi F (2020) A simple planning problem for COVID-19 lockdown. NBER Working Paper 26981

Arrhenius G (2013) The Repugnant conclusion. In: La Follette H (ed) International encyclopedia of ethics. Blackwell, Oxford

Atkeson A (2020) What will be the economic impact of COVID-19 in the US? Rough estimates of disease scenarios. NBER Working Paper 26867

Becker G, Philipson T, Soares R (2005) The quantity and quality of life and the evolution of world inequality. Am Econ Rev 95:277–291

Bentham J (1789) An Introduction into the principles of morals and legislation. London

Blackorby C, Bossert W, Donaldson D (2005) Population issues in social choice theory, welfare economics and ethics. Cambridge University Press, Cambridge

Bosi S, Camacho C, Desmarchelier D (2021) Optimal lockdown in altruistic economies. J Math Econ 93:102488

Deb P, Furceri D, Ostry J, Tawk N (2020) How the Great Lockdown saved lives. IMF Blog June 2020

Eichenbaum M, Rebelo S, Trabandt M (2020) The macroeconomics of epidemics. NBER Working Paper 26882

Euronews (2020) Coronavirus: half of humanity now on lockdown as 90 countries call for confinement. Article 04/03/2020

Fleurbaey M (2009) Fairness, responsibility and welfare. Oxford University Press, Oxford

Fleurbaey M, Leroux ML, Ponthiere G (2014) Compensating the dead. J Math Econ 51:28–41

Fleurbaey M, Maniquet F (2004). In: Arrow K, Sen A, Suzumura K (eds) Compensation and responsibility. Handbook of social choice and welfare, North Holland

Fleurbaey M, Ponthiere G (2013) Prevention against equality? J Public Econ 103:68–84

Fleurbaey M, Ponthiere G (2019) Measuring well-being and lives worth living. PSE Discussion Paper 2019–65

Garriga C, Manuelli R, Sanghi S (2020) Optimal management of an epidemic: an application to COVID-19. Mimeo

Gollier C (2020) Cost-benefit analysis of age-specific deconfinement strategies. Mimeo

Hammond P (1979) Equity in two-person situations. Econometrica 47:1127–1136

Lee R (2003) The demographic transition: three centuries of fundamental change. J Econ Perspect 17:167–190

Mill JS (1863) Utilitarianism. New Edition by Mary Warnock, London

Moore GE (1903) Principia Ethica. Cambridge University Press, Cambridge

OECD (2020) Evaluating the initial impact of COVID-19 containment measures on economic activity. OECD, June 2020

Parfit D (1984) Reasons and persons. Oxford University Press, Oxford

Piguillem F, Shi L (2020) Optimal COVID-19 quarantine and testing policies. CEPR Discussion Papers 14613

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.