Student: Max Floris van Geuns
Supervisor: Max Nielsen
Course: Applied Environmental and Natural Resource Economics

How much effort should we devote to controlling the spread of COVID-19?

Abstract: This paper presents a theoretical framework of a cost-benefit analysis (CBA) in a lockdown following the COVID-19 pandemic, focusing on the value of a statistical life (VSL). While the use of this parameter is deemed appropriate in environmental economics, it remains uncertain whether and how it can be applied in the context of a COVID-19 lockdown. As this has consequences for determining the best policy, this paper answers the following research question: what are the implications of differences between the application of VSL in usual environmental projects and a lockdown following the COVID-19 pandemic? Reviewing sensitivity analyses and in contrast to the homogeneous application in recent literature, it seems to be more reasonable to let VSL vary with age: the costs of a lockdown fall disproportionally heavy on people below the age of 50, whereas most lives saved are in higher age categories. This study suggests a more appropriate VSL of below $10 million, which is lower than most values applied so far. How much lower it should be is yet to be determined by future research, although a lockdown might even require a new CBA framework based on total wellbeing. Regarding policy implications, some sort of 'intermittent' social distancing would be most justified, both economically and socially.

Keywords: COVID-19, VSL, CBA, lockdown, social distancing

UNIVERSITY OF COPENHAGEN

Table of Contents

1. Introduction

The global coronavirus (COVID-19) pandemic is one of the greatest public health threats in modern history. First detected in late November 2019 in Wuhan, China, it has now been reported in almost every country in the world. As of 7 June 2020, more than 7 million cases and over 400 thousand deaths have been confirmed (Worldometers.info, 2020).

The world has not seen a disease epidemic of the current scale since the 1918-1919 influenza pandemic, when there was also no vaccine readily available. The current response of public health authorities resembles that of global communities a hundred years ago: the implementation of non-pharmaceutical interventions (NPIs), aimed at reducing contact rates in the population and thereby reducing transmission of the virus.

The leading principle here is to save as many lives as possible; the costs are of secondary importance. However, bearing in mind that funds (and healthcare) are finite and the quality of life is profoundly important to many, this raises the question how much effort we should devote to controlling the spread of COVID-19. As rational as this question might seem, it is also controversial; before moving on, it should thus be emphasised that a more philosophical debate between utilitarianism and deontology is beyond the scope of this paper. Only the costs and benefits of a lockdown and the associated live(year)s saved will be assessed.

Ferguson et al. (2020) assessed two fundamental strategies of NPIs, mitigation and suppression, for the COVID-19 pandemic in the United Kingdom (UK) and the United States (US). The researchers show that while suppression might be effective in the short term, it is socially and economically cost-intensive, and any intervention package would need to be maintained until a vaccine becomes available (potentially 18 months or longer). The researchers conclude that mitigation with intermittent social distancing, focussing on slowing but not necessarily stopping epidemic spread, might be a more realistic strategy.

Using the Ferguson et al. (2020) simulation model of COVID-19's spread and mortality impacts in the US, Greenstone and Nigam (2020) developed and implemented a method to monetise the impact of moderate social distancing in the early stages of the pandemic. They projected that 3-4 months of distancing would save 1.7 million lives, resulting in mortality benefits of $8 trillion, or $60,000 per US household. Several other social planner problems and cost-benefit analyses have already been assessed for the COVID-19 lockdown in the US, resulting in similar results: Thunström et al. (2020) found net benefits of about $5.2 trillion and Wilson (2020) estimated the total costs of unconstrained spread of the virus at $8 to $60 trillion.

These studies have applied distinguishable methods to monetise the impacts of lockdown policies, but all of them used one critical parameter to assess the costs of extra casualties due to uncontrolled virus spreading or, rather, the benefits of social distancing: the value of reduced mortality risk, or the so-called value of a statistical life (VSL).

VSL terminology was first introduced by economist Thomas Schelling in his essay 'The Life You Save May Be Your Own' in 1968. Focussing on risk rather than individual lives, it calculates the value of the reduction of some mortal hazard to some part of the population. It has become a widely used concept in cost-benefit analyses by government agencies around the world to place monetary values on changes in premature deaths. For example, the estimation of benefits of highway traffic safety measures and reductions in air pollution rely on VSL applications (Banzhaf, 2014).

While the use of this parameter is deemed appropriate in environmental economics, it remains uncertain whether and how it can be applied in the context of a COVID-19 lockdown. As this might have serious consequences for determining which lockdown policy and mitigation measurements are most beneficial to society, this paper will answer the following research question: what are the implications of differences between the application and estimates of VSL in usual environmental projects and a lockdown following the COVID-19 pandemic? The answer to this question could provide us with a response to the earlier mentioned much simpler, yet more provoking question: how much effort should we devote to controlling the spread of COVID-19?

In addition to VSL, Quality-Adjusted Life Years (QALYs) are often mentioned in regular environmental policy analyses. The value of one QALY also considers someone's general health and shape; people who are in a critical condition due to a serious illness, for example, have a lower QALY than those who are fit and healthy. However, in most CBA's of the COVID-19 lockdown, these issues are not fully being accounted for. The hypothesis of this paper is that this results in potential overestimations of the VSL. The reason for this is that most people who die from COVID-19 are older, ill or unhealthy people, with lower QALY values.

This study first presents a simplified theoretical and mathematical framework of a cost-benefit analysis (CBA) in a moderate lockdown following the COVID-19 pandemic. After the application and estimates of VSL are explained in this context, the differences between the parameter in a lockdown and environmental projects are discussed. The absence of the application of Quality-Adjusted Life Years (QALY) in recent COVID-19 literature are also discussed and criticised. Lastly, this paper concludes with recommendations on the appropriate use of VSL in a CBA of the COVID-19 lockdown.

2. Methods: CBA Framework for a COVID-19 Lockdown

This paper analyses a number of studies that have already been published after the outbreak of the COVID-19 pandemic, so it does not just focus on one single method or model. However, most of the researchers who have been working on CBAs or lockdown planning in the context of COVID-19 do depart from the same earlier studies and their models. More specifically, the simple SIR model of the progression of COVID-19 in the US by Atkeson et al. (2020) and the impact estimations of NPIs to reduce COVID-19 mortality and healthcare demand by Ferguson et al. (2020) are among the most cited works in this newly established field.

The main model presenting the progression of any epidemic is the SIR model. This is a Markov model, assuming that future states in randomly changing systems depend only on the current state rather than on the events that occurred before. It describes the spread of the epidemic in a population that is divided into categories of being susceptible to the disease (S), actively infected with the disease (I), and recovered or dead and thus no longer contagious (R). How the disease evolves over time, is determined by the transition rates between S, I, and R; hence the name of the SIR model. It allows for quantitative statements regarding the trade-off between the severity and timing of suppression of the epidemic with social distancing, and the progression of the disease among the population (Atkeson, 2020).

In the model by Atkeson et al. (2020), which is based on the model presented in Wang et al. (2020), the population is set to N and normalised to one. All results should thus be interpreted in fractions of the relevant population. At each moment in time, N is divided into the four categories S, E, I and R, of which E stands for those exposed to the disease. These fractions of the population evolve over time and follow these equations:

$$dS/dt = -\beta_t(S/N)I \tag{1}$$

$$dE/dt = \beta_t(S/N)I - \sigma E \tag{2}$$

$$dI/dt = \sigma E - \gamma I \tag{3}$$

$$dR/dt = \gamma I \tag{4}$$

$$\beta_t = R_t\gamma \tag{5}$$

The logical reasoning behind these equations will now be explained. Parameter β_t is the rate at which infected people meet others and 'shed' the virus onto them. If this is applied to the people actively infected with disease (I) and the people being susceptible to the disease as a fraction from the total population (S/N), this will decrease the total amount of people susceptible to the disease (equation 1). The parameter σ governs the rate at which those exposed become infected, which reflects the estimated incubation period. When applying this to the people exposed (E) and subtracting both from the newly exposed people ($\beta_t(S/N)I$), this together increases with the total amount of exposed people (equation 2). Parameter γ governs the rate per day at which infected people recover or die, which reflects the estimated duration of illness. When applying this to the people infected (I) and subtracting that term from the people exposed that have been infected (σE), this together increases with the number of people infected (equation 3). The people infected who recover or die (γI) increases with the number of people who are no longer contagious (equation 4). Parameter R_t is the ratio of the 'meeting rate' β_t and the recovery/death-rate γ on day t, and governs the ratio of the rates at which those susceptible become infected and those ill either recover or die (equation 5). It varies over time and is controlled by social distancing and quarantine measures, as the remainder of this paper demonstrates (Atkeson et al., 2020).

The parameter R_0 stands for the transmission of the disease with no mitigation efforts and is critical for the economic costs of mitigation, hence of high importance for this paper. Different studies consider values of R_0 ranging from 2.2 (Fauci et al., 2020) to 3.25 (Remuzzi and Remuzzi, 2020). Strict social distancing and heavy mitigation measures have already shown to significantly reduce this transmission rate R_t from above 3 to close to or even under 1, whereas milder social distancing measures result in a reduction in transmission from 3.0 to 1.6 (Atkeson et al., 2020).

Another influential paper in this new field of planning problems following COVID-19 that has been cited by many authors already is Ferguson et al. (2020). They modified an individual-based simulation model to support pandemic influenza planning, exploring scenarios for COVID-19 in the UK. In their transmission model, some influential assumptions were made (as their model has been adopted by most of the other studies), the most important one being that R_0 is 2.4 (Ferguson et al., 2020).

Based on these assumptions, Ferguson et al. (2020) consider the impact of five different NPIs implemented individually and in combination: case isolation in the home, voluntary home quarantine, social distancing of those over 70 years of age, social distancing of the entire population, and closure of schools and universities. They find that without any control measures, a peak in mortality occurs after approximately 3 months and 81% of GB and US populations would become infected. This would result in approximately 510,000 casualties in GB and 2.2 million in the US, not accounting for the potential deaths following overwhelmed health systems (Ferguson et al., 2020).

However, if social distancing measures would be implemented, this would reduce *R* to below 1. The results of the study suggest that population-wide social distancing would have the largest impact. The researchers conclude that the minimum policy for effective suppression is combining this with other interventions, notably home isolation of cases and school and university closure, as this could suppress transmission below the threshold of $R = 1$, which would be required to rapidly reduce case incidence. This is also the moderate lockdown policy with social distancing that most authors base their analysis on (hence also this paper). Lastly, as this policy might have to be maintained for 18 months or more (until a vaccine becomes available), they do emphasise that "it is not at all certain that suppression will succeed long term; no public health intervention with such disruptive effects on society has been previously attempted for such a long duration of time" (Ferguson et al., 2020).

Analysing both the simulations of the SIR model and the detailed analysis in Ferguson et al. (2020), Atkeson et al. (2020) concludes the following:

> *(…) paint a grim picture of the choices regarding public health that policymakers face in mitigating the impact of the COVID-19 pandemic. What is urgently needed is an economic analysis of the economic consequences of the mitigation steps currently being implemented and contemplated going forward so that economic tradeoffs between public health and the economy can be considered quantitatively.*

This analysis Atkeson et al. (2020) asked for has been carried out by a number of authors, among whom are Thunström et al. (2020), Greenstone and Nigam (2020) and Layard et al. (2020). The remainder of this methodology section will briefly touch upon some of their findings and adjustments of the model.

The research conducted by Thunström et al. (2020) resembles the SIR model provided by Atkeson et al. (2020), but worked with equations of motion that focus on the total number of people in the population (instead of fractions). Also, they included a term that accounts for the fact that some infected people die, hence these are removed from the infected people. To account for the possibility of overwhelming the health care system, the researchers made a critical assumption that "the system has sufficient resources to provide adequate treatment for one half of the maximum number of individuals who would be infected at any one time in an uncontrolled scenario, with no social distancing to slow the spread of the virus" (Thunström et al., 2020). Regarding the benefits of a lockdown, they took the VSL to be $10 million, consistent with US federal agency guidelines. The costs to social distancing were measured as lost GDP, assuming that it would have grown at a constant rate of 1.75% (CBO, 2020). They assume that for the controlled scenario, GDP will fall by 6.2% in 2020, after which it grows 5.5%, 3.5%, 2.0%, and 1.75% (continued) in the years after (OECD, 2003). For the uncontrolled scenario, they choose an immediate decline of 2.0% for their benchmark case. Finally, the net present value of social distancing is written as the VSL of the lives saved, minus the present value of GDP lost:

$$NPV = VSL \times (D_1 - D_2) - \sum(Y_{1t} - Y_{2t})\ (1+r)^{-t} \quad (6)$$

Here, D_1 is the number of deaths without social distancing and D_2 with social distancing (estimated with the SIR model), Y_{1t} is the forecasted level of GDP in year t without social distancing and Y_{2t} with social distancing, r is the discount rate, and T is the planning horizon. With a discount rate of 3% per year and a planning horizon of 30 years for the benchmark case, the result is that the benefits of social distancing are $12.4 trillion (1.24 million lives saved with a VSL of $10 million), whereas the costs are $7.21 trillion (GDP losses of $13.7 trillion with social distancing, minus GDP losses of $6.49 trillion without the policy). Their main result is that social distancing generates net benefits of $5.16 trillion. An important assumption is that the measures buy enough time to develop and distribute COVID-19 treatments or vaccines, should a second wave occur. Also, the sensitivity analysis of Thunström et al. (2020) shows some interesting and important results regarding VSL, which will be discussed in the results section of this paper (Thunström et al., 2020).

A novel feature was added by Greenstone and Nigam (2020). They estimated not only the direct deaths of COVID-19, but also accounted for the potential shortages in the supply of hospital intensive care services. According to their analysis, social distancing reduces COVID-19 caused deaths by 1.76 million, of which 630,000 are ‘overflow' deaths because of denied patients. Also, Greenstone and Nigam (2020) allowed the VSL to vary with age, but required the average for people 18 and over to equal $11.5 million. As a result of the lower VSL of higher age categories (among which there are more deaths caused by COVID-19), their conclusion is that the reduction in COVID-19 caused deaths by 1 October is worth only $7.9 trillion; less than the benefits estimated by Thunström et al., yet still more than their estimated costs. This estimation of the costs of a lockdown was not made by Greenstone and Nigam (2020) themselves, and is also beyond the scope of this paper.

One last interesting contribution that is important to discuss regarding the CBA model of a lockdown, before going into discussing the individual parameter of VSL, is the wellbeing framework provided by Layard et al. (2020). It is based on evaluating all changes in impact on individuals' satisfaction with their lives, in terms of Wellbeing-Years (WELLBYs). Based on this metric, the researchers estimated the net benefits of releasing the lockdown in the UK on a stated date rather than one month later. This CBA framework goes beyond other COVID-19 literature, as it does not merely consider the value of lives saved and losses in GDP. Their full framework and estimates is shown below (Layard et al., 2020).

Lockdown release date	**May 1**	**June 1**	**July 1**	**August 1**
Benefits				
Income (increase)	48	66	86	103
Unemployment (reduction)	79	82	84	86
Mental health (improvement)	20	23	26	30
Confidence in government	9	13	22	32
Schooling (more)	5	5	3	0
Costs				
COVID-19 deaths (more)	-158	-158	-158	-158
Road deaths (more)	-5	-5	-5	-5
Commuting (more)	-10	-10	-10	-10
CO_2 emissions (more)	-7	-7	-7	-7
Air quality (worse)	-8	-8	-8	-8
Net benefits	**-27**	**1**	**33**	**63**

Table 1: net benefits compared to releasing the UK lockdown one month later, in terms of 10,000 WELLBYs (Layard et al., 2020).

Clearly, as it is purely illustrative and subject to great uncertainty, this framework is hard to compare with the CBAs that work with monetary values instead of WELLBYs. Nevertheless, its variables seem to be more all-encompassing (as it also takes other important costs and benefits into account), and the result suggests that it might be more beneficial to society if the lockdown is (partly) released at some point, even in the absence of a medicine or vaccine. Also here, assumptions about the value of life play an important role and will be analysed and critically discussed in the following chapters. A summary of the full methodology applied in this all-encompassing CBA framework is presented on the next page.

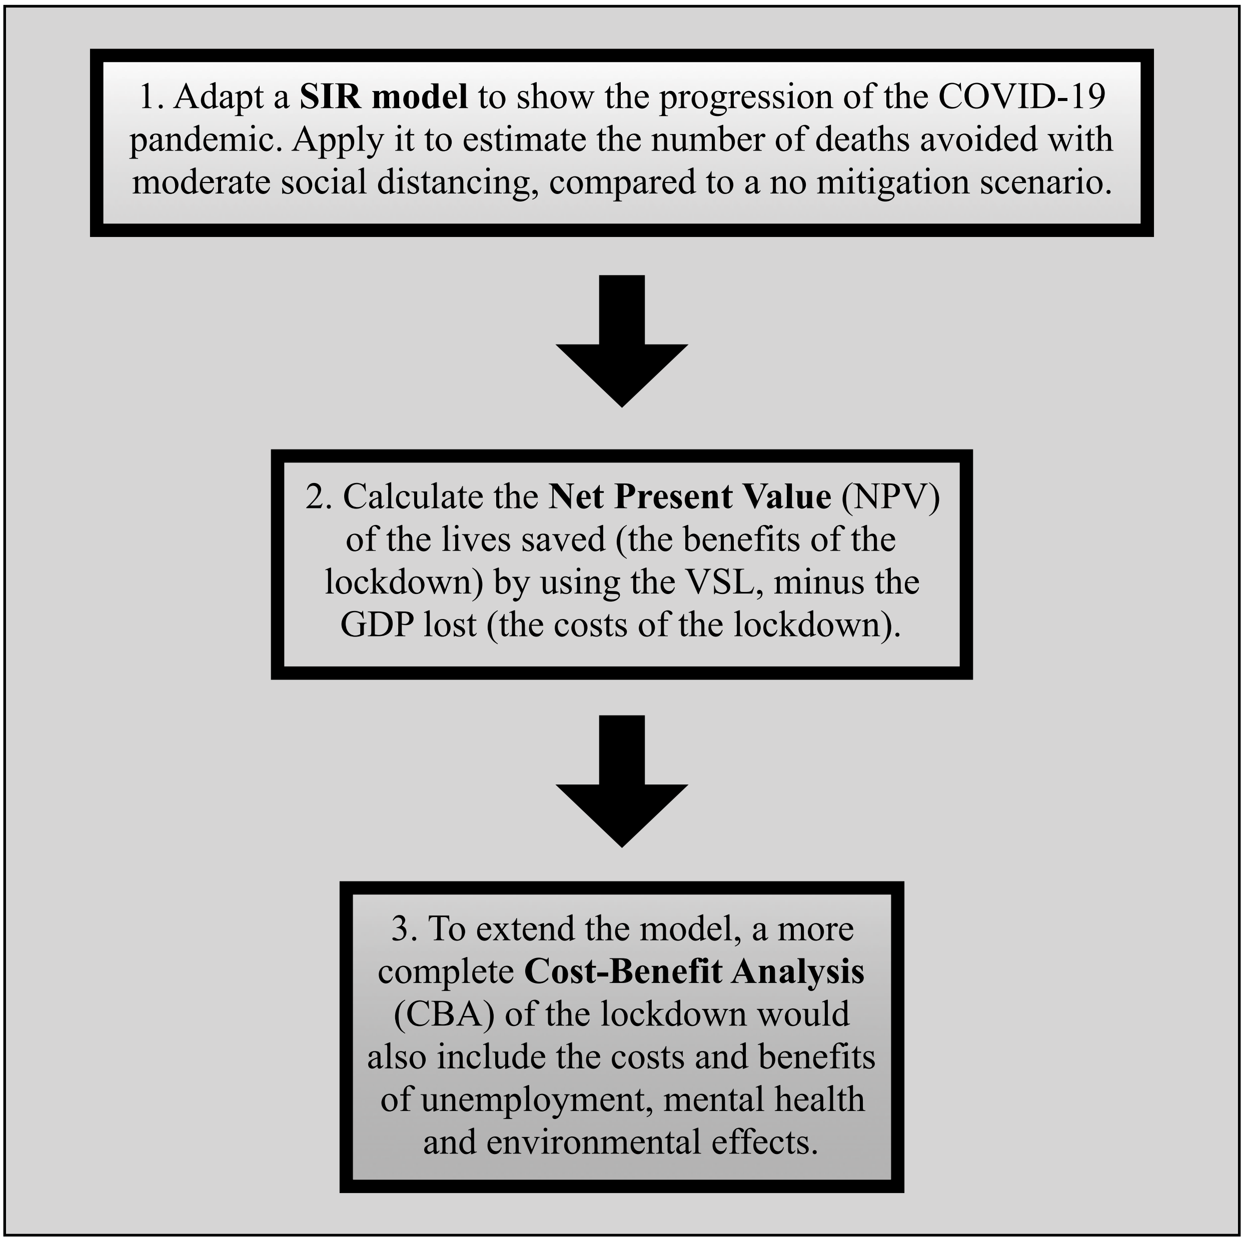

Figure 1: the methodology of a CBA framework of the COVID-19 lockdown in 3 steps.

3. Results: the Value of a Statistical Life (VSL)

The results chapter of this paper focuses on the Value of a Statistical Life (VSL) within the CBA framework of a lockdown following the COVID-19 pandemic. However, first it should be emphasised that this parameter assesses *statistical* and not *concrete* lives. Ask any doctor about the life of someone in particular, and he or she will always make sure to do everything needed to save that life. Treating the ill to the best of one's ability is the Hippocratic Oath all physicians are legally bound to, but most other people would also want to do anything to save the lives of the ones they love, or even know. It should thus be clear that VSL does not indicate 'who shall live and who shall die'. Rather, it is an instrument to calculate which policies and projects are worth the investment to reduce the risk of dying — and nobody can say that we cannot have any risks, because we take them every day when doing sports, driving our cars or even crossing a road by foot. However, we do reduce those risks by wearing certain equipment, implementing traffic laws and constructing safe crossings.

Before going into the details and sensitivity analyses of VSL in the different papers reviewed, some general history and background information on VSL, as well as its basic theory and mathematics, are vital for a proper understanding of this topic.

3.1 History and background theory

As it is the aim of many environmental projects to enhance the health of individuals, policymakers need some measure to monetise these consequential reductions in mortality risks and to facilitate a comparison of all benefits and costs associated. Before, economists adapted the human capital approach, calculating the present value of the loss in income and medical costs associated with the death of an individual. Apart from the ethical question whether people with higher incomes are worth more than poor people, this approach becomes problematic when estimating the value of lives of very young, unemployed, and elderly people (as the present value of their life income is zero, or unknown). Also, it fails to capture many features of what people value about their life, such as their consumption of non-market goods (Greenstone and Nigam, 2020). Focussing on trade-off rates between fatality risk and money became a more accepted economic approach, as introduced by Schelling (1968). While he laid the foundation of the theory, it was not until the 1980s that empirical data led to reliable estimates of the VSL. Since then, the two principal ways of realising this have been revealed preferences, based on (labour) market evidence, and stated preferences, based on survey data with hypothetical decisions rather than actual behaviour (Kniesner and Viscusi, 2019).

Stated preference studies are particularly useful for countries in which less (suitable) risk and employment data is available and such a country is not within the scope of this paper (as the US has a wide range of data available), so only revealed preference evidence will be elaborated upon. The basic theory of the revealed VSL is that more dangerous or hazardous jobs are less desirable to obtain and thus require higher wages, *ceteris paribus*. Likewise, safer jobs are more expensive for employers to offer, so they will have to pay lower wages to make this possible (while employees accept this lower wage in a safer working environment). Economists call this collection of wage and health risk combinations the 'hedonic equilibrium locus' and describe it by the equation $w = w(p)$, where w stands for hourly wage and p is the probability of a fatal health hazard exposure at work. To account for the fact that less safe jobs pay a higher wage, $w' = dw/dp > 0$. The VSL is derived from the value of the hedonic wage locus' slope: if one needs to be paid $1,000 more per year to accept a job where there is one more casualty per 10,000 workers every year, a collective of 10,000 employees would earn $10,000,000 more as a group, hence these additional wage payments comprise the VSL. Furthermore, the hedonic locus generally is nonlinear so that $w'' = (d^2w)/(dp^2) \neq 0$; employers with the lowest marginal cost of providing safety hire employees who value this the most, whereas firms with higher costs to improve safety hire workers who are willing to take risks (Kniesner and Viscusi, 2019).

Based on this revealed preference method in the labor market, an estimation of the VSL is derived from a multivariate regression that is typically referred to as the Mincer equation:

$$ln(w_{i,j,k}) = \alpha + \beta p_{i,j,k} + X_{i,j,k}\Gamma + u_{i,j,k} \qquad (7)$$

Here, for worker i in industry j and occupation k the dependent variable is the natural logarithm of the real wage, p stands for the work-related fatality rate, X is a vector with demographic control and job characteristic variables, Γ is a coefficient vector, and $u_{i,j,k}$ is an error term. In this wage equation, the estimates are equal to the proportionate change in the wage per unit change in the independent variables, which means that $\beta = (\partial w/\partial p)(1/w)$. The typical US fatality risk measure is deaths per 100,000 employees, so the VSL with a typical work year of h hours is calculated as follows:

$$VSL = (\beta \times w) \times h \times 100{,}000 \qquad (8)$$

Using this formula, the mean VSL for US men has been estimated at $7-8 million in $2001 (the dollar value in 2001, accounting for inflation). Even though this estimate has been growing due to inflation and it depends on the population sample, the risk measure, and the equation specification, its "general order of magnitude has remained quite stable" (Kniesner and Viscusi, 2019). A more recent study in the US, including 1,025 VSL estimates from 68 publications, resulted in a mean of $12 million and a median of $9.7 million in $2015 (Viscusi, 2018).

Above, the VSL has been treated as being homogenous rather than heterogenous. However, it is likely to vary across several individual characteristics, the most important ones being income and age. To account for the possibility that higher income workers are less sensitive to costs of a product (or job), a positive income elasticity can be added to the multivariate regression. To account for the fact that different groups value riskiness and costs differently, the risk variable β can be interacted with an age-group indicator, showing to which cohort an individual belongs (O'Brien, 2017).

O'Brien (2017) estimated a VSL with these additions to the framework by examining the risk-income tradeoff that individuals face in their automobile purchase decisions, rather than in their labor market decisions. By comparing the relative importance of automobile costs and safety for various age groups, he generated VSL estimates for seven age cohorts ranging from the age of 18 up to the age of 85. By calculating a separate willingness to pay for reduced mortality for individuals in these different age groups, he found a significant inverted-U shape to the age-VSL function that ranges from $1.5 million (for the ages 25-34) to $19.2 million (ages 55-64), back to $8.15 million (ages 75-85) in $2009. The overall weighted average was $9.24, with the age cohorts of 45-74 above and the others below this average (O'Brien, 2017).

3.2 Estimates and sensitivity analyses

The remainder of this chapter discusses the VSL estimates and sensitivity analyses of the considered research conducted on COVID-19. As stated earlier, Thunström et al. (2020) took the VSL to be $10 million, consistent with the US Federal Agency guidelines. However, they acknowledge that the death rate due to COVID-19 is substantially higher for older people, indicating that age is an important factor to take into account in a sensitivity analysis. According to their model, if the average VSL would fall below $5.85 million due to a lower VSL that is appropriate for older individuals, their benchmark result would change. Hence, this is their benchmark value for which the benefits of socials distancing measures would not be worth the costs anymore (Thunström et al., 2020). Whether this is truly the case is beyond the scope of this paper, because it would require an elaborate assessment of the costs within the CBA as well.

Another study that did not distinguish between age and VSL is Wilson (2020), although he does seem to criticise his own approach: “A 90-year-old man with a life expectancy of 4.1 years has the same VSL as a 1-year-old girl with a life expectancy of 80.4 years in our analysis” (Wilson, 2020).

One of the studies that did use age-varying values of reduced mortality risk is Greenstone and Nigam (2020). Similar to O'Brien (2017), they divided the reduction in fatalities from the mitigation scenario into 9 age categories. However, they required the average for people 18 and over to be equal to $11.5 million, matching the US Government's VSL for adults. Because of this, the benefits of social distancing measures are still worth the costs in their CBA, but the impacts are “strikingly heterogeneous” across age categories: all people under the age of 50 only get 11 percent of the total benefits of $7.9 trillion. When using an updated version of the VSL estimated by Ashenfelter and Greenstone (2004), the total social benefits of Greenstone and Nigam (2020) are much lower ($2.4 trillion), demonstrating the importance of assumptions about the VSL (Greenstone and Nigam, 2020).

Hall, Jones and Klenow (2020) used a substantially lower VSL, derived from the US Environmental Protection Agency's recommendation of $7.4 million in $2006, equivalent to $9.5 million today. Apart from a higher VSL of $11.4 million, they also included a lower VSL of $7.6 million in their sensitivity analysis, adding that “those at greatest risk of dying from COVID-19 are those with the fewest years of life remaining. This means fewer life years are at stake than if the virus struck down all ages with the same probability or was particularly lethal for younger people” (Hall, Jones and Klenow, 2020).

Following this reasoning, Alvarez, Argente and Lippi (2020) also took into account that the majority of the victims of the virus have a below average life expectancy. They took the VSL to be 20 times annual GDP per capita, which would only be about $1.3 million using the World Bank estimate of $63,000 in 2018. For this parameter, they estimated that the optimal lockdown level is to peak at 60% (of the population in lockdown) about 10 days after it starts and to decrease to below 10% by the sixteenth week. Their optimal lockdown duration is about 4 months. In other words: according to their CBA and these parameters, it would not be worth it to have a more severe or longer lockdown (Alvarez, Argente and Lippi, 2020).

Study	Mean VSL	Lower VSL	Higher VSL
Thunström et al. (2020)	$10 million	None	None
Wilson (2020)	$11.2 million	$7 million	$15.4 million
Greenstone and Nigam (2020)	$11.5 million	$3.5 million	None
Hall, Jones and Klenow (2020)	$9.5 million	$7.6 million	$11.4 million
Alvarez, Argente and Lippi (2020)	$1.26 million	$630,000	$8.82 million

Table 2: an overview of the mean VSL used in the different studies discussed, as well as the lower and higher VSL estimates in the sensitivity analyses of these studies.

In general, it thus seems more reasonable to let VSL vary with age, rather than to give the same value to every individual regardless of their life expectancy. This is particularly important in the context of COVID-19, as the costs of a lockdown fall disproportionally heavy on people below the age of 50, whereas most lives saved are in higher age categories (with a lower life expectancy). This suggests that the lower ranges of VSL considered (below $10 million) are more appropriate to apply to a CBA of a COVID-19 lockdown than the higher ranges (above $10 million). However, it is still uncertain if it should be anywhere near this initial $10 million estimate of Thunström et al. (2020), or preferably closer to the extremely low estimate of $1.26 million by Alvarez, Argente and Lippi (2020).

Lastly, to conclude this chapter, an important criticism of VSL is that it does not take into account quality of life (apart from the age-varying life expectancy). Most people would not value a year of life recovering from an ICU facility as much as a year of life in excellent health (Conover, 2020). It might thus be more preferable to estimate values in any healthcare CBA based on Quality Adjusted Life Years (QALYs), allowing the value of life to vary based on subjective assessments of the quality of life. One study that took this into account was conducted in the Netherlands. The researchers of Gupta Strategists found that the Dutch COVID-19 care response has led to a gain of 13 to 21 thousand healthy life years, but the loss of healthy life years due to the reduction in regular care alone is estimated to be 100 to 400 thousand. They calculated that the financial burden is equally disproportional: the costs per healthy life year gained is estimated to be 100 to 250 thousand euros, which is much higher than the upper cost-effectiveness reference norm of 80 thousand euros used by Dutch authorities (Gupta Strategists, 2020).

Commonly used QALY values in the UK are in the range of $24,300 to $36,500 (converted from British Pounds with the conversion rates at the time of writing), but even when applying the higher end of the range of $50,000 to $150,000 that is used in the US, the value of a life of a 70-year old with a remaining life expectancy of 15 years is only $2.25 million, which is still much lower than Greenstone and Nigam (2020) assign to a life of someone aged 70-79 (Freeman and Groom, 2020). Taking into account the quality of life, it thus seems more reasonable to apply a VSL closer to $1 million than to $10 million, but more research on QALYs in the context of COVID-19 is needed to assess this.

4. Conclusion

This paper presents a simplified theoretical framework of a CBA in a moderate lockdown following the COVID-19 pandemic. First, a simple SIR model of the progression of the disease in the US is applied to a simulation model with impact estimations of NPIs that reduce mortality and healthcare demand, calculating the number of deaths avoided with moderate social distancing compared to a no mitigation scenario. Second, the net present value of the lives saved minus the GDP lost (the costs of the lockdown) is calculated, even though an elaborate assessment of these costs is not within the scope of this paper. Third, to extend the model, a more complete CBA of the lockdown would also include the costs and benefits of unemployment, mental health and environmental effects.

Within this framework, the study focuses on the Value of a Statistical Life (VSL), a measure that has been widely used by economists and policymakers to monetise reductions in mortality risks and to facilitate a comparison of all benefits and costs associated with an environmental project or a health policy. Derived from a multivariate regression of wages, the mean VSL in the US has recently been estimated between $9 million and $12 million, depending on whether characteristics such as income and age are taken into account.

Estimates and sensitivity analyses of the research conducted on the VSL in a lockdown following the COVID-19 pandemic indicates that it is more reasonable to let VSL vary with age: the costs of a lockdown fall disproportionally heavy on people below the age of 50, whereas most lives saved are in higher age categories. This suggests that it is more appropriate to apply a VSL of below $10 million to the CBA. How much lower it should be is still to be determined by future research. Lastly, the most important criticism of VSL in general is that it does not take into account quality of life, such as QALYs do.

5. Discussion

Can we now conclude that VSL estimates in CBA's of a COVID-19 lockdown have been too high? Not necessarily, because some important yet uncertain features (that are beyond the scope of this paper) might actually ask for a larger rather than a smaller VSL estimate. For example, people could have a high willingness-to-pay (WTP) for a smaller infection risk, because people are farsighted, paying particular attention to the possibility of vaccine discovery, and because they put extra weight on worst-case scenarios (Echazu and Nocetti, 2020). Also, evidence shows that VSL is increasing for non-marginal changes in fatality risk and it seems plausible that the changes considered here are non-marginal, meaning that a larger VSL should be used (Greenberg et al., 2020). On the other hand, there is also more evidence that a lower VSL is more appropriate. For example, the median wrongful death jury award is only $2.6 million, the median 9/11 compensation was $2.4 million, and the average life-time earnings of college graduates is $2.8 million (Wilson, 2020).

Assuming that the VSL for a COVID-19 lockdown should be lower than applied so far, will this change the conclusion of most authors that the social distancing measures are economically justified? Also this cannot be concluded after merely changing the VSL; pandemics are "characterised by a great deal of uncertainty related to epidemiological features, health-care constraints, the timing of vaccine discovery, and others" (Echazu and Nocetti, 2020). For more certainty about the correct CBA, sensitivity analyses should also be conducted for each of the other assumptions in the model, including the lockdown period and effectiveness, congestion of the healthcare system, testing of the recovered, welfare costs under the optimal policy, and the applied discount rate (Alvarez, Argente and Lippi, 2020). Furthermore, there is the potential for social distancing to reduce the rates of non-fatal sickness experienced by the population, adding more benefits to the CBA (Greenstone and Nigam, 2020). All of these uncertainties could bring us back to the initial conclusion of most researchers, as put in the following words by Echazu and Nocetti (2020):

> *There are arguably many other costs beyond reduced economic activity, but to the extent that economic restrictions are not extended for a long period, all measures indicate that the costs will be below the expected benefits in terms of lives saved and disease costs avoided — at least on an economic/statistical sense.*

Finally, we return to the provoking title of this paper: how much effort should we devote to controlling COVID-19? As a thought experiment (rather than making hard conclusions based on actual numbers), we simply assume that the CBA will result in no net benefits of social distancing because of the lower VSL. With pure economic reasoning, this would indicate that the benefits of social distancing are not worth the costs and the lockdown should thus be lifted. High unemployment will have disastrous consequences and it is the responsibility of a government to avoid this. In addition, lives will be hurt or shortened if economic conditions deteriorate, so social and economic distress also translates into the loss of human lives. However, the CBA approach fundamentally clashes with our moral intuitions and other ethical standards, as well explained by Colonomos (2020):

> *It also undermines values such as dignity and lacks robustness in logical terms because it evades the issue of political responsibility. Increasing national wealth is not the primary objective of a state; ensuring the security of its population is. In the logic of political responsibility, life must be prioritised over money.*

Nevertheless, even if the lockdown measures are justified, most of its benefits evidently go to a relatively small and older group in society, whereas the burden of the costs falls disproportionally on the younger population. This implies large redistributions of income and quality of life from healthy to vulnerable people (Echazu and Nocetti, 2020). It could thus be advocated that this needs to be undone with other redistribution policies. What might be helpful in this respect is so-called 'intermittent' social distancing, as proposed by Ferguson et al. (2020). This may allow interventions to be relaxed temporarily, with the measures reintroduced if case numbers rebound. Such a 'partial deconfinement' with room for back-and-forth movements will probably even become a necessary solution, given that it could last a year or more until a medicine or vaccine becomes available (Colonomos, 2020).

Regarding suggestions for further research, especially more empirical analysis is needed to evaluate people's WTP to reduce morbidity and mortality risks during a pandemic. Other potentially important consequences should also be considered, such as "the associated decline in levels of pollution, a possible spike in domestic violence and suicides, and other non-market social impacts" (Thunström et al., 2020). According to Freeman and Groom (2020), it is even likely that neither VSL nor QALY are even applicable to COVID-19:

> *The VSL is generally calibrated using people's stated willingness to pay to avoid a small increase in the risk of death. The much lower estimates of a cost/QALY ratio in the UK reflects*

the observed health effects of relatively small changes in NHS expenditure over recent years, rather than the ethical question of what society 'ought' to pay to improve health. Given the scale of the impact of the pandemic on how the NHS functions and the decision to dramatically increase public expenditure these estimates are unlikely to be a useful guide in current circumstances. The world we now find ourselves in is entirely new, and calls for new methods in healthcare CBA.

The WELLBY-framework proposed by Layard et al. (2020), as presented in table 1 in the methods section, might eventually be the most all-encompassing and correct CBA in this situation. As their numbers are purely illustrative, the last suggestion for further research is to develop and refine the WELLBY-model, so that policymakers can adopt this approach in reaching their decision on when to end the lockdown.

This research has been conducted within the MSc Environmental and Natural Resource Economics at the University of Copenhagen. Although VSL is especially used in health economics, the importance for the field of environmental economics should finally be highlighted as well. For every environmental project or policy that involves a reduction in mortality risk, such as air quality improvements and climate change mitigation, the VSL facilitates a CBA to estimate whether the benefits outweigh the costs. Even if not directly used, VSL estimates from different situations (such as a lockdown following the COVID-19 pandemic) contributes to the wider literature on VSL in environmental economics.

References

Alvarez, F., Argente, D. and Lippi, F. (2020). A Simple Planning Problem for COVID-19 Lockdown. University of Chicago, Becker Friedman Institute for Economics. *Working Paper No. 2020-34*.

Ashenfelter, O. and Greenstone, M. (2004). "Using Mandated Speed Limits to Measure the Value of a Statistical Life." *Journal of Political Economy, 112*(S1), 226-267.

Atkeson, A. et al. (2020). "What will be the economic impact of COVID-19 in the US? Rough estimates of disease scenarios." *Staff Report 595*, Federal Reserve bank of Minneapolis.

Banzhaf, H. S. (2014). "Retrospectives: The Cold-War Origins of the Value of Statistical Life." *Journal of Economic Perspectives*, 28 (4): 213-26.

Colonomos, A. (2020). "Pricing Lives in the Time of COVID-19." Retrieved from www.globalpolicyjournal.com/blog/04/05/2020/pricing-lives-time-covid-19

Congressional Budget Office (CBO) (2020). "The budget and economic outlook: 2020 to 2030."

Conover, C. (2020). "How Economists Calculate The Costs And Benefits Of COVID-19 Lockdowns." Retrieved from www.forbes.com/sites/theapothecary/2020/03/27/how-economists-calculate-the-costs-and-benefits-of-covid-19-lockdowns/#58836256f630

Echazu, L. and Nocetti, D. (2020). Willingness to Pay for Morbidity and Mortality Risk-Reductions during An Epidemic. Theory and Preliminary Evidence from COVID-19.

Fauci, A., Lane, C., and Redfield, R. (2020). "Covid-19 — navigating the uncharted." *New England Journal of Medicine*.

Ferguson, N. M. et al. (2020). "Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand." London: Imperial College COVID-19 Response Team.

Freeman, M. and Groom, B. (2020). "The Cold Equations - Covid-19: When The Value Of Statistical Life Changes All Our Lives." Retrieved from www.longfinance.net/news/pamphleteers/cold-equations-covid-19-when-value-statistical-life-changes-all-our-lives/

Greenberg, K., Greenstone, M., Ryan, S., and Yankovich, M. (2020). "The Value of Statistical Life: Evidence from Military Retention Incentives and Occupation-Specific Mortality." Mimeograph.

Greenstone, M. and Nigam, V. March, (2020). "Does Social Distancing Matter?" *University of Chicago, Becker Friedman Institute of Economics.*

Gupta Strategists (2020). "COVID goes Cuckoo: How the March-April 2020 COVID-19 surge overwhelmed Dutch hospitals and undermined regular care."

Hall, R. E., Jones, C. I. and Klenow, P. J. (2020). "Trading Off Consumption and COVID-19 Deaths." Mimeo, Stanford University.

Kniesner, T. J. and Viscusi, W. K. (2019). "The Value of a Statistical Life." *Oxford Research Encyclopedia of Economics and Finance*; *Vanderbilt Law Research Paper No. 19-15.*

Layard, R. et al. (2020). "When to Release the Lockdown? A Wellbeing Framework for Analysing Costs and Benefits." IZA Institute of Labor Economics.

OECD (2003). The World Economy: Historical Statistics, OECD Development Centre.

O'Brien, J. (2017). "Age, autos, and the value of a statistical life." *Journal of Risk and Uncertainty. 57.*

Remuzzi, A. and Remuzzi, G. (2020). "Covid-19 and Italy: what next?" *The Lancet*.

Schelling, T. C. (1968). "The life you save may be your own." In *S. B. Chase, Jr. (Ed.), Problems in public expenditure analysis*. Washington, DC: Brookings Institution.

Thunström, L., Newbold, S., Finnoff, D., Ashworth, M. and Shogren, J. (2020). "The Benefits and Costs of Flattening the Curve for COVID-19." *Journal of Benefit-Cost Analysis.*

Viscusi, W. K. (2018). Best estimate selection bias of estimates of the value of a statistical life, *Journal of Benefit-Cost Analysis*, *9(2), 205–246.*

Wang, H. et al. (2020). "Phase-adjusted estimation of the number of coronavirus disease 2019 cases in Wuhan, China." *Cell Discovery, 6(1):10.*

Wilson, L. (2020). "Estimating the Life Expectancy and Value of a Statistical Life (VSL) Losses from COVID-19 Infections in the United States." *University of Louisiana*.

Worldometers.info (2020). Coronavirus Update (Live). Retrieved on 7 June 2020 from https://www.worldometers.info/coronavirus

Reflections

Before starting the course Applied Environmental and Natural Resource Economics at the University of Copenhagen, my initial idea was to conduct research on the 'true price' of flight tickets or meat; I wanted to estimate an appropriate tax on these products with high negative environmental externalities. But when the coronacrisis hit, I could not get my head around anything else than COVID-19. Also, I had just finished the course Economic Valuation Methods (one of my favourite subjects so far), in which we briefly discussed the topic of the Value of a Statistical Life (VSL). Both of these events, as well as media coverage and quite some new scientific literature, triggered me and surprisingly led to my new paper topic on the appropriate use of VSL in a COVID-19 lockdown.

The moment I proposed my topic to the course teachers, I have to admit that I was first a bit disappointed in their response. One of them was pretty enthusiastic, but another teacher was quite critical and pointed out the fact that my topic was not suitable for Applied Environmental and Natural Resource Economics. Nevertheless, I took this criticism very serious, and tried to incorporate this. We e-mailed back and forth about the topic selection, and I adjusted my proposal bit by bit. After sending in my final proposal, I was even more surprised when that critical teacher became my own supervisor: Max Nielsen. It soon became clear that my worries about my topic selection and his criticism were unnecessary. In the end, I only benefited from the fact that 'the critical teacher' — which, in hindsight, was exaggerated from my part — became my supervisor. His guidance was just right: committed and critical, but also very considerate and constructive.

Regarding my own research and writing process, if I do say so myself, I am quite satisfied. I already collected the right reading list at an early stage, and I have made a tight but realistic time scheme that I managed to follow almost perfectly. The only adjustment I made along the way, was that I initially also wanted to analyse and conduct sensitivity analysis on a second parameter: the discount rate. However, it became clear that this was not only too ambitious regarding time management, but also (the prove for) my hypothesis about this parameter was not strong enough. Luckily, my supervisor and I concluded this soon enough, so I could still leave that part out. This would be something to learn from for my master thesis next year: not aiming too high and having a realistic research question to answer, from the beginning on. So, maybe I will have to do something less ambitious than determining 'the true price of meat' or 'the appropriate tax on flight tickets' after all…