

Utilitarianism on the front lines: COVID-19, public ethics, and the “hidden assumption” problem

Charles Shaw¹ & Silvio Vanadia²

Abstract

How should we think of the preferences of citizens? Whereas self-optimal policy is relatively straightforward to produce, socially optimal policy often requires a more detailed examination. In this paper, we identify an issue that has received far too little attention in welfarist modelling of public policy, which we name the “hidden assumptions” problem. Hidden assumptions can be deceptive because they are not expressed explicitly and the social planner (e.g. a policy maker, a regulator, a legislator) may not give them the critical attention they need. We argue that ethical expertise has a direct role to play in public discourse because it is hard to adopt a position on major issues like public health policy or healthcare prioritisation without making contentious assumptions about population ethics. We then postulate that ethicists are best situated to critically evaluate these hidden assumptions, and can therefore play a vital role in public policy debates.

Keywords: utilitarianism, COVID-19, ethical analysis, health resources, health care rationing, public health, public ethics

Introduction

The COVID-19 outbreak, which has resulted in restrictions on transportation, work, and ordinary life, has affected billions of people and cost the global economy trillions of dollars as a result of a possible hazard to many lives. It has also stimulated wide-ranging public debate over complex social dilemmas. This is not surprising because in order for a society to make choices, we need a set of rules or principles to follow. Medical, governmental, and political choices must be made, understood, and applied. We expect decisions which affect the public to be grounded in presuppositions which have been consciously chosen and critically assessed, to a robust philosophical standard. Can we rely on ethics, especially population ethics, to evaluate the merits and demerits of a government’s response to a pandemic? Whilst moral and ethical theories are occasionally referenced, the contribution of ethicists to public policy in times of crisis, including the COVID-19 pandemic, has been unclear.

The aim of this article is twofold. Motivated by the Italian experience, our primary goal is to discuss the approach that national authorities have taken, as well as the specific value frameworks that they have used or proposed to establish the priority of treatment of patients in intensive care. We examine the ethical issues emerging in the context of a public crisis, such as the COVID-19 pandemic.

Our secondary goal is to examine decision-making behind resource allocation and some of the ethical difficulties that society and healthcare systems face in the context of a public emergency. In regular circumstances, it is the job of hospital ethics committees to assist in the resolution of ethical quandaries in the setting of inpatient treatment. Clinicians are accustomed to similar quandaries: setting care objectives, evaluating patients’ decision-making competence, handling treatment refusals, and so on. Ordinarily, clinicians would address these quandaries with a patient-centered approach, but these are not ordinary circumstances.

¹ Independent economist from London (United Kingdom); charles@fixedpoint.io; ORCID: 0000-0002-3287-8251

² Independent philosopher from Milan (Italy); ORCID: 0000-0002-9808-5836

It is not our aim to provide a framework for the resolution of extant moral quandaries or even to provide a tentative roadmap for doing so. This is not least because the crisis is still ongoing. Instead, we aim to shed light on some of the practical challenges regarding triage ethics, resource allocation, and medico-ethical decision-making under crisis.

Discussions on social issues are frequently related to policy options, which necessitate evaluation. What does it mean to state that one policy option is superior than another one? Principles must be able to balance gains and losses because policy changes seldom result in benefits for everyone. Some policies result in a greater population that is less well-off than others. As a result, there is an implicit ethical dimension to policy decisions. Welfarist evaluation principles are the primary focus of our analysis. Welfarism is founded on the premise that if two policy scenarios (with the same population) result in everyone getting the same amount of reward, they are both equally beneficial (Sen, 1987).

In accordance with welfarist ideas, values like individual liberty and autonomy are instrumental and desirable because they contribute to well-being. As well as this, virtues and fair practices may be useful. Since this is the case, a holistic approach to well-being is essential. For example, it seems appropriate to focus on lifetime well-being and include enjoyment; pleasure and the absence of pain; autonomy; liberty; and good health. We presume that people who are fully informed and autonomous have self-regarding preferences that align with their well-being.

Where should we begin our investigation of the moral underpinnings of lockdown rules in the case of a pandemic? It is well understood that COVID-19 is an example of a pandemic that has had a significant impact on global welfare and economic prosperity. On the other hand, there is substantial disagreement on whether there exists an optimal policy approach when social interaction is crucial to both illness transmission and economic activity. In the face of excessive demand, health systems have attempted to develop rules in this vein for which patients should be treated. Enquiring minds may want to know the extent to which broad and intuitive analysis of social policy decisions are made as a result of such rules and the corresponding outcomes. Whereas self-optimal policy is relatively straightforward to produce, socially optimal policy, which is arguably of more practical importance, requires a more detailed examination.

The remainder of the paper is organised as follows. Section 2 briefly outlines some of the standard philosophical and ethical frameworks. Section 3 demonstrates that amidst the public discussion over pandemic response, one ethical framework – utilitarianism – has been both referenced and criticised. Section 4 introduces a so-called axiomatic approach, which consists of formulating and solving problems related to social choice, as well as analysing their underlying axioms. Section 5 describes how, in order to account for the feedback between the status of the pandemic and people's behaviour, formal (mostly game-theoretic) models of social distance have gained ascendancy. Section 6 examines the question of assumptions that appear to be hidden or implicit in models used to guide epidemiological, economic, or public health related decision making. Section 7 concludes and provides tentative recommendations for future research.

Some standard philosophical approaches

Before we begin our analysis, we will briefly outline some of the standard ethical frameworks as follows.

- “Act utilitarianism” demands that each of our individual decisions maximise quantifiable good. By imposing limits (ostensibly bad), governments have attempted to prevent the spread of the pandemic (good), enhance public health (good), and save lives (good), maybe

at the price of other lives (bad). It will only be possible to evaluate the strategy and timing of 'applying the brakes' once the pandemic is over. Although empirical calculations may favour isolation in retrospect, the likely sacrifice of some innocent lives will nevertheless leave some moral concerns unaddressed.

- “Rule utilitarianism” advises us to devise rules to act on in specific instances. The use of Quality Adjusted Life Years (QALYs) in decision making is one such concept. If individuals who die are extremely elderly, they will not have as great of an impact on the national QALY aggregate as the young and healthy.
- “Kantian ethics” dictates that one must never treat other human beings as means to an end but only as ends in and of themselves. Did, for example, the Italian government follow this rule? This can be argued in either direction. Yes, since the lives were not sacrificed in one given (possibly isolated) region for the sake of health advantages elsewhere. No, since the weak were placed in danger to help others. The Kantian formula leaves considerable room for interpretation.
- “Natural law” bans us from violating or disregarding our fundamental human rights, which are, in the traditional formulation, survival, health, shelter, having and rearing children, and pursuing knowledge, particularly about God. This method easily supports the government’s determination to safeguard lives and health. However, since the same family of ideas gives rise to the theory of twofold impact for exceptions, it is not always apparent how the principles should be implemented.
- “Moral legalism” is not an ethical system in and of itself, but it influences people’s thinking in several regions of the globe, including Italy. To be moral, we merely need to follow the rules of the land. This is consistent with the legal positivist jurisprudential credo, which simply holds that the law is the law because it is the law, and that no further explanation is conceivable or desired. The argument, on the other hand, is incompatible with natural law theory, which holds that morality is superior to current law and that certain laws are harmful and should not be observed.
- “Virtue ethics”, with roots in its Aristotelian formulation, instructs us to find a balance between doing too much or feeling too strongly and doing too little or feeling too weakly. This concept works particularly well in hazardous situations: instead of being cowardly (e.g. too little action) or rash (e.g. too much action) in response to fear, we should follow the golden mean and be courageous (appropriate action). It does not, however, give precise guidance for making policy decisions.

Now that we have outlined some of the standard approaches relevant to the question at hand, let us continue with a typical challenge of how to distribute limited resources. Successfully addressing this problem can be achieved in a variety of ways, such as advocating for the distribution of the items in question to those who need them, deserve them, or can afford them. In our setting, heterogeneity in choice sets may result in fundamentally different types of state governance. There is temporal variation of choice sets to consider too, since the impacts of a choice on resource allocation may be long-lasting or subject to hysteresis. In light of these considerations, any solution to a given distributive justice issue ought to be welcomed. However, we will argue that this is a naive view. We will show that in the face of such complex societal issues, ethics alone (and even ethics in general) may struggle to provide definitive solutions. On the other hand, ethics has a direct role to play in public discourse because it is hard to adopt a position on major issues like public health policy or healthcare prioritisation without making contentious assumptions about population welfare.

“Ruthless” vs “soft” utilitarianism

Health systems have a responsibility to manage finite resources during a pandemic (Emanuel et al., 2020). This guideline, widely accepted in countries such as the United States (New York State Department Of Health, 2015) and the United Kingdom (Royal College Of Physicians, 2020), mandates the allocation of resources in order to “maximise the number of patients who survive therapy with a realistic life expectancy.” Ethical dilemmas might arise in the process of obtaining such goals. For example, it is “justifiable” to remove a patient from an ICU bed or a ventilator “since maximising benefits is crucial,” according to experts. “Patients who are fairly regarded to have the ability to benefit fast should be given access to ventilators,” the British Medical Association recommends (British Medical Association, 2020).

Should the optimal strategy, if one exists, be to reduce infections until a vaccine and a treatment are available, or is herd immunity preferable? When a pandemic shows indications of abating, how quickly should the limits on economic activity be lifted, either organically or as a result of a vaccination campaign? Whilst there is considerable uncertainty in both the answer to this question and the factors that affect this response, moral arguments continue to be shaped by the pandemic. Many of them centre on a more holistic view of ethics, which emphasises justice and care for the vulnerable. Amidst public discussion regarding the response to the pandemic, one ethical framework – utilitarianism – has been both referenced and criticised.

In March 2020, the civil rights office of the U.S. Department of Health and Human Services stated that:

[P]ersons... should not be put at the end of the line for health services during emergencies. Our civil rights laws protect the equal dignity of every human life from ruthless utilitarianism...HHS is committed to leaving no one behind during an emergency, and this guidance is designed to help health care providers meet that goal.

Roger Severino, Office of Civil Rights Director, U.S. Department of Health and Human Services also outlined that healthcare providers that apply triage protocols may be subject to legal liability as he announced the intention of his office to investigate those who apply them (Fink, 2020).

This view seems to conflict with one proposed by the Italian Society of Anesthesia Analgesia Resuscitation and Intensive Care (SIAARTI), who in March 2020 (in one of the most serious phases of the Italian emergency) presented a list of recommendations, offered operational guidelines, and examined ethical concerns in order to assist clinicians involved in the care of COVID patients (Ricconi et al., 2020, pp. 207–211). It advised making choices in the best interests of the greatest number of people. This idea has various practical implications, including the recommendation that resources be directed towards younger patients, who may survive longer following recovery than older ones. They were also urged to utilise resources on patients for whom they would have a greater impact and on those for whom the treatment would most likely be shorter (thereby making the resource available to aid another patient). In practice, it is a utilitarian approach to resource allocation that aims to maximise the overall number of years of life preserved. It entails making trade-offs between years of life preserved for some and years forfeited for others. Specifically, SIAARTI stated that:

An age limit for the admission to the ICU may ultimately need to be set. The underlying principle would be to save limited resources which may become extremely scarce for those

who have a much greater probability of survival and life expectancy, in order to maximize the benefits for the largest number of people (Shaw, Harvey & Gardiner, 2020).

The protocol developed by the Piedmont region’s civil protection department goes even further, stating that

The requirements for access to intensive treatment in times of emergency must include age of less than 80 or a score on the Charlson comorbidity Index of less than five (Shaw, Harvey & Gardiner, 2020).

Peterson, Largent & Karlawish (2020) call this a “soft” utilitarian approach, and point out that it is already being used in areas of Italian healthcare where resources are severely limited, such as organ transplants. By looking at both the likelihood of survival and “maximum life expectancy,” as well as projected length of ICU stay and the consequent use of intensive care resources, the Italian authorities use this strategy to maximise potential value in terms of additional years of life (Cillo et al., 2015, pp. 2552–2561). However, some authors have highlighted that this strategy is in direct contrast with the egalitarianism of the Italian healthcare system (Craxì et al., 2020, pp. 325–330).

Whilst terms such as “soft” utilitarianism may sound appealing, they offer limited scope for understanding, let alone estimating, the welfare loss of a pandemic. During a pandemic, one of the simplest methods for people to lessen their risk of infection is to limit their contact with infected persons. These notions do not include social distancing strategies, which are behavioural modifications that reduce contact rates between those who are at risk of contracting a disease and those who are already ill. The severity of a pandemic can be reduced by social distancing techniques, but the advantages of social distancing rely on the extent to which individuals employ it. When people do not want to pay the expenses associated with social separation as an effective form of control, it can be less successful.

Clearly, such seemingly ad-hoc categorisation of utilitarianism as “soft” or “ruthless” lacks applicability. Should we, for example, interpret “ruthless” utilitarianism as pure utilitarianism? Probably not. Pure utilitarianism (the non-rule variety) proves very difficult to apply on the level of policy, as evidenced by the many paradoxes which arise when it is taken to its logical conclusion. So, what kinds of utilitarianism, if they are at all comparable, did the US authorities find so objectionable and Italian authorities advocate? The Italian authorities were arguably more explicit with their recommendations.

The SIAARTI recommendations can be summarised as follows (Vergano et al., 2020a; 2020b):

1. When the availability of resources is overwhelmed by their need, a decision to deny access to one or more life-sustaining therapies, solely based on the principle of distributive justice, may ultimately be justified.
2. Criteria for allocation should be flexible and adapted locally in response to available resources, the potential for patient transfer, and the ongoing or foreseen number of admissions.
3. An age limit for admission to the ICU may ultimately need to be set.
4. Together with age, the comorbidities and functional status of any critically ill patient should be carefully evaluated.
5. Every admission to the ICU should be considered and communicated as an “ICU trial”. The appropriateness of life-sustaining treatments should be re-evaluated daily.

The extant literature examining the Italian response seems to focus on the early reactions (disbelief and inactivity) of both the Italian national and regional administrations, as well as the general population. Later in the crisis, however, the national government took rapid and decisive action to address the health crisis because of the early designation of an emergency state. Could Italian authorities have followed China's example and implemented much stricter controls? Lockdown measures like those imposed by China in the middle of the crisis would arguably have been difficult to implement because they necessitated a trade-off between individual rights and the need to limit, or at least minimise, the spread of the virus, which was unprecedented in a democratic government. Non-essential production operations would arguably have had to be shut down as a result of a lack of support from both the business industry and union officials for lockdown measures (Bosa et al., 2020).

A wide-ranging debate remains as to whether the Italian government's reaction, as well as other variables, played a significant causal role in the high death rate. Given Italy's lack of centralization, readiness and containment may have been compromised. Italy's age structure, its geographical concentration, and slow and poor decision-making all contributed to the spread of the virus into Italy's care-homes and other vulnerable populations, as did a dearth of knowledge of the virus and its effects, as well as a lack of both national and regional coordination, which allowed the virus to enter the care-home sector. Lockdown had become the norm throughout Europe and the rest of the globe by the time Italy started lifting restrictions in May. As a result, the government was able to maintain control over the virus, and a considerable period of time (spring and summer) followed before the second wave started to present new issues.

In terms of the "standard framework" of ethics (see section 2), the SIAARTI recommendations may plausibly be interpreted as something akin to "rule utilitarianism". Translated into utility theory, it states that although preferences are pre-defined, choices are included in the model.³ But does it do this in a logically satisfactory way? In other words: does it subscribe to e.g. "rule utilitarianism" in a way that is internally consistent and logically valid? An alternative is what Thomson & Lensberg (2006) refer to as the "axiomatic approach".

Axiomatic approach

The axiomatic approach consists in formulating and then solving problems related to social choice. The problems are then checked to see whether they are compatible: that is, whether or not there are solutions which fulfil all of the axioms. The axioms are the key intellectual challenge in this framework. The difficulties arise, inter alia, in the process of achieving internal consistency, a desiderata of formal-axiomatic systems. For example, the axioms may conflict. The desire to impose more axioms than are mutually compatible is not unheard of either.

If a solution meets all of the axioms on a one-by-one basis, it may turn out to be unsatisfactory in a specific context. Then, the logical next step is to identify the class of cases exemplified and create an extra axiom indicating how the solution should behave on that class, and lastly to discover the biggest subset of axioms from the original list that are consistent with the new axiom. There may be several unique subsets that are compatible with each other. Maximal compatibility analysis reveals the "cost" of each axiom and the "trade-offs" between axioms through methodical research. Ultimately, progress in addressing the problems of social choice may be achieved on the basis of such facts and one's intuition about the various axioms.

³ Which is obviously very different from e.g. categorical imperative where choices are/should be predefined in the model. Needless to say, public policy subscribes to "rule utilitarianism" far more frequently.

This allows for a richer and more concrete taxonomy, such as one proposed by (Blackorby, Bossert & Donaldson, 2005) and includes critical-level generalized utilitarianism, restricted critical-level generalized utilitarianism, number-sensitive critical-level generalized utilitarianism, restricted number-sensitive critical-level generalized utilitarianism, number-dampened generalized utilitarianism (Hurka, 2000), and rank-discounted critical-level generalized utilitarianism (Asheim & Zuber, 2014, pp. 629–650; Spears & Stefánsson, 2021). The main properties of the axiomatic population ethics framework are tabulated in Table 1, which is a non-exhaustive list.

Table 1: Principles of population ethics (Blackorby, Bossert & Donaldson, 2005)

		Utility Independence	Existence Independence	Negative Expansion	Avoidance of the Repugnance Conclusion	Priority for Lives Worth Living
1	CU	✓	✓	✓	✓	
2	CLU	✓	✓	✓	✓	
3	RCLU			✓	✓	✓
4	NCLU	✓		✓	✓	
5	RNCLU			✓	✓	✓
6	AU				✓	✓
7	RAU			✓	✓	✓
8	NDU					✓
9	RNDU					✓
10	NDU					✓
11	RNDU			✓	✓	
12	NDU				✓	✓
13	RNDU			✓	✓	✓
14	RDCLU			✓	✓	✓

The abbreviations in Table 1 can be explained as follows:

1. Classical utilitarianism (CU)
2. Critical-level utilitarianism: positive critical level.
3. Restricted critical-level utilitarianism: positive critical level parameter
4. Number-sensitive critical-level generalized utilitarianism
5. Restricted number-sensitive critical-level generalized utilitarianism: restricted version of 4
6. Average utilitarianism (AU)
7. Restricted average utilitarianism
8. Number-dampened utilitarianism: general case
9. Restricted number-dampened utilitarianism: general case
10. Number-dampened utilitarianism: ratio of critical level to average utility is a positive constant less than one
11. Restricted number-dampened utilitarianism: restricted version of 10
12. Number-dampened utilitarianism: ratio of critical level to average utility is positive, nondecreasing, and approaches one as population size increases
13. Restricted number-dampened utilitarianism: restricted version of 12
14. Rank-discounted critical-level generalized utilitarianism

For formal definitions of all SWO functions in Table 1 the interested reader is referred to (Thomson & Lensberg, 2006; Blackorby, Bossert & Donaldson, 2005; Asheim & Zuber, 2014) and references therein. In the below illustrative examples, we provide some formal definitions and state some important conditions. We follow Asheim and Zuber's notation (Asheim & Zuber, 2014).

Let $\mathbf{X} = \bigcup_{n \in \mathbb{N}} \mathbb{R}^n$ be the set of possible finite allocations of lifetime well-being. For each $n \in \mathbb{N}$, each allocation $n \in \mathbb{R}^n$ determines the finite population size, $n(\mathbf{X}) = n$, and the distribution of well-being, $\mathbf{x} = (x_1, \dots, x_{n(\mathbf{x})})$. A social welfare relation (SWR) on the set \mathbf{X} is a binary relation \succeq , where for all $\mathbf{x}, \mathbf{y} \in \mathbf{X}$, $\mathbf{x} \succeq \mathbf{y}$ implies that the allocation \mathbf{x} is deemed socially at least as good as \mathbf{y} . A complete, reflexive, and transitive SWR is called a social welfare order (SWO). For each $\mathbf{x} \in \mathbf{X}$, let $\mathbf{x}_{[\cdot]} = x_{[1]}, \dots, x_{[r]}, \dots, x_{[n(\mathbf{x})]}$ denote the nondecreasing allocation that is a reordering of \mathbf{x} .

The following axioms are uncontroversial in population ethics.

Axiom 1 (Order). *The relation \succeq is complete, transitive, and reflexive on \mathbf{X} .*

Axiom 2 (Continuity). *For all $n \in \mathbb{N}$ and all $\mathbf{x} \in \mathbb{R}^n$, the sets $\mathbf{y} \in \mathbb{R}^n : \mathbf{y} \succeq \mathbf{x}$ and $\mathbf{y} \in \mathbb{R}^n : \mathbf{x} \succeq \mathbf{y}$*

Axiom 3 (Suppes-Sen). *For all $n \in \mathbb{N}$ and all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, if $\mathbf{x}_{[\cdot]} > \mathbf{y}_{[\cdot]}$ then $\mathbf{x} > \mathbf{y}$.*

Axiom 4 (Existence independence of the best off). *For all $n \in \mathbb{N}$ and all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, and $z \in \mathbb{R}$ satisfying $z \geq \max\{x_{[n]}, y_{[n]}\}$, $(\mathbf{x}, z) \succeq (\mathbf{y}, z)$ if and only if $\mathbf{x} \succeq \mathbf{y}$.*

Axiom 5 (Existence independence of the worst off). *For all $\mathbf{x}, \mathbf{y} \in \mathbf{X}$, and $z \in \mathbb{R}$ satisfying $z \geq \min\{x_{[1]}, y_{[1]}\}$, $(\mathbf{x}, z) \succeq (\mathbf{y}, z)$ if and only if $\mathbf{x} \succeq \mathbf{y}$.*

Axiom 6 (Existence of a critical level). *There exists $c \in \mathbb{R}_+$ and $n \in \mathbb{N}$ such that, for all $\mathbf{x} \in \mathbb{R}^n$ satisfying $x_{[n]} \leq c$, $(\mathbf{x}, c) \sim \mathbf{x}$*

All axioms above are also satisfied by ordinary critical-level generalized utilitarianism. However, as Asheim & Zuber (2014) points out, the CLU SWO yields the Repugnant Conclusion when $c = 0$ and leads to the Very Sadistic Conclusion if $c > 0$. The authors propose a weaker axiom as follows and argue that axiom 7 is key to avoiding the Repugnant and Very Sadistic Conclusions, while not by itself contradicting these conclusions:

Axiom 7 (Existence of egalitarian equivalence). *For all $\mathbf{x}, \mathbf{y} \in \mathbf{X}$, if $\mathbf{x} \succeq \mathbf{y}$, there exists $z \in \mathbb{R}$ such that, for all $N \in \mathbb{N}$, $\mathbf{x} > (z)_n > \mathbf{x}$ for some $n \geq N$*

Axiom 8 (Existence of egalitarian equivalence). *For all $n, m \in \mathbb{N}$, for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, and for all $\mathbf{u}, \mathbf{v} \in \mathbb{R}^m$, $(\mathbf{y}, \mathbf{u}) \succeq (\mathbf{x}, \mathbf{v}) \Leftrightarrow (\mathbf{x}, \mathbf{v}) \succeq (\mathbf{y}, \mathbf{u})$*

It is possible to define established generalized utilitarian SWOs of population ethics.

Definition 4.1 (AU). *An SWR \succeq is an average utilitarian⁴ (AU) SWO if there exists a continuous and increasing function $u : \mathbb{R} \Rightarrow \mathbb{R}$ such that for all $\mathbf{x}, \mathbf{y} \in \mathbf{X}$,*

$$\mathbf{x} \succeq \mathbf{y} \Leftrightarrow \frac{1}{n(\mathbf{X})} \sum_{r=1}^{n(\mathbf{X})} u(x[r]) \geq \frac{1}{n(\mathbf{Y})} \sum_{r=1}^{n(\mathbf{Y})} u(y[r])$$

⁴ Also called average generalized utilitarian.

Definition 4.2 (CLU). *An SWR \succeq on \mathbf{X} a critical-level generalized utilitarian (CLU) SWO if there exist $c \in \mathbb{R}_+$ and a continuous and increasing function*

$$\mathbf{x} \succeq \mathbf{y} \Leftrightarrow \frac{1}{n(\mathbf{X})} \sum_r^{n(\mathbf{X})} u(x[r]) \geq \frac{1}{n(\mathbf{Y})} \sum_r^{n(\mathbf{Y})} u(Y[c])$$

The CLU SWO with $c = 0$ is the total generalized utilitarian (TU) SWO (Asheim & Zuber, 2014, pp. 629–650). CLU SWO with $c > 0$ also leads to the Very Sadistic Conclusion that there is an egalitarian allocation with positive well-being that is worse for every allocation with negative well-being. The RDCLU SWO may be attractive since it offers an escape from the Repugnant and Very Sadistic Conclusions.

Definition 4.3 (RDCLU). *An SWR \succeq on \mathbf{X} is a rank-discounted critical-level generalized utilitarian (RDCLU) SWO if there exist $c \in \mathbb{R}_+$, $\beta \in (0, 1)$ and a continuous and increasing function $u : \mathbb{R} \Rightarrow \mathbb{R}$ such that for all $\mathbf{x}, \mathbf{y} \in \mathbf{X}$,*

$$\mathbf{x} \succeq \mathbf{y} \Leftrightarrow \frac{1}{n(\mathbf{X})} \sum_{r=1}^{n(\mathbf{X})} \beta^r (u(x[r]) - u(c)) \geq \frac{1}{n(\mathbf{Y})} \sum_{r=1}^{n(\mathbf{Y})} \beta^r (u(Y[r]) - u(c))$$

According to the generalised utilitarian criteria, the function u transforms lifelong well-being into transformed values. The term “rank-discounted” refers to the fact that the utility weights are discounted according to rank by a geometrically decaying function, rather than just being rank-dependent. This is similar to the time-discounted utilitarian criteria of intertemporal social choice, which uses “time-discounted utilities” rather than “time-dependent utilities.” The parameter c is referred to as a critical-level parameter. The recognised generalised utilitarian SWOs of population ethics can now be defined. To be consistent with the literature on intertemporal social choice, we refer to the converted values as utility and β as a rank utility discount factor.

Axioms 1–7 produce Rank-discounted critical-level generalized utilitarianism (RDCLU).⁵ Similarly, axioms 1–6 together with axiom 8 produce Critical-level generalized utilitarianism (CLGU). Spears & Zuber (2021) present new axiomatic characterizations of utilitarian (that is, additively separable) social welfare functions that generalize Blackorby, Bossert & Donaldson (1998). Expected Critical-Level Generalized Utilitarianism (ECLGU). Spears & Stefánsson (2021) further generalise the results of (Asheim & Zuber, 2014, pp. 629–650) and explain that those who are drawn to Axioms 1–6 but wish to keep the common policy-evaluation practice of same-number independence might be tempted to use CLGU, which incorporates:

1. priority for the worse-off (for concave u), and
2. utilitarianism as special case (for linear u).

More recently, Pestieau & Ponthiere (2022) provide a discussion of the Repugnant Conclusion in the context of the current COVID-19 pandemic and the design of an optimal lockdown policy. They propose that some forms of utilitarianism lead to maximum containment (albeit under certain conditions): for each lockdown with low-quality life periods, there must be a harsher lockdown that is considered as superior, even if it lowers the quality of life periods even more.

⁵ This is achieved via Lemmas 1–4 in (Asheim & Zuber, 2014), proof is presented in Asheim & Zuber (2014, p. 642).

Game-theoretic models

In order to account for the feedback between the status of the pandemic and people's behaviour, formal (i.e. game-theoretic) models of social distance have gained ascendancy. Despite their differences, these models tend to have a similar message: the outbreak is not as serious as expected without taking behavioural response into consideration. Individuals tend to return to their old habits when the pandemic subsides and there is less chance of exposure. Due to self-limiting feedback, the pandemic can persist for a long time in an intermediate-severity state. This is not so awful that all individuals take it seriously enough to separate themselves, but maintains enough of a threat that some people do.

Game-theoretic study of agents' motivations to adopt precautionary steps to avoid infection during an epidemic was offered in the 1980s by Fine & Clarkson (1986, pp. 1012–1020). Following Philipson & Posner (1993) and Kremer (1996), as well as Geoffard & Philipson (1996), the 1990s saw more advanced dynamic analysis. Because of the negative feedback between illness prevalence and the motivation to take precautions, there is a limit to what may be done by voluntary preventative measures. It is impossible to eliminate diseases that spread by chance encounters: for example, the effect of vaccination diminishes as the disease approaches eradication (see, for instance, Geoffard & Philipson, 1997, pp. 222–230)

The types of models referenced are well established and have their origins in the early epidemiology literature, such as the seminal studies of Kermack & McKendrick made as far back as 1926. Indeed, many public health professionals use standard epidemiology models of the kind originally introduced by Kermack & McKendrick in 1926 (Kermack & McKendrick, 1927, pp. 700–721)⁶ and recently updated by, inter alia, (Alvarez et al., 2020) to study the effects of an optimal lockdown policy in the context of COVID-19. It is possible, using arguments from economics and epidemiology, to investigate the ideal lockdown approach for a social planner who wishes to limit pandemic mortality while reducing lockdown costs to output. To formalise the planner's dynamic control problem, it is common to employ a standard SIR-type (susceptible-infected-removed) epidemiological model and a linear economy, calibrating the model to actual data. The best policy is then determined by the proportion of infected and vulnerable people in the population. The models have certain variations and special cases, such as SI, SIR/SIS/SIRS, SCIR, and SCIS, see also the recent survey article of (McAdams, 2020b).⁷ McAdams helpfully presents a tabulated summary of the literature (page 12) – all from either the economics, epidemiology, or public health literatures. However, none of the papers presented therein satisfactorily address their modelling assumptions. Either the authors make an informal comment about using a "utilitarian welfare function" or, for the most part, they do not make explicit their assumptions. In general, the discussion of welfarist axiology is rare outside of narrowly specialised academic research, despite welfarist assumptions commonly grandfathered in or baked into the policy modelling.

Probing latent assumptions of policy models

Let us now turn to the question of assumptions that are hidden or implicit in models used to guide epidemiological, economic, or public health-related decision making. The illustrative example we provide later in this paper, which will undoubtedly be valuable to ethicists and applied philosophers alike, borrows from Alvarez, Argente, & Lippi (2021), a well-cited and recent paper from the economics literature published in one of the world's best known economics journals. In the context of COVID-19, it serves as the foundational study for a line of research articles

⁶ See (Anderson, 1991, pp. 1–32) for a historical discussion.

⁷ Republished as (McAdams, 2020a), see page 11.

examining dynamic interventions. Alvarez, Argente, and Lippi rely heavily on (Hall, Jones & Klenow, 2020), who use a utilitarian criterion to value the extra years of life lost among those likely to die due to the infection.

Next, we will explain how such SIR-type models, composed of three connected nonlinear ordinary differential equations, inform decision-making. Since citizens under lockdown are expected to be inactive, it is common to assume that the goal of a policy planner is to reduce both the present discounted value of deaths and the output costs of lockdown. A common framework for analysis of such questions in epidemiology involves using the SIR epidemiology model, which describes how the virus spreads from infected agents to those who are susceptible, as well as the rates at which infected agents either recover or die. In this framework, a pandemic planner has just one tool at their disposal: the lockdown of citizens.

It is the planner's dilemma to balance the costs of locking down, which increase as the number of diseased and vulnerable individuals grows, against the costs of a pandemic, which increase as the number of infected and susceptible individuals rises. This framework lends to examination of how the severity and length of a lockdown are affected by the cost of deaths, which is assessed by how much a life is worth statistically, as well as how successful the lockdown is (i.e. how many contacts are reduced when residents are advised to remain at home). Increasing the fatality rate (chance of dying if infected) increases the policy maker's motivation for lockdown, as is expected to occur after the hospital capacity is surpassed. However, in the absence of non-linear costs, an ideal gradual lockdown may occur, as the costs and benefits of the policy depend on the current state of the system.

Using quantitative analysis, it is possible to demonstrate these ideas using simulations that enable us to explore the issue from a practical point of view. Whilst there are many factors to consider while designing the best lockdown strategy, quantitative findings that draw from such modelling may help to determine which ones are most significant. Let us illustrate our argument with a fundamental disease transmission model. The model is composed of three connected nonlinear ordinary differential equations, none of which have an explicit formula solution. However, basic calculus techniques enable us to extract a wealth of information about solutions. Along the way, we demonstrate how this basic model contributes to the theoretical underpinning of public health interventions and illuminates numerous public health foundations. The population at time t may be classified as susceptible to infection $S(t)$, infected $I(t)$, and recovered $R(t)$

$$N(t) = S(t) + I(t) + R(t) \text{ for all } t \geq 0 \quad (6.1)$$

In the above specification:

- $S(t)$: susceptible persons who are not now infected, but who may become infected in the future. There is a chance that a person who is vulnerable to infection may get infected. Over time, as the virus spreads from its point of origin or as new sources emerge, the number of people at risk will rise (surge period).

- $I(t)$: persons who have been exposed to the virus. These are people who have previously been exposed to the virus and can pass it on to others. When an infected person is taken out of the contaminated group, he or she might either recover or die.
- $R(t)$ This is the group of people who have recovered from the virus and are presumed to be immune, $R(t)$ or have died, $D(t)$.

Such an analytical framework is beneficial both numerically and qualitatively because the SIR model simplifies many of the difficulties involved with understanding viral dissemination in real time. The population N is normalised to one, hence all findings should be regarded as fractions of the relevant population. The model relies on related ODEs to describe populations. The social planner can make a decision to lockdown at fraction $L(t) \in [0, \bar{L}]$ of $S(t) + I(t)$ with effectiveness $\theta \in (0,1)$. Those who are susceptible become infected when they come into contact with I at rate $\beta(1 - \theta L(t))^2$ (a single dot above a letter indicates the first differential with respect to time):

$$\dot{S}(t) = -\beta S(t) \times I(t)(1 - \theta L(t))^2 \quad (6.2)$$

Among those not under lockdown, β is the number of vulnerable agents per unit of time to whom an infected agent can transmit the virus. In the context of this model, it is assumed that the SIR time scale is short enough to ignore births and deaths (except from those caused by the virus) while at the same time allowing for a minimal number of fatalities from the disease. Amongst the infected, a fraction γ recovers:

$$\dot{I}(t) = \beta S(t) \times I(t)(1 - \theta L(t))^2 - \gamma I(t) \quad (6.3)$$

A percentage of those infected die at rate of $0 < \phi(I) \leq \gamma$ per unit of time

$$\dot{N}(t) = D(t) = \phi(I(t))I(t) \quad (6.4)$$

It is worth noting that, in this epidemiological model, isolating a portion of the population, while economically costly, can be quite effective in reducing the rate at which the vulnerable become infected. This is because the product of the infected and susceptible determines the number of new infections per unit of time. As a result, reducing each person's number of contacts reduces the number of new infections by its square.

Then, a social planner chooses path $\{L(t)\}$ to minimize the present discounted value:

$$V(S, I) = \min_{L(t)} \int_0^{\infty} e^{-(r+v)t} \left(\underbrace{wL_t[\tau(S_t + I_t) + 1 - \tau]}_{\text{GDP loss}(t)} + \underbrace{\phi(I(t))I(t) \times [vsl]}_{\text{deaths}(t) \times \text{value of life}} \right) \quad (6.5)$$

by choosing the fraction of population $L(t)$ to lockdown.

According on the parameter τ , testing is either available or not. The parameter τ therefore corresponds to testing of the recovered i.e., there are two scenarios:

- with testing of the recovered ($\tau = 1$): lockdown to (S+I)
- without testing of the recovered ($\tau = 0$): lockdown to (S+I +R)

The social planner solves the Bellman equation, where the value function $V(S, I)$ such that $0 \leq S + I \leq 1$ solves

$$(\tau + \nu)V(S, I) = \min_{L \in [0, L]} [\tau(S + I) + 1 - \tau] + I\phi(I)\left[\frac{w}{r} + \chi\right] + \quad (6.6)$$

$$+ [\beta S(t)(1 - \theta L(t))I(t)(1 - \theta L(t))] \partial_S V(S, I) + \quad (6.7)$$

$$+ [\beta S(t)(1 - \theta L(t))I(t)(1 - \theta L(t)) - \gamma I(t)] \partial_I v(S, I) \quad (6.8)$$

Here, $V(S, I)$ might be read as the minimal predicted discounted loss in output units for implementing the optimal strategy. After setting, the initial conditions $N(0) = 1$, $I(0) = \epsilon$, $S(0) = 1 - \epsilon$, and boundary conditions:

- Boundary condition at $I = 0$, $\forall S \in (0, 1) : V(S, 0) = 0$
- Boundary condition at $S = 0$, $\forall I \in (0, 1) : V(0, I) = 0 = \nu S I \times \left(\frac{\phi \gamma}{r + \kappa + \gamma I} + \frac{\phi \gamma}{r + \nu + 2\gamma} \right) \times I$

We then use data from the World Health Organization, alongside key economic variables to parameterise the model. Whilst our methodology is consistent with (Alvarez, Argente & Lippi, 2021), we calibrate the parameters to Italian data where practicable.

Even if a person has already recovered from the illness, the lockdown is still in effect. In this scenario, locking down the population is less efficient since the recovered are also locked down, with the cost of limiting productivity without the advantage of reducing viral transmission. The lockdown effectiveness decreases significantly in the event of no testing compared to the benchmark condition. In both circumstances, the lockdown costs the same amount of time in terms of lost productivity, because the lockdown length is shorter in the absence of testing, but the lockdown affects a higher percentage of workers (recovered agents are also in lockdown). It is the most important element of a situation when a test is unavailable that the lockdown comes to an abrupt conclusion sooner. As time goes on, the percentage of people who have recovered grows, and therefore the lockdown becomes less effective at stopping the virus from spreading by locking down an increasing number of those who do not have it.

A natural question that arises is how to set the fatality rate. Extensive testing has been done in two cases: the Diamond Princess cruise ship (which yields an age-adjusted fatality rate) and the Italian city of Vo' Euganeo (which yields a lower bound mortality rate). Consistent with both cases, we set the mortality rate $\phi = 0.01 \times \gamma$. We then set $\kappa = 0.05\gamma$ so that the fatality rate is 3% when 40 % of the population is infected. The discount factor used by the social planner is set at 5% annual interest. The likelihood ν per unit of time that a vaccine or treatment will be discovered means that these medical breakthroughs take an average of one and a half years to become available. In accordance with (Alvarez, Argente & Lippi, 2021), we normalise output $w = 1$. In accordance with (Hall, Jones & Klenow, 2020), we chose the benchmark value for the additional cost of dying χ to 0. A utilitarian criterion is used by these authors to value the additional years of life lost by individuals who are more likely to die as a result of the illness, and they arrive at a cost of roughly 30 times annual consumption per capita, which is fairly close to our benchmark value of 20 times annual GDP. We set $\chi = 0$ as the penalty death threshold, which is consistent with (Hall, Jones & Klenow, 2020; Alvarez, Argente & Lippi, 2021, pp. 367–382) for a brief discussion).

Our results are as follows. Total welfare costs are almost three times greater due to the cost of deaths. The paths obtained are under medium effectiveness (i.e. $\theta=0.5$). The best strategy reduces

the overall number of deaths by around 1% in the long run.⁸ Whilst the numerical magnitude of the results obtained might merit examination, we would like to focus the reader's attention on the assumptions involved. Specifically, Alvarez, Argente & Lippi (2021, pp. 367–382) outline some of the assumptions made in such modelling, arguing that these are done in order to make the model tractable. Specifically,

Assumption 1. *The first assumption is that a lockdown is limited to $\bar{L} \leq 1$. This is due to the fact that some sectors, such as health and fundamental services, cannot be shut down even in the most severe lockdown.*

Assumption 2. *Members of the population who are sick but not in lockdown are assumed to be able to work as much as those who are susceptible or have recovered. After locking down those exhibiting symptoms, the rest of the group should be isolated.*

Assumption 3. *Population that is put under lockdown will not produce.*

Assumption 4. *Except for risk of dying following virus contraction, we assume that all agents are infinitely long-lived. Given the limited time frame of the problem, this simplification is acceptable. However, the lack of a defined age structure results in unrealistic impacts on mortality risk.*

Assumption 5. *The initial proportion of the population susceptible is assumed to be 98%, or $S(0) = 0.98$*

Whilst these assumptions are made explicit by the authors, there is a more nuanced (and hidden) assumption which is grandfathered from Hall, Jones & Klemow (2020). Herein lies our “hidden assumptions” problem. In this example, Hall, Jones, and Klemow use a (classic) utilitarian criterion to value the additional years of life lost by individuals who are more likely to die as a result of the virus.

Whether philosophers use them, understand them, or appreciate their role in informing public policy-making and advising scientific bodies, such models are important tools of public policy (see, for example, (Griffiths, et al., 2021) and the references cited therein). SIR-type models are the workhorse tools of the COVID-19 pandemic. These models are being utilised by public health experts and policy makers to develop scenarios that are used to guide judgments about recommending or mandating even more stringent mitigation measures on economies (globally) in response to its fast growth. Ethicists and moral philosophers operate at a distance from the front lines.

However, ethicists and moral philosophers could have an important role to play in the policy debate by offering principled, philosophical guidance on how to address the extent and bounds of individual liberty in the context of a pandemic; how to think about collective interests constraining individual interests; and how to assess the soundness of assumptions that underlie mathematical models used in policy modelling. But because philosophers appear to be largely unfamiliar with these models – along with the assumptions underlying them – they are not fully involved in the policy debate over social distancing measures. To understand the trade-off between public health and the policy implications of these mitigation and social distance measures, there is a case to be made that philosophers ought be aware of the logical consequences of these models. It should be stressed, however, that such models are particularly sensitive to seemingly innocuous assumptions regarding the dynamics of infectious diseases.

⁸ Full results, as well as the Matlab code for numerical calculations, are available on request.

What other insights, if any, are we able to glean from economics, or one of its sub-fields, such as economic epidemiology? Arguably, economic analysis has much to offer to the extent that we envisage avenues for further research where applied philosophy and economics are seen as complements rather than substitutes. For example, the literature on welfare economics is useful for elucidating the externality elements of infectious diseases as well as broad policy options. The assumptions about the social health planner, on the other hand, are accompanied by a set of well-known objections. These include concerns regarding knowledge of the social welfare function, the availability of proper incentives which results in resource allocation that maximises social welfare, and the lack of unexpected consequences that weaken welfare efforts. The traditional method places the social health planner outside of society and separates policy-making from economic analysis.

Nevertheless, there are some meaningful insights to be obtained. Whilst the literature is both vast and actively developing, we point the interested reader to the recent work by Toxvaerd & Rowthorn (2020) who, building on the economic epidemiology literature, compare socially optimum treatment and immunisation strategies to the old non-economic approach. Non-economic policies not only have poorer social welfare than socially optimal policies, but they may even be outperformed by fully non-responsive policies that do not maximise any goal at all, according to the authors. These findings show that in order to develop effective and welfare-enhancing medicinal therapies, researchers must first examine how to effectively fulfil well-defined social goals. When policies clearly consider the costs and benefits of developing population immunity, they may help to strike the correct balance and increase social welfare.⁹

In line with our “Rule utilitarianism” approach mentioned earlier, a method borrowed from health economics might be used to minimise the effect of pandemics on years of well-being by assessing the utility loss from restrictions on economic activities in terms of QALYs. The value of gains in well-being over the course of a year is equal for all persons, and the value of losses in well-being is likewise equal for all people when this technique is used. This strategy, although common, is not without its drawbacks. On the basis of the utilitarian principle and the premise that individuals have linear utility, it implicitly depends on this approach. One dollar earned here equals one dollar gained somewhere else if utilities are linear (and same across persons). This assumption makes it simple to replace monetary values everywhere. Even though this classic utilitarian approach, using linear utility functions, is often employed in health care cost-effectiveness analysis, it is clear that some kind of value judgement is made regarding the way in which expenses and benefits should be distributed.

Conclusion

In this discussion, we focused on the context of the Italian reaction to the COVID-19 crisis in order to create a useful resource for researchers and policymakers alike. Italy was the first country in Europe to be affected by the COVID-19 outbreak and this is significant to our discussion for the following reason: whilst European policymakers had the opportunity to learn from other countries’ responses and adapt their policies accordingly, Italian policymakers – effectively being Europe’s “patient zero” – did not have this relative advantage. Therefore, *rules* rather than *data* asserted path dependence and informed decision-making in the early stages of the crisis.

The most helpful ethical theories help to answer society’s questions, many of which have arisen due to the extraordinary challenges brought about by the pandemic. Classic utilitarianism, which

⁹ The interested reader is invited to see (Chen & Toxvaerd, 2014; Gersovitz, 2003) and (Gersovitz, 2011) for a richer discussion.

employs the sum of utilities to evaluate options, and average utilitarianism, which utilises average utility, are the two most well-known welfarist principles. These concepts are often rejected since classical utilitarianism leads to some sub-optimal results, while average utilitarianism occasionally deems the *ceteris paribus* addition of a person whose life is not worth living to be beneficial. However, there are many principles that eschew both of these characteristics. The central question of this paper was as follows: how should we think of citizens' preferences? In a democracy, difficult choices involving trade-offs of lives saved vs other economic and societal issues must represent voters' preferences. Is there a role for ethicists and philosophers in this process? Undoubtedly, the answer is yes. A philosopher could assist policy choices by defining the moral principles that people are ready to follow when difficult decisions must be made, and to advise politicians about the trade-offs that would represent these preferences.

Utilitarian welfare functions are appealing because they are consistent with the assumption that social preferences satisfy the von Neumann Morgenstern axioms (Harsanyi, 1955) and the Strong Pareto assumption.¹⁰ But there are, of course, other ways to think about social welfare. For example, prioritarianism prioritises outcomes with lower degrees of well-being by calculating a strictly increasing and concave transformation of well-being. This prioritises those with lower degrees of well-being. Credited to moral philosopher Derek Parfit, prioritarianism is a relatively new notion (Parfit, 2000, pp. 81–125). In public policy, it is standard to define social welfare as maximising a person's expected present value of utility over the frequency with which he or she interacts with others and over all time periods. An individual utility function is the aggregation of unit-comparable functions. It is necessary to standardise individual utility functions so that transitory utility prior to the COVID-19 pandemic is the same for all persons.

While utilitarianism is not without its detractors, utility theory appears to be well-equipped to deal with a number of theoretical challenges, such as questions around redistribution. As demonstrated in this paper, the axiomatic approach requires minimal conditions for normative plausibility with only one axiom sufficient to avoid both Repugnant and Very Sadistic Conclusions.

Both the public policy and economic epidemiology literatures appear to be silent on the relationship between axiomatic properties of social welfare relations and evaluating output costs of the lockdown. The aim of this paper was to address this research gap and provide tentative suggestions for future research.

There exists a lacuna between two literatures, namely between economic epidemiology and population ethics. The critical evaluation of hidden assumptions does not go hand-in-hand with expertise in formal modelling. As a result, important assumptions that underpin policy models are either incomplete or omitted. A concealed assumption is not always incorrect; it may be perfectly appropriate. However, until we are aware that it exists, we cannot assess it and so cannot reach a complete conclusion regarding the soundness of the model or its output.

We therefore envisage avenues for further research where applied ethics and public policy are seen as close complements, working towards a plausible, balanced, and consistent welfarist axiology that can rigorously support models within economics and public health.

References

ALVAREZ, F., ARGENTE, D. & LIPPI, F. (2021): A simple planning problem for COVID-19 lock-down, testing, and tracing. In: *American Economic Review: Insights*, 3(3), pp. 367–382.

¹⁰ It means that society prefers one allocation over another if all individuals weakly prefer it and at least one individual strictly prefers it.

ANDERSON, R. M. (1991): Discussion: the Kermack-McKendrick pandemic threshold theorem. In: *Bulletin of mathematical biology*, 53(1), pp. 1–32.

ASHEIM, G. B. & ZUBER, S. (2014): Escaping the repugnant conclusion: Rank-discounted utilitarianism with variable population. In: *Theoretical Economics*, 9(3), pp. 629–650.

ATKESON, A. (2020): What will be the economic impact of COVID-19 in the US? Rough estimates of disease scenarios (No. w26867). National Bureau of Economic Research. [online] [Retrieved March 18, 2021]. Available at: <http://www.nber.org/papers/w26867>

BLACKORBY, C., BOSSERT, W., DONALDSON, D. (1998): Uncertainty and critical-level population principles. In: *Journal of Population Economics*, 11(1), pp. 1–20.

BLACKORBY, C., BOSSERT, W. & DONALDSON, D. J. (2005): *Population issues in social choice theory, welfare economics, and ethics*. Cambridge: Cambridge University Press.

BOSA, I., CASTELLI, A., CASTELLI M., CIANI, O., COMPAGNI, A., GAROFANO, M., GIANNONI, M., MARINI, G. & VAINIERI, M. (2020): *Cambridge Core Blog*. [online] [Retrieved December 14, 2021] Available at: <https://www.cambridge.org/core/blog/2020/04/16/italys-response-to-the-coronavirus-pandemic/>

BRITISH MEDICAL ASSOCIATION (2020): *Covid-19 ethical issues: a guidance note*. [online] [Retrieved January 20, 2022] Available at: <https://www.bma.org.uk/media/2226/bma-covid-19-ethics-guidance.pdf>

CHEN, F. H. & TOXVAERD, F. (2014): The economics of vaccination. In: *Journal of theoretical biology*, 363, pp. 105–117.

CILLO, U., BURRA, P., MAZZAFERRO, V., BELLI, L., PINNA, A. D., SPADA, M., NANNI COSTA, A. & TONIUTTO, P. on behalf of I-BELT (Italian Board of Experts in the Field of Liver Transplantation) (2015): A multistep, consensusbased approach to organ allocation in liver transplantation: Toward a “blended principle model”. In: *American Journal of Transplantation*, 15(10), pp. 2552–2561.

CRAXÌ, L., VERGANO, M., SAVULESCU, J. & WILKINSON, D. (2020): Rationing in a pandemic: Lessons from Italy. In: *Asian Bioethics Review*, 12(3), pp. 325–330.

EMANUEL, E. J., PERSAD, G. & UPSHUR, R. (2020): Fair allocation of scarce medical resources in the time of Covid-19. In: *The New England Journal of Medicine*, 382(21), pp. 2049–2055.

FINE, P. E. M. & CLARKSON, J. A. (1986): Individual versus public priorities in the determination of optimal vaccination policies. In: *American journal of epidemiology*, 124(6), pp. 1012–1020.

FINK, S. (2020): U.S. Civil Rights Office rejects rationing medical care based on disability, age. In: *The New York Times*. [online] [Retrieved December 20, 2021] Available at: <https://www.nytimes.com/2020/03/28/us/coronavirus-disabilities-rationing-ventilators-triage.html>

GEOFFARD, P. Y. & PHILIPSON, T. (1996): Rational epidemics and their public control. In: *International Economic Review*, 37(3), pp. 603–624.

GEOFFARD, P. Y. & PHILIPSON, T. (1997): Disease eradication: Private versus public vaccination. In: *American Economic Review*, 87(1), pp. 222–230.

GERSOVITZ, M. (2003): 25 births, recoveries, vaccinations, and externalities. In: R. J. Arnott, B. Greenwald, R. Kanbur, & B. Nalebuff (eds.): *Economics for an imperfect world: Essays in honor of Joseph E. Stiglitz*. Cambridge, MA: MIT Press, pp. 469–483.

GERSOVITZ, M. (2011): The economics of infection control. In: *Annual Review of Resource Economics*, 3(1), pp. 277–296.

- HALL, R. E., JONES, C. I. & KLENOW, P. J. (2020): Trading off consumption and COVID-19 deaths. In: *Quarterly Review*, 42(1), pp. 2–13.
- HARSANYI, J. (1955): Cardinal welfare, individualistic ethics and interpersonal comparisons of utility. In: *Journal of Political Economy*, 63(4), pp. 309–321.
- HURKA, T. (2000): Comment on ‘Population principles with number-dependent critical levels’, unpublished manuscript, Department of Philosophy, University of Calgary.
- KERMACK, W. O. & MCKENDRICK, A. G. (1927): A contribution to the mathematical theory of pandemics. In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 115(772), pp. 700–721.
- KREMER, M. (1996): Integrating behavioral choice into epidemiological models of AIDS. In: *Quarterly Journal of Economics*, 111(2), pp. 549–573.
- MCADAMS, D. (2020a): Economic epidemiology in the wake of Covid-19. In: *Covid Economics*, 48, pp. 1–45.
- MCADAMS, D. (2020b.): Economic epidemiology in the wake of COVID-19. In: *Economics*, Issue 82120, pp. 1–40.
- NEW YORK STATE DEPARTMENT OF HEALTH (2015): *VENTILATOR ALLOCATION GUIDELINES*. [online] [Retrieved February 10, 2022] Available at: https://www.health.ny.gov/regulations/task_force/reports_publications/docs/ventilator_guidelines.pdf
- PANOVSKA GRIFFITHS, J., KERR, C. C., WAITES, W. & STUART, R. M. (2021): Mathematical modeling as a tool for policy decision making: Applications to the COVID-19 pandemic. In: *Handbook of Statistics*, 44, pp. 291–326.
- PARFIT, D. (2000): Equality or Priority? In: M. Clayton & A. Williams (eds.): *The ideal of equality*. Houndmills: Palgrave, pp. 81–125.
- PESTIEAU, P. & PONTIÈRE, G. (2022): Optimal lockdown and social welfare. In: *Journal of Population Economics*, 35(1), pp. 241–268.
- PETERSON, A., LARGENT, E. A. & KARLAWISH, J. (2020): Ethics of reallocating ventilators in the covid-19 pandemic. In: *BMJ*, 369, p. m1828.
- PHILIPSON, T. J. & POSNER, R. A. (1993): *Private choices and public health: The AIDS pandemic in an Economic Perspective*. Cambridge, MA: Harvard University Press.
- RICCIONI, L., BERTOLINI, G., GIANNINI, A., VERGANO, M., GRISTINA, G., LIVIGNI, S., MISTRALETTI, G. & PETRINI, F. (2020): Raccomandazioni di etica clinica per l’ammissione a trattamenti intensivi e per la loro sospensione, in condizioni eccezionali di squilibrio tra necessità e risorse [Clinical ethics recommendations for admission to intensive treatments and their suspension, in exceptional conditions of imbalance between needs and resources]. In: *Recenti Progressi in Medicina*, 111(4), p. 207–211.
- ROYAL COLLEGE OF PHYSICIANS (2020): *Ethical dimensions of COVID-19 for frontline staff*. [online] [Retrieved March 10, 2022] Available at: replondon.ac.uk/file/20551/download
- SEN, A. (1987): *The Standard of Living*. Cambridge: Cambridge University Press.
- SHAW, D., HARVEY, D. & GARDINER, D. (2020): Don’t let the ethics of despair infect the intensive care unit. In: *Journal of Medical Ethics Blog*. [online] [Retrieved February 10, 2022] Available at: <https://blogs.bmj.com/medical-ethics/2020/03/26/dont-let-the-ethics-of-despair-infect-the-intensive-care-unit/>.
- SPEARS, D. & STEFÁNSSON, H. O. (2021): Additively-separable and rank-discounted variable-population social welfare functions: A characterization. In: *Economics Letters*, 203, p. 109870.

SPEARS, D. & ZUBER, S (2021): Foundations of utilitarianism under risk and variable population. In: *IZA Discussion Paper*, No. 14515. [online] [Retrieved February 10, 2022]. Available at SSRN: <https://ssrn.com/abstract=3879363> or <http://dx.doi.org/10.2139/ssrn.3879363>

THOMSON, W. & LENSBERG, T. (2006): *Axiomatic theory of bargaining with a variable number of agents*. Cambridge: Cambridge University Press.

TOXVAERD, F. & ROWTHORN, R. (2020): *On the management of population immunity*. Cambridge: University of Cambridge.

VERGANO, M., BERTOLINI, G., GIANNINI, A., GRISTINA, G. R., LIVIGNI, S., MISTRALETTI, G., RICCIONI, L. & PETRINI, F. (2020a): SIAARTI recommendations for the allocation of intensive care treatments in exceptional, resource-limited circumstances. In: *Minerva Anestesiologica*, 86(5), pp. 469–472.

VERGANO, M., BERTOLINI, G., GIANNINI, A., GRISTINA, G. R., LIVIGNI, S., MISTRALETTI, G., RICCIONI, L. & PETRINI, F. (2020b): Clinical ethics recommendations for the allocation of intensive care treatments in exceptional, resource-limited circumstances: The Italian perspective during the COVID-19 pandemic. In: *Critical Care*, 24(165), pp. 1–3.

WANG, H., WANG, Z., DONG, Y., CHANG, R., XU, C., YU, X., ZHANG, S., TSAMLAK, L., SHANG, M., HUANG, J. & WANG, Y. (2020): Phase-adjusted estimation of the number of coronavirus disease 2019 cases in Wuhan, China. In: *Cell discovery*, 6(1), pp. 1–8.