What Should We Spend to Save Lives in a Pandemic?

A Critique of the Value of Statistical Life

Matthew D. Adler,[1] Duke University

June 25, 2020

Abstract: The value of statistical life (VSL) is a risk-to-money conversion factor that can be used to accurately approximate an individual's willingness-to-pay for a small change in fatality risk. If an individual's VSL is (say) $7 million, then she will be willing to pay approximately $7 for a 1-in-1-million risk reduction, $70 for a 1-in-100,000 risk reduction, and so forth.

VSL has played a central role in the rapidly emerging economics literature about COVID-19. Many papers use VSL to assign a monetary value to the lifesaving benefits of social-distancing policies, so as to balance those benefits against lost income and other policy costs. This is not surprising, since VSL (known in the U.K. as "VPF": value of a prevented fatality) has been a key tool in governmental cost-benefit analysis for decades and is well established among economists.

Despite its familiarity, VSL is a flawed tool for analyzing social-distancing policy—and risk regulation more generally. The standard justification for cost-benefit analysis appeals to Kaldor-Hicks efficiency (potential Pareto superiority). But VSL is only an approximation to individual willingness to pay, which may become quite inaccurate for policies that mitigate large risks (such as the risks posed by COVID-19)—and thus can recommend policies that fail the Kaldor-Hicks test.

This paper uses a simulation model of social-distancing policy to illustrate the deficiencies of VSL. I criticize VSL-based cost-benefit analysis from a number of angles. Its recommendations with respect to social distancing deviate dramatically from the recommendations of a utilitarian or prioritarian social welfare function. In the model here, it does indeed diverge from Kaldor-Hicks efficiency. And its relative valuation of risks and financial costs among groups differentiated by age and income lacks intuitive support. Economists writing about COVID-19 need to reconsider using VSL.

1

I.   Introduction

Fatality risk regulation is the bread and butter of the U.S. regulatory state. Administrative agencies in the U.S. government are required to prepare written documents accompanying proposed major rules that evaluate the proposed rules using cost-benefit analysis (CBA)[2]; and the monetized benefit of fatality risk reduction, as per these documents, is the largest category of monetized benefit.[3]

In particular, fatality risk reduction is central to the mission of the Environmental Protection Agency (EPA). Anti-pollution rules enacted by this agency, over the last five decades, have yielded large decreases in individuals' annual fatality risks; and much of the monetized fatality-risk-reduction benefit from federal regulation is attributable to the EPA.[4] But other federal agencies also enact regulations that reduce fatality risks and evaluate these regulations by means of CBA, including a monetization of the risk-reduction benefit. These include the Department of Transportation, the Occupational Safety and Health Administration, the Food and Drug Administration, the Department of Health and Human Services, and others.[5]

In general, CBA evaluates governmental policy by predicting, to the extent feasible, a policy's effects on the components of individual well-being (income, health, fatality risk, environmental quality, leisure, etc.) and then measuring these effects on a monetary scale.[6] The value of statistical life (VSL) is, in turn, the linchpin concept for monetizing risk reduction.[7] VSL can be thought of as a conversion factor for translating an individual's risk change into a monetary equivalent. Assume Suzy's VSL is $7 million. Then Suzy's willingness to pay for a small risk reduction is approximately the reduction multiplied by $7 million (so that she is willing to pay $7 for a 1-in-1 million reduction, $70 for a 1-in-100,000 reduction, and $700 for a 1-in-10,000 reduction).

Although VSL is, in principle, heterogeneous—Suzy's VSL might well be different from Jana's or Wu's—the practice in the U.S. government is to use a population-average VSL, now on the order of $10 million.[8] Note that the assumption of a single VSL leads to a simple formula for valuing lifesaving. If a policy reduces expected deaths in the U.S. population next year by ΔD, then the sum total of individual risk reduction is ΔD; and if a single VSL is being used, the aggregate willingness to pay for the policy's risk-reduction benefit is just $\Delta D \times$ VSL. This calculation now occurs in numerous agency cost-benefit documents.

---

[2] Executive Order No. 12,866, 3 C.F.R. 638 (1994).
[3] Office of Management and Budget (2017, p. 11).
[4] *Id.* at 10-16; Robinson (2007).
[5] See Viscusi (2018, ch. 2); Robinson (2007).
[6] See, e.g., Adler (2019); Adler and Posner (2006).
[7] See generally Viscusi (2018).
[8] On the U.S. government's use of population-average VSL, see Robinson (2007); U.S. Department of Health and Human Services (2016); U.S. Department of Transportation (2016); U.S. Environmental Protection Agency (2010). For the $10 million figure, see Viscusi (2018, p. 28).

Other governments also employ VSL in governmental cost-benefit analysis, including the U.K. The concept is the same, but the terminology may differ. In the U.K., the term "value of a prevented fatality" is used instead of "value of statistical life," and the corresponding abbreviation is "VPF" rather than "VSL."[9]

Although U.S. governmental monetization of risk policy via VSL was, initially, controversial, that controversy has largely faded, and VSL is quite firmly entrenched in U.S. regulatory practice. It is also very well established in the academic literature. Many articles in applied economics employ behavioral evidence, in particular wage differentials for riskier occupations, or stated-preference evidence, to estimate VSL.[10] There is also a large theory literature about VSL.[11]

The COVID-19 pandemic has brought VSL to the fore. Among the flood of academic working papers posted in immediate response to the pandemic, quite a number employ VSL to estimate the benefits of social-distancing requirements.[12] Many opinion pieces and news articles have also discussed VSL.[13] This is not unexpected. Since VSL is now a cornerstone of U.S. governmental risk analysis, and the focus of decades of academic work, it is hardly surprising that academic and public discussion about COVID-19 policy would resort to VSL.

But VSL—despite its familiarity—is flawed. VSL is flawed as a component of administrative routines for evaluating proposed regulations. It gives us a misleading picture, for example, of the benefits of EPA's anti-pollution regulations or the Department of Transportation's safety requirements. VSL is also a flawed basis for thinking about COVID-19 policy. Or so I shall argue here.

The deficits are the same in the two cases. VSL's shortcomings as applied to COVID-19 policy are no different from its shortcomings in evaluating anti-pollution rules, vehicle safety requirements, workplace safety regulations, and so forth. In the years before the pandemic, various scholars, including myself, published careful academic critiques of VSL.[14] These critiques didn't prompt intensive academic or public conversation, because fatality risk

---

[9] The U.K. also uses a considerably lower value for VSL than the U.S. See Viscusi (2018, p. 38) on international differences in VSL.

[10] For overviews of the empirical literature on VSL, see Aldy & Viscusi (2007); Cropper, Hammitt and Robinson (2011); Kniesner and Viscusi (2019); Krupnick (2007); OECD (2012); Viscusi and Aldy (2003); Viscusi (2018).

[11] See, e.g., Eeckhoudt and Hammitt (2001); Evans and Smith (2010); Hammitt (2000, 2007); Hammitt et al. (2020); Johansson (2002); Jones-Lee et al. (2015).

[12] See, e.g., Alvarez, Argente and Lippi (2020); Bairoliya and Imrohoroğlu (2020); Barnett-Howell and Mobarak (2020); Béland, Brodeur and Wright (2020); Bethune and Korinek (2020); Greenstone and Nigam (2020); Gros et al. (2020); Pretnar (2020); Robinson, Sullivan and Shogren (2020); Scherbina (2020); Thurnstrom et al. (2020); Ugarov (2020); Wilson (2020). But see Pindyck (2020), criticizing the use of VSL; Hall, Jones and Klenow (2020), using VSL to calibrate a utilitarian social welfare function.

[13] See, e.g., Corcoran (2020); Henderson and Lipow (2020); Hilsenrath and Armour (2020) Ingraham (2020); Jenkins (2020); Masur and Posner (2020); Sunstein (2020).

[14] See Adler (2016a); Adler (2017); Adler (2019, ch. 5); Adler, Hammitt and Treich (2014); Adler, Ferranna, Hammitt and Treich (2019); Broome (1978); Bronsteen, Buccafusco and Masur (2013); Broughel (2020); Dolan et al. (2008); Grüne-Yanoff (2009).

regulation—in the years before the pandemic—wasn't generally the topic of intensive academic and public conversation. Certainly it was not for the first two decades of the 21st century, up until the terrible coronavirus outbreak of 2020. The topics on the front burner were instead, e.g., terrorism, globalization, inequality, climate change, and democratic breakdown.

But the pandemic's death toll, and the massive economic losses of shutdowns and social distancing, force us to think about balancing lives against livelihood. Risk policy is now at the center of our conversations. And so, too, we should use this moment to think about how we think about risk policy. Is VSL really the best we can do? No, it is not—or so I'll claim.[15]

I'll consider three different versions of VSL: (1) textbook VSL, (2) population-average VSL, and (3) the "value of statistical life year" (VSLY). Textbook VSL is as presented in the academic literature on CBA and VSL. VSL, here, is an individual's marginal rate of substitution between survival probability and money, and varies among individuals. As mentioned above, U.S. governmental agencies depart from VSL and instead monetize VSL with a constant VSL, specifically a population average. VSLY is a construct proposed by some scholars as a way to circumvent the flaws of textbook and population-average VSL.[16]

Various kinds of policy choices have bedeviled governments during the COVID-19 pandemic, but perhaps the most important and salient such choice concerns social distancing policy. Measures to reduce close physical proximity between individuals in the population can be expected to lower the fatality (and also morbidity) impacts of the pandemic, but at the cost of lost income as well as losses with respect to other dimensions of well-being. This Article's strategy will be to illustrate the flaws of the three versions of VSL—textbook, population-average, and VSLY—with reference to social-distancing policy. (As already mentioned, the academic papers that have appeared in response to the pandemic and that employ VSL mainly do so with respect to social distancing policy.)[17]

My main benchmark for evaluating VSL will be *utilitarianism*. Utilitarianism has been a dominant school of ethical thought for hundreds of years, since the work of Jeremy Bentham. Utilitarianism has endured through generations of academic and public debates about ethics, and continues to play a vibrant role at both levels. To be sure, utilitarianism has hardly gone unchallenged. It has endured over centuries of ethical *debate*. In recent years, the chief opponents of utilitarianism within academic ethics have been non-consequentialists, in particular

---

[15] See also Hammitt (2020), critically assessing the use of VSL to evaluate social-distancing policy.

[16] On VSLY, see Aldy and Viscusi (2007); Hammitt (2007); Hammitt et al. (2020); Jones-Lee et al. (2015); Kniesner and Viscusi (2019); Viscusi (2018, ch. 5).

[17] A second type of policy choice that was much discussed early in the pandemic was triage—in particular, how to allocate lifesaving equipment to seriously ill COVID-19 patients (namely, ventilators) under scarcity. The academics engaged in that debate were principally health ethicists, who tend not to favor CBA; and thus CBA and VSL played little (if any) role in their analyses. Still, VSL-based CBA is a global methodology for assessing *any* kind of policy involving fatality risk and certainly could be applied to triage. Thus applied, it would be no less problematic than VSL-based CBA applied to social distancing. So as not to try the reader's patience in belaboring the difficulties of the three types of VSL with respect to both social distancing and triage, I focus here on the former.

"contractualists" and deontologists. However, I'm going to leave aside non-consequentialism and focus on a consequentialist critique of VSL—specifically, a critique from the vantage point of utilitarianism.[18]

Utilitarianism and CBA are *not* the same. There are key, structural, dissimilarities between the two that can yield divergent recommendations in many different policy domains, including fatality risk policy.

A secondary benchmark for my assessment of VSL will be *prioritarianism*.[19] Prioritarianism is a variation on utilitarianism that has emerged in ethics over the last several decades, and (as suggested by the name) gives extra weight to the well-being of the worse off.[20] It is appealing to those who are impressed by utilitarianism's consequentialist structure and attention to individual well-being, but believe that utilitarianism's exclusive focus on the sum total of well-being is too narrow. Instead, prioritarianism considers *both* the sum total *and* the distribution of well-being.

In developing my critique of VSL-based CBA from the standpoint of utilitarianism and prioritarianism, I will rely upon a simulation model that is built upon the U.S. population survival curve and income distribution, and that will be used as a testbed for balancing the benefits and costs of social distancing policy.

The standard defense of CBA appeals to Kaldor-Hicks efficiency (potential Pareto superiority).[21] Assume that CBA picks policy option P^*, as opposed to a second option P— which could be the status quo, or alternatively the option chosen by a competing policy framework such as utilitarianism or prioritarianism. Then there is in principle a change ΔT to the tax-and-transfer system such that P^* together with ΔT is Pareto superior to P. ΔT would take the form of increased taxes upon (or lowered transfer payments to) those better off with P^* than P, and increased transfer payments to (or lowered taxes upon) those better off with P than P^*. In recent years, work by Louis Kaplow has extended the Kaldor-Hicks defense of CBA, by showing that ΔT exists not merely in a lump-sum tax system but even in a system (such as ours) where taxes are levied on income and thus increased taxes have a distortionary effect.[22]

---

[18] Non-consequentialists have struggled to formulate a systematic account of the ethics of risk. For a recent attempt, see Frick (2015). Utilitarianism experiences no such difficulty. Further, CBA itself is consequentialist. It will be especially illuminating, I hope, to set forth a critique of VSL that is not merely a rehash of the criticisms that non-consequentialists level against all forms of consequentialism—but instead is a critique from within the consequentialist camp, namely from the standpoint of utilitarianism and, secondarily, prioritarianism.

[19] Utilitarianism and prioritarianism are operationalized for policy assessment as social welfare functions (SWFs). On utilitarian and prioritarian SWFs, see generally Adler (2019). On these SWFs applied specifically to risk regulation, see Adler (2016a); Adler (2017); Adler (2019, ch. 5); Adler, Hammitt and Treich (2014); Adler, Ferranna, Hammitt and Treich (2019).

[20] Parfit (2000) is the seminal philosophical text on prioritarianism. For defenses, and overviews of the subsequent philosophical literature, see Adler (2012, ch. 5); Holtug (2010, 2017).

[21] See Adler and Posner (2006), discussing and criticizing this defense.

[22] See Kaplow (1996; 2004; 2008).

The Kaldor-Hicks defense of CBA has been vigorously challenged. These challenges are set forth in a substantial academic literature debating the Kaldor-Hicks criterion.[23] Perhaps the strongest challenge is this: the change ΔT to the tax system is a purely *potential* change. Either ΔT is *actually* implemented together with P^*, in which case *any* Pareto-respecting assessment methodology (e.g., utilitarianism) will prefer P^* plus ΔT to P; or ΔT is *not* implemented together with P^*, in which case some are worse off with the policy and the Pareto criterion is not applicable. Indeed, in actual U.S. practice, administrative rules and new statutory provisions are rarely coupled with tax-and-transfer changes designed to compensate those who are made worse off; and regulatory agencies are instructed to engage in CBA regardless of whether they anticipate such a compensatory change.

For purposes of this Article, however, I will place to one side these well-known critiques of Kaldor-Hicks efficiency. Assume that the reader *accepts* the Kaldor-Hicks criterion. Should she, therefore, accept VSL-based CBA as a tool for evaluating fatality risk regulation and, more specifically, COVID-19 policy?

No. VSL-based CBA can *deviate* from the Kaldor-Hicks criterion. An individual's VSL, multiplied by the change in fatality risk that would result from a given policy, is only an *approximation* to her willingness to pay for that change. This approximation becomes poorer as the change becomes larger. Thus CBA with textbook VSL, used to assess policies that involve significant changes in individuals' risks, can readily favor policies that are *not* Kaldor-Hicks efficient relative to alternatives that CBA disprefers. And CBA with population-average VSL or VSLY can deviate from Kaldor-Hicks efficiency even for small changes in individuals' risks. We'll see, specifically, that all three approaches assign a positive score to a range of social distancing policies, preferring these policies to the status quo—even though the policies are not Kaldor-Hicks efficient relative to the status quo.

The structure of the Article is as follows. Part II provides the conceptual framework. Part III uses the simulation model based upon the U.S. survival curve and income distribution to illustrate the implications of textbook VSL (as defined in Part II) with respect to social distancing policy, as compared to utilitarianism and prioritarianism, and to assess whether VSL conforms to the Kaldor-Hicks criterion. Part IV does the same for CBA using population average VSL, and Part V for VSLY.

The aim of the Article, it should be stressed, is *not* to provide guidance with respect to social distancing policy. Rather, its aim is *methodological*: to argue that VSL-based CBA provides flawed guidance. The simulation model is used as a diagnostic tool, to bring to light the deficiencies of VSL-based CBA—to demonstrate how the guidance it provides is different from that of utilitarianism, prioritarianism, and Kaldor-Hicks efficiency.

---

[23] See Adler (2012, pp. 98-104), citing sources. For a response to Kaplow, see Adler (2019, pp. 225-33); Adler (2017).

II.   The Concept of VSL

Roughly speaking, VSL is an individual's marginal rate of substitution between money (income, consumption, wealth) and survival probability. It captures how money and survival trade off in terms of the individual's utility.

Providing a more precise definition of VSL depends on how utility is conceptualized. The conceptual framework I'll be using in this paper, a discrete-time and multi-period model, is as follows. An individual's life is divided into periods numbered 1, 2, … T, with T, the maximum number of periods that any individual lives. An individual either dies immediately after the beginning of the period or does not die and survives, at least, until the end of the period.

The number of the current period, for individual i, is A_i. (For example, if periods are years, and individual i is currently age 20, then $A_i = 21$.)

In the status quo, individual i has a lifetime profile of survival probabilities $(p_i^{A_i}, ..., p_i^{T})$. These are conditional probabilities: namely, p_i^{t} is the probability that individual i survives to the end of period t, given that she is alive at the beginning.

Individual i also has a status quo lifetime profile of income and non-income attributes. $((y_i^{1}, b_i^{1}), ..., (y_i^{A_i-1}, b_i^{A_i-1}), (y_i^{A_i}, b_i^{A_i}), ..., (y_i^{T}, b_i^{T}))$. If t is a past period ($t < A_i$), y_i^{t} is the income that individual i earned in period t and b_i^{t} denotes the non-income attributes that she had in that period. If t is the current period or a future period ($t \geq A_i$), then y_i^{t} is the income which i will earn in period t if she survives to its end, and b_i^{t} is the bundle of non-income attributes which she'll have in that period if she survives to its end.

For short, let $(\mathbf{p}_i, \mathbf{y}_i, \mathbf{b}_i)$ denote individual's status quo lifetime profile of survival probabilities, incomes, and non-income attributes.

Individual i has a lifetime utility function. Let $U_i(\cdot)$ denote expected utility. Individual i's status quo expected utility, at present, is
$U_i((p_i^{A_i}, p_i^{A_i+1}, ..., p_i^{T}), ((y_i^{1}, b_i^{1}), ..., (y_i^{A_i}, b_i^{A_i}), ..., (y_i^{T}, b_i^{T})))$ or, for short, $U_i(\mathbf{p}_i, \mathbf{y}_i, \mathbf{b}_i)$.

Note that the framework is quite general. Periods can be arbitrarily long or short. Non-income attributes of any sort (health, leisure, happiness, public goods, etc.) can be included in the period bundles. An individual's consumption in each period might be set equal to her income; or, instead, the individual might be supposed to have access to intertemporal financial markets, allowing her to save and to borrow against anticipated income, and to engage in saving and borrowing in either a perfectly or imperfectly rational manner. Finally, the lifetime utility function might be temporally additive or non-additive.

Assume that individual i's survival probability in period s is increased, relative to the status quo, by Δp_i^{s}. Period s is either the current period or a future period ($s \geq A_i$). Let $m(\Delta p_i^{s})$

denote the reduction to individual i's current status quo income that just suffices to make her indifferent to this improvement in survival probability.[24] Then VSL_i^s , individual's i's VSL with respect to changes in survival probability in period s, is the limit of the ratio of $m(\Delta p_i^s)$ to Δp_i^s: $VSL_i^s \equiv \lim\limits_{\Delta p_i^s \to 0} \dfrac{m(\Delta p_i^s)}{\Delta p_i^s}$. Equivalently, VSL_i^s is the marginal rate of substitution between survival probability in period s and current income.

$$(1) \qquad\qquad\qquad VSL_i^s = \dfrac{\partial U_i / \partial p_i^s}{\partial U_i / \partial y_i^{A_i}}\Big|_{(\mathbf{p}_i, \mathbf{y}_i, \mathbf{b}_i)}$$

The numerator in this fraction is the status quo marginal utility of period s survival probability ; the denominator is the status quo marginal utility of current income. VSL_i^s is the ratio of these marginal utilities.[25]

CBA is standardly defined as the sum of individual "compensating variations" or "equivalent variations."[26] A given governmental policy is some departure from the status quo. An individual's "compensating variation" for the policy is the change to her current income with the policy—the current income that she would have, were the policy to be implemented—that would just suffice to make her indifferent between the policy and the status quo. An individual's "equivalent variation" for the policy is the change to her current status quo income that would just suffice to make her indifferent between the policy and the status quo.[27]

VSL is a useful construct, for purposes of CBA, because it can be used to define an individual monetary valuation for a policy that is a good approximation to her compensating variation or equivalent variation—*if* the policy is a small change from the status quo. In general, I'll use the symbol "MV" to denote this VSL-based monetary valuation of a policy: one that well

---

[24] In other words, if s is the current period, $m(\Delta p_i^s)$ is such that

$U_i((p_i^{A_i} + \Delta p_i^s, p_i^{A_i+1}, ..., p_i^T), ((y_i^1, b_i^1), ..., (y_i^{A_i} - m(\Delta p_i^s), b_i^{A_i}), ..., (y_i^T, b_i^T))) = U_i((p_i^{A_i}, ..., p_i^T), ((y_i^1, b_i^1), ..., (y_i^{A_i}, b_i^{A_i}), ..., (y_i^T, b_i^T)))$
. If s is a future period, $m(\Delta p_i^s)$ is such that:

$U_i((p_i^{A_i}, ..., p_i^s + \Delta p_i^s, ..., p_i^T), ((y_i^1, b_i^1), ..., (y_i^{A_i} - m(\Delta p_i^s), b_i^{A_i}), ..., (y_i^T, b_i^T))) = U_i((p_i^{A_i}, ..., p_i^T), ((y_i^1, b_i^1), ..., (y_i^{A_i}, b_i^{A_i}), ..., (y_i^T, b_i^T)))$

[25] In one-period models, VSL is often defined as the marginal rate of substitution between survival probability and *wealth* rather than income. See, e.g., Eeckhoudt and Hammitt (2001). However, in a multiperiod model in which individuals have exogenous incomes in each period, it is more straightforward to define VSL in terms of this exogenous attribute rather than to construct an endogenous wealth attribute and then define VSL in terms of that. See Adler et al. (2019). If individuals are modelled as behaving myopically (consuming income and other resources available in each period, rather than saving and borrowing), then income equals wealth except perhaps in the first period (an inheritance). If individuals save and borrow, then VSL defined in terms of income will reflect opportunities to save and borrow.

Further, the sum of compensating variations whether defined in terms of income or wealth signals a potential Pareto improvement (Kaldor-Hicks efficiency), and thus VSL whether defined in terms of income or wealth coheres with the traditional justification for CBA (namely that VSL for small changes is a good approximation to the compensating variation).

[26] See, e.g., Adler and Posner (2006); Boadway (2016); Freeman (2003, ch. 3).

[27] See Appendix D for formulas for the equivalent and compensating variations.

approximates the individual's compensating and equivalent variations if the policy is a small change from the status quo. And I'll use the term "VSL-CBA" to mean the version of CBA that assesses policies by summing these MV values.

In particular, imagine that a policy increases individual i's period s survival probability by Δp_i^s. Then $MV_i = (\Delta p_i^s)VSL_i^s$. Alternatively, imagine that a policy increases individual i's period s survival probability by Δp_i^s and increases her current income by $\Delta y_i^{A_i}$. Then $MV_i = (\Delta p_i^s)VSL_i^s + \Delta y_i^{A_i}$.[28]

These formulas show why VSL can be thought of as a conversion factor that translates risk changes into monetary equivalents. If a policy changes i's period s survival probability by Δp_i^s, the monetary equivalent for that risk change is $(\Delta p_i^s)VSL_i^s$. This monetary equivalent is added to the monetary equivalents for the other policy impacts on individual i to determine her overall monetary valuation (MV_i) of the policy. In the simple case where the policy only changes i's period s survival probability, MV_i is equal to the monetary equivalent of the change: $(\Delta p_i^s)VSL_i^s$. Consider next a policy that changes i's current income by $\Delta y_i^{A_i}$. MV_i in this case is nothing other than $\Delta y_i^{A_i}$.[29] Finally, in the case where a policy changes both an individual's survival probability and her current income, we have the formula $MV_i = (\Delta p_i^s)VSL_i^s + \Delta y_i^{A_i}$. The first term is the monetary equivalent for the risk change Δp_i^s, i.e., the risk change multiplied by VSL_i^s; the second term is the monetary equivalent for the income change $\Delta y_i^{A_i}$, which is, trivially, just $\Delta y_i^{A_i}$. MV_i is the sum of these two.

Several other points about VSL are worth noting at this juncture. First, VSL is heterogeneous. An individual's VSL with respect to a current or future period depends upon her age, current and future (and perhaps past) income, and current and future (and perhaps past) non-income attributes. These can all, of course, vary among individuals, and so VSL can vary as well.

Second, an individual's MV is only an approximation to her compensating variation or equivalent variation, and this approximation may well be quite poor for policies that produce a large change in the individual's survival probabilities or non-income attributes. In particular, the component of MV for valuing a change to an individual's period s survival probability—$(\Delta p_i^s)VSL_i^s$—reflects the marginal impact of income on the individual's expected utility. The individual's willingness to pay for a non-marginal reduction in risk, or her willingness to accept in return for a non-marginal increase in risk, will not generally be equal to or even well

---

[28] Similar formulas can be defined for policies that change non-income attributes, but because the model of social distancing below will focus on changes to risks and incomes, I won't spell out such formulas here.
[29] Note that MV_i in this simple case is not merely an approximation for the compensating and equivalent variation, but indeed exactly equal to both.

approximated by $(\Delta p_i^s)VSL_i^s$ if her marginal utility of income is not constant and Δp_i^s is sufficiently large.

For an illustration of this point, consider a simple version of the conceptual framework here in which each individual is endowed with a lifetime vector of survival probabilities and incomes. The period length is one year; non-income attributes are ignored; and the individual consumes her income each period. Using this model, and the most recent available U.S. survival curve, I calculated equivalent variations, compensating variations, and MV values for a 50 year old individual earning \$60,000/year with an additive lifetime utility function equaling the sum of log income.[30] I assumed that the individual's current-year survival probability was as per the survival curve (.99597), minus 0.01. In short, in the status quo the individual's ordinary current-year survival probability has been lowered by 1% by virtue of some significant, short-term fatality risk (such as a pandemic). Under all the assumptions just stated, the individual's VSL for a current-year change in survival probability is \$9,081,485.

Table 1 displays the individual's equivalent variations, compensating variations, and MVs for a policy that improves her survival probability by amounts ranging from 1 in 1 million to 1%. (MV is simply \$9,081,485 times the risk change.) As Table 1 shows, MV becomes an increasingly poor approximation to the equivalent or compensating variation as the risk change from the policy increases. At the 1% level (a policy that wholly eliminates the acute current-year risk), the MV is \$90,814.85, while her equivalent variation is more than twice that amount, and her compensating variation roughly half.

Table 1: Equivalent Variations, Compensating Variations, and
VSL-based Monetary Valuations (MVs)

Probability Change	Equivalent Variation	Compensating Variation	MV	Diff. between Equiv. Variation and MV as % of MV	Diff. between Comp. Variation and MV as % of MV
1 in 1 million	\$9.08	\$9.08	\$9.08	.01%	.01%
1 in 100,000	\$90.88	\$90.75	\$90.81	.08%	.08%
1 in 10,000	\$915.06	\$901.22	\$908.15	.76%	.76%
1 in 1000	\$9804.79	\$8419.70	\$9081.48	7.96%	7.29%
5 in 1000	\$67,885.45	\$31,742.11	\$45,407.42	49.50%	30.09%
1 in 100	\$212,578.16	\$46,590.54	\$90,814.85	134.08%	48.70%

The third and final observation about VSL worth noting at this juncture is closely related to the second. Because an individual's MV is only an *approximation* to her equivalent and compensating variations, VSL-CBA—the sum of MVs—need not conform to Kaldor-Hicks

---

[30] Temporal additivity and having utility be logarithmic in income are both quite standard assumptions in economics. Appendix D gives explicit formulas for equivalent and compensating variations applicable to the example here.

efficiency. A policy may be ranked above the status quo by VSL-CBA and yet not be Kaldor-Hicks efficient relative to the status quo.

To continue with the above example: Imagine that there are a certain number of at-risk individuals in the population similarly situated to the individual above; and that, for each such individual, there are 1000 potential cost-bearers who aren't at risk but would bear the cost of any policy to reduce or mitigate the risk. Each at-risk individual has an MV of $90,814.85 for a policy that eliminates the 1% risk, but a compensating variation of only $46,590.54. She is willing to pay no more than $46,590.54 to eliminate the risk. Imagine, now, that the per-capita cost of a policy to eliminate the risk is between $46.59 and $90.81. Such a policy is approved by VSL-CBA (it has a positive sum of MVs), but is *not* Kaldor-Hicks efficient. There is no scheme of compensating payments from the at-risk individuals to the cost-bearers which, if adopted together with the policy, makes everyone better off than the status quo.

The conceptual framework set out here can also be used to highlight the difference between VSL-CBA and utilitarianism. While VSL-CBA ranks policies according to the sum of individuals' MV_i values, utilitarianism assigns each policy a score equaling the sum of individuals' expected utilities and ranks policies in the order of these scores.[31] Let ΔU_i denote the difference in i's expected utility between a given policy and the status quo. Utilitarianism says that the policy is better than the status quo if the sum of these ΔU_i values is positive; by contrast, VSL-CBA says that the policy is better than the status quo if the sum of the MV_i values is positive.

To illustrate the difference, consider a policy that changes i's period s survival probability by small amount Δp_i^s and her current income by a small amount $\Delta y_i^{A_i}$. In this case:

$$(2) \qquad \Delta U_i \approx (\Delta p_i^s) \frac{\partial U_i}{\partial p_i^s}\Big|_{(\mathbf{p}_i, \mathbf{y}_i, \mathbf{b}_i)} + (\Delta y_i^{A_i}) \frac{\partial U_i}{\partial y_i^{A_i}}\Big|_{(\mathbf{p}_i, \mathbf{y}_i, \mathbf{b}_i)}$$

The first term in this equation is the approximate change in expected utility that results from the risk change; it is Δp_i^s multiplied by individual i's status quo marginal utility of period s survival probability. The second term is the approximate change in expected utility that results from the income change; it is $\Delta y_i^{A_i}$ multiplied by the status quo marginal utility of current income. Consider now the formula for MV_i. Recalling that VSL is the ratio of the marginal utilities of survival probability and income, we have:

---

[31] I focus in this Article on utilitarianism understood as the sum of individuals' von-Neumann/Morgenstern (vNM) utilities. $U_i(\cdot)$ is a vNM utility function. If individuals have homogeneous preferences with respect to lifetime bundles of income and non-income attributes, and lotteries over these bundles, then utilitarianism is implemented by choosing any one of the vNM utility functions representing the common preferences and setting $U_i(\cdot)$ to be this function. If individuals have heterogeneous preferences, then we take a vNM utility function representing each preference in the population and rescale it using scaling factors. In this case, $U_i(\cdot)$ is the rescaled vNM utility function representing the preferences of individual i. See Adler (2019, ch. 2 and Appendix D); Adler (2016b).

$$ (3) \qquad MV_i = \Delta p_i^s \frac{\partial U_i / \partial p_i^s}{\partial U_i / \partial y_i^{A_i}} \big|_{(\mathbf{p}_i, \mathbf{y}_i, \mathbf{b}_i)} + \Delta y_i^{A_i} $$

Comparing the formulas for ΔU_i and MV$_i$, two differences emerge. (1) While the ΔU_i formula multiplies Δp_i^s by the marginal utility of period s survival probability, the MV$_i$ formula does the same but then divides by the marginal utility of current income. Consider, for example, two individuals facing the very same small change in current survival probability. The relative valuation of these two changes, according to utilitarianism, depends upon the individuals' relative marginal utilities of current survival probability. By contrast, according to VSL-CBA, the relative valuation of the two changes depends upon these marginal utilities of survival probability *and* the two individuals' relative marginal utilities of current income. (2) While the ΔU_i formula multiples $\Delta y_i^{A_i}$ by marginal utility of current income, the MV$_i$ formula simply adds $\Delta y_i^{A_i}$ without adjustment. This difference illustrate why utilitarianism is sensitive to income distribution,[32] while VSL-CBA is not. Imagine that a policy redistributes a small Δy in current income from individual i to individual j. VSL-CBA sees this policy as a wash: MV$_i$ decreases by Δy, while MV$_j$ increases by Δy, and thus the sum of these monetary valuations does not change. By contrast, the relative magnitudes of ΔU_i and ΔU_j depend upon the individuals' relative marginal utilities of current income. If individual i has lower marginal utility of income than j, ΔU_i will be smaller in magnitude than ΔU_j, and so utilitarianism will see the income redistribution as an improvement.

Finally, the difference between *prioritarianism* and VSL-CBA can also be illustrated in the conceptual model being deployed here. Prioritarianism sums expected *transformed* utilities, using a *concave* transformation function (so as to give greater weight to the worse off).[33] Let ΔG_i denote the difference in i's expected transformed utility between a given policy and the status quo; the prioritarian score for the policy is the sum of these ΔG_i values. Consider once more a policy that changes i's period s survival probability by small amount Δp_i^s and her current income by a small amount $\Delta y_i^{A_i}$. In this case:

$$ (4) \qquad \Delta G_i \approx (\Delta p_i^s) \frac{\partial G_i}{\partial p_i^s} \big|_{(\mathbf{p}_i, \mathbf{y}_i, \mathbf{b}_i)} + (\Delta y_i^{A_i}) \frac{\partial G_i}{\partial y_i^{A_i}} \big|_{(\mathbf{p}_i, \mathbf{y}_i, \mathbf{b}_i)} $$

---

[32] See Robinson, Hammitt and Zeckhauser (2016), documenting inattention to distribution in agency cost-benefit analyses of environmental, health and safety regulations.

[33] This is so-called "ex post" prioritarianism, as opposed to "ex ante" prioritarianism, a different version of prioritarianism under uncertainty. This Article focuses on the "ex post" approach. Ex ante prioritarianism violates a stochastic dominance axiom, and for that reason is quite problematic. See generally Adler (2019, chs. 3-4).

While the formula above for MV$_i$ values a risk change (Δp_i^s) by multiplying by the term $\dfrac{\partial U_i}{\partial p_i^s}$ and then dividing by a marginal income utility term $\dfrac{\partial U_i}{\partial y_i^{A_i}}$, the prioritarian formula instead multiplies by the term $\dfrac{\partial G_i}{\partial p_i^s}$ —this captures the effect of the risk change on expected transformed utility – and does not divide by a marginal income utility term. Further, like utilitarianism, and unlike VSL-CBA, prioritarianism is sensitive to the distribution of income. The income change for an individual, $\Delta y_i^{A_i}$, is multiplied by a term $\dfrac{\partial G_i}{\partial y_i^{A_i}}$, capturing the effect of income on expected transformed utility.

III.   **VSL and Pandemic Policy: A Simulation Model**

A.   *A Simulation Model*

In order to compare VSL with utilitarianism and prioritarianism in assessing pandemic policy, I construct a simulation model based upon the U.S. population survival curve and income distribution. The population is divided into five income quintiles (denoted as "low," "moderate," "middle," "high," and "top"). I assume that an individual remains in the same quintile her entire life. An individual, at birth, is endowed with a survival curve and lifetime income profile. The survival curve for the middle quintile is the U.S. population survival curve, while the curves for the other quintiles are adjusted up or down to match observed differences in life expectancy by income.

The lifetime income profile, for a given quintile, is constructed as follows. Based upon U.S. governmental data, I estimate average after-tax-and-transfer individual income, by quintile, to be as follows: \$21,961; \$30,118; \$41,349; \$57,538; and \$134,840. An individual's income at a given age is then scaled up or down from this baseline amount, in accordance with data concerning the time path of earnings.

I assume that the current population consists of seven age groups: individuals aged 20, 30, 40, 50, 60, 70, and 80. Each age group, in turn, has five income quintiles. Thus the population consists of 35 age-income cohorts. The proportions of the population in the seven age bands are chosen to match the actual age distribution of the current U.S. population.

Although in reality the COVID-19 pandemic increases individuals' fatality risks over multiple years, I simplify the analysis by assuming that this risk manifests itself as an increased fatality risk only in the current year. The baseline current-year survival probability for each age-income cohort is taken from the survival curve for her income quintile, *except* that her survival chances have been made worse by the COVID-19 pandemic. The baseline (status quo) captures the case in which the population is faced with the pandemic and government does nothing to

address it; social distancing policies will, then, reduce the COVID-19 risk, at some cost to individual incomes.

The COVID-19 risk for each age group is taken from the now-famous Imperial College report of March 2020, whose alarming estimates of the potential deaths from the virus triggered the adoption of shutdown policies in the U.K. and U.S. This report estimated infection fatality ratios (IFRs) by age group and also estimated that 81% of the population in each country would be infected absent governmental intervention.[34] Using the Imperial College IFRs by age and the 81% infected estimate, I assigned the current year COVID-19 risks given in Table 2 to the seven age groups.

Table 2: Risk of Dying from COVID-19 Absent Policy to Suppress or Mitigate Pandemic

Age group	Baseline (No Intervention) Risk of Dying from COVID-19 in Current Year
20	.024%
30	.065%
40	.122%
50	.486%
60	1.782%
70	4.131%
80	7.533%

I ignore non-income attributes and assume that an individual's consumption in each year equals her income. I assume a common, temporally additive, logarithmic, lifetime utility function: utility is the sum of the logarithm of consumption each year, minus a term to reflect the subsistence level of consumption.

In this setup, the utilitarian value of a policy is the sum of individuals' expected utilities (according to the common utility function just described). For prioritarianism, I use an Atkinson social welfare function with a moderate level of priority to the worse off, γ (the priority parameter) = 1.5.

More details about the simulation model, the lifetime utility function, and the utilitarian and prioritarian social welfare functions are provided in the Appendix.

VSLs for the 35 age-income cohorts, for a current-year change in survival probability, were calculated using the above assumptions about survival curves and lifetime income profiles for each of the cohorts and the common lifetime utility function, and are displayed in Table 3:

---

[34] Ferguson et al. (2020).

	Income: Low	Moderate	Middle	High	Top
Age: 20	$1,379,256	$2,181,761	$3,391,051	$5,240,842	$15,369,839
30	$3,163,107	$5,005,939	$7,784,849	$12,026,744	$35,235,573
40	$3,254,671	$5,191,073	$8,122,805	$12,593,803	$37,122,979
50	$2,454,637	$3,967,977	$6,274,797	$9,791,316	$29,183,439
60	$1,494,663	$2,458,852	$3,943,065	$6,207,337	$18,786,610
70	$755,476	$1,272,810	$2,080,198	$3,313,544	$10,231,995
80	$330,329	$570,613	$951,330	$1,533,270	$4,826,901

The average VSL across the 35 cohorts (weighted for the different proportions of the age groups) is $8,635,355. This is reasonably consistent with estimates of the U.S. population average VSL, typically in the range of $10 million. Observe also that the pattern of VSL in Table 3 has two features observed in much empirical work: VSL increases with income;[35] and its time profile has a "hump" shape—first increasing and then decreasing with age.[36]

The VSL information in Table 3 can be expressed in a different way. Consider an increment Δp in a given cohort's current-year fatality risk (increment meaning either a reduction or increase). The social value of that increment, according to VSL-CBA, is $MV_c = (\Delta p)VSL_c$. Let c^+ be some reference cohort. Then the relative value of an increment in fatality risk for cohort c, as compared to the value for the reference cohort, is just VSL_c/VSL_{c+}. Table 4a expresses the VSL information in this fashion, with an 80-year old, low income group as the reference cohort.

Table 4a: Social value of risk increment according to VSL-CBA
(relative to social value of risk increment for 80-year-old, low-income cohort)

	Income: Low	Moderate	Middle	High	Top
Age: 20	4.2	6.6	10.3	15.9	46.5
30	9.6	15.2	23.6	36.4	106.7
40	9.9	15.7	24.6	38.1	112.4
50	7.4	12.0	19.0	29.6	88.3
60	4.5	7.4	11.9	18.8	56.9
70	2.3	3.9	6.3	10.0	31.0
80	1.0	1.7	2.9	4.6	14.6

For example, Table 4a has the number 15.7 in the cell for the 40-year-old, moderate income cohort. This indicates that the MV for a given risk increment Δp accruing to a member of the 40-

---

[35] See, e.g., Kniesner and Viscusi (2019); Viscusi (2018, ch. 6).
[36] See, e.g., Aldy and Viscusi (2007); Viscusi (2018, ch. 5).

year-old, moderate-income cohort is 15.7 times the MV for the very same risk increment Δp accruing to a member of the 80-year-old, low-income cohort. The entry in the cell for the 20-year-old, top-income cohort is 46.5. So the MV for a risk increment of Δp to this cohort is 46.5 times that for the same increment to an 80-year-old, low-income individual; and 46.5/15.7 = 2.96 times that for the same increment to a member of the 40-year-old, moderate income cohort.

Tables 4b and 4c display the analogous information for utilitarianism and prioritarianism.[37] These show the relative social value of an increment Δp in each cohort's current-year fatality risk, again with 1 indicating its value to an 80-year-old low-income individual, according to utilitarianism (Table 4b) and prioritarianism (Table 4c). The social value of a risk increment to a cohort member as per utilitarianism is not the individual's MV, but rather the change in the individual's expected utility; the social value as per prioritarianism is the change in the individual's expected transformed utility.

Table 4b: Social value of risk increment according to utilitarianism
(relative to social value of risk increment for 80-year-old, low-income cohort)

	Income: Low	Moderate	Middle	High	Top
Age: 20	8.6	9.9	11.2	12.5	15.6
30	7.4	8.6	9.7	10.8	13.5
40	5.9	6.9	7.9	8.8	11.0
50	4.4	5.2	6.0	6.7	8.6
60	3.0	3.6	4.3	4.8	6.2
70	1.9	2.3	2.7	3.1	4.2
80	1.0	1.3	1.6	1.8	2.5

Table 4c: Social value of risk increment according to prioritarianism
(relative to social value of risk increment for 80-year-old, low-income cohort)

	Income: Low	Moderate	Middle	High	Top
Age: 20	31.5	29.6	27.9	26.3	23.2
30	19.7	18.6	17.7	16.7	14.8
40	11.9	11.4	11.0	10.5	9.5
50	7.0	6.9	6.8	6.5	6.0
60	4.0	4.0	4.0	4.0	3.7
70	2.1	2.2	2.3	2.2	2.2
80	1.0	1.1	1.1	1.1	1.1

---

[37] The numbers in tables 4a, 4b and 4c as well as 5a, 5b, and 5c are rounded to 1 decimal place (except for prioritarian income numbers, in Table 5c, for top-income cohorts)

As emerges from these tables, the relative values of risk increment assigned by VSL are quite different from those assigned by either utilitarianism and prioritarianism.[38] Within each income quintile, the utilitarian value of a risk increment decreases monotonically as individuals get older; saving a younger individual produces a larger gain in life expectancy and thus (within each income quintile) a larger gain in expected utility. Prioritarian values also decrease monotonically—yet more sharply than utilitarian values, reflecting priority to younger cohorts (who are worse off with respect to expected lifetime utility.) By contrast, the VSL-based values do not decrease monotonically with age. Rather, they have the classic "hump" shape. VSL-based values reflect not merely the gain to expected utility, but also the expected marginal utility of income—which changes over time because income changes over time within each quintile, first increasing with age and then decreasing. (See Table 6.)

Within each age band, the utilitarian value of a risk increment increases moderately with income (social values for top-income individuals are roughly twice those of low-income individuals of the same age). Because the utility of each year alive increases with income, the utilitarian value of extending a higher-income individual's life by a year is greater than that of extending a lower-income individual's life by a year. VSL-based values also increase with income in each age band, but much more dramatically than utilitarian values. VSL-based values are more highly skewed to the rich than utilitarian values, because the denominator of VSL is expected marginal utility of income—which decreases within each age band as income increases. Prioritarian values generally decrease or stay constant with income, reflecting priority for the worse off: lower-income individuals within an age band are worse off with respect to lifetime utility.

Tables 5a, 5b, and 5c display the same information on the income side. These show the relative social value of a $1 increment in a cohort's current income, again with an 80-year-old, low-income cohort as the reference cohort. The number in each cell is the social value of a $1 income increment, as a multiple of the value of a $1 income increment for an 80-year-old, low-income individual.[39]

---

[38] See Adler (2019, ch.5), for similar tables showing the relative risk valuations of utilitarianism, prioritarianism, and CBA, and for a longer discussion of why the three methodologies have different such valuations.

[39] Because the effect of an income increment on an individual's expected utility and expected transformed utility is non-linear, the relative value of income increments for the cohorts displayed in Tables 5b and 5c depends upon the magnitude of the increment (here, $1). MV values are linear in income increments, and so the constant pattern in Table 5a will hold true for any increment, not merely $1.

By contrast, it can be shown that the effect of a current-year change in survival probability on individuals' expected utility and expected transformed utility *is* linear. Thus the relative values stated in Tables 4b and 4c are independent of the size of the increment. MV too, by construction, is linear in risk reduction, and thus the relative values in Table 4a (cohort VSLs, normalized by dividing by the VSL of the 80-year-old, low-income cohort) are also independent of the size of the increment.

Table 5a: Social value of $1 income increment according to VSL-CBA
(relative to social value of $1 increment for 80-year-old, low-income cohort)

	Income: Low	Moderate	Middle	High	Top
Age: 20	1	1	1	1	1
30	1	1	1	1	1
40	1	1	1	1	1
50	1	1	1	1	1
60	1	1	1	1	1
70	1	1	1	1	1
80	1	1	1	1	1

Table 5b: Social value of $1 income increment according to utilitarianism
(relative to social value of $1 increment for 80-year-old, low-income cohort)

	Income: Low	Moderate	Middle	High	Top
Age: 20	2.1	1.5	1.1	0.8	0.3
30	0.8	0.6	0.4	0.3	0.1
40	0.6	0.4	0.3	0.2	0.1
50	0.6	0.4	0.3	0.2	0.1
60	0.7	0.5	0.4	0.3	0.1
70	0.8	0.6	0.4	0.3	0.1
80	1.0	0.7	0.5	0.4	0.2

Table 5c: Social value of $1 income increment according to prioritarianism
(relative to social value of $1 increment for 80-year-old, low-income cohort)

	Income: Low	Moderate	Middle	High	Top
Age: 20	3.2	1.8	1.1	0.7	0.20
30	1.1	0.6	0.4	0.2	0.07
40	0.8	0.5	0.3	0.2	0.05
50	0.7	0.4	0.3	0.2	0.05
60	0.8	0.5	0.3	0.2	0.05
70	0.9	0.5	0.3	0.2	0.06
80	1.0	0.6	0.4	0.2	0.08

The VSL-based social value of an income increment is constant (the number 1 is in every cell.) Recall Equation 3 above. An increment Δy to an individual's income produces the very same change to her MV (namely, by Δy), regardless of who receives it. By contrast, utilitarian values decrease with income within each age band; this reflects the diminishing marginal utility

of income. Prioritarian values decreases yet more dramatically, reflecting both the diminishing marginal utility of income and priority for the worse off.

Given the fairly stark differences between the relative valuations of risk increment and income increment, as per VSL-CBA, and those relative valuations as per utilitarianism and prioritarianism, it is not surprising that the recommendations of VSL-CBA with respect to social distancing policy are different from those of the latter two methodologies—as we'll now see.

B.   *Social Distancing Policy*

The current-year incomes of the various cohorts are displayed in Table 6. To model social distancing policy, I assume that an 80% reduction in "GDP" (the sum of total current income across the cohorts) will completely eliminate the current-year COVID-19 risk (as given in Table 2).[40]

Table 6: Cohort Incomes

	Income: Low	*Moderate*	*Middle*	*High*	*Top*
Age: 20	$8,331	$11,425	$15,686	$21,827	$51,152
30	$22,098	$30,306	$41,607	$57,896	$135,680
40	$28,426	$38,984	$53,522	$74,476	$174,536
50	$28,681	$39,334	$54,003	$75,145	$176,103
60	$24,930	$34,189	$46,939	$65,316	$153,069
70	$19,719	$27,043	$37,128	$51,664	$121,075
80	$14,757	$20,238	$27,784	$38,662	$90,605

I posit a concave function from GDP reduction to the reduction in COVID-19 risk.[41] Specifically, I set the COVID-19 risk reduction to be the square root of the GDP reduction, scaled so that a 80% reduction wholly eliminates the risk. (See Appendix.). A concave rather than linear or convex function is employed on the assumption that GDP reductions have a diminishing marginal impact on COVID-19 risk. The first 1% GDP loss produces a larger drop in COVID-19 risk than the second 1%, and so on.

The assumptions just described are the base case for analysis. To allow for the possibility that the base case may be too pessimistic or optimistic about the costs of eliminating COVID-19 risk, I consider two alternative cases: a 40% reduction in GDP wholly eliminates the COVID-19

---

[40] This assumption is loosely based on Acemoglu et al. (2020, Figure 7). Note: after the analyses for this Article were run, Acemoglu et al. released a new version of their paper, which (as of the date of this Article, June 25, 2020) is the posted version. (Acemoglu et al. 2020b). The new Acemoglu et al. estimates of the GDP reduction needed to wholly eliminate the COVID-19 risk are more optimistic—closer to 40%. See Acemoglu et al. (2020b, Figure 5.1). Because the intention of this Article is not to provide guidance with respect to COVID-19 policy, but rather to compare VSL-based CBA with utilitarianism and prioritarianism, I have retained the 80% figure roughly based on Acemoglu (2020 Figure 7) as the base case for analysis. In any event, the 40% figure is covered here as the "optimistic case" and still shows significant differences between VSL-based CBA, on the one hand, and utilitarianism and prioritarianism, on the other.
[41] See Acemoglu et al. (2020 Figure 7).

risk (optimistic), and an 80% reduction only reduces the COVID-19 risk by half. I also consider a third alternative ("convex case"): a 80% reduction in GDP wholly eliminates the COVID-19 risk, and the function from GDP reduction to COVID-19 reduction is convex (increasing marginal impact of GDP reduction on the risk reduction) rather than concave.

On the cost side, I consider two alternative cases: "regressive incidence" and "proportional incidence." "Regressive incidence" means a π% reduction in GDP is borne more heavily (in proportional terms) by the three lowest income groups within each age, than by the two higher quintiles.[42] Proportional incidence means that individuals in all five quintiles incur a π% reduction in income when GDP is reduced by π%.

The causal linkages between social distancing policy and GDP reduction, on the one hand, and between social distancing policy and reduction in COVID-19 risk, on the other, are quite uncertain—and certainly were highly uncertain at the beginning of the pandemic, which is the timing of the policy choice being modeled here. (The full spectrum of COVID-19 risk-reduction options, including the option of wholly eliminating the risk via sufficiently stringent social-distancing measures, were only available to policymakers at the outset of the pandemic.) The assumptions I have adopted are not meant as best estimates, but rather as plausible, stylized facts that will illustrate the properties of the various assessment tools with respect to social distancing policy. As already explained, the aim of this Article is methodological: not to provide guidance with respect to social-distancing policy, but to use the simulation model to illustrate the features of VSL.

As illustrated in Table 7, the various cohorts have starkly opposing interests with respect to social distancing policy. Table 7 shows the cohort breakeven GDP reduction (the largest reduction such that its expected utility is not lower than baseline[43]), depending on whether the cost incidence of the reduction is regressive or proportional. (In each cell, the number above the slash is the cohort breakeven GDP reduction assuming regressive incidence; the number below the slash is for proportional incidence.)

---

[42] Specifically, my assumptions regarding regressive incidence are as follows. A π% reduction in GDP lowers income for individuals in the first three income quintiles (Low, Moderate and Middle) by $\pi/80$ of the difference between baseline income and subsistence income ($1000). Thus, an 80% reduction in GDP has the benefit of wholly eliminating the COVID-19 risk, but at the cost of reducing individuals in the first three quintiles to subsistence income. As for the two higher-income groups: a π% reduction in GDP reduces their income by $(\pi/80)p$, where p is chosen so that—given the above assumption about losses by the first three income quintiles—the overall reduction in GDP is π%. With the income data being used, $p = 0.7187$.

On these assumptions, each of the first three income quintiles experiences a larger percentage reduction in income than the top two quintiles, for $0 < \pi \leq 80$. "Regressive" is a slight misnomer, since the lowest quintile experiences a slightly lower percentage reduction than the second, in turn slightly lower than the third.

[43] Results for the simulation model were calculated for integer values of GDP reduction, and so the cohorts may be better off than baseline at these integer "breakeven" values. It would also interesting to know the cohort-*optimal* values of GDP reduction, but determining these values for all thirty-five cohorts would be quite laborious.

<div align="center">Table 7: Cohort Breakeven GDP Reductions</div>

	Low	*Moderate*	*Middle*	*High*	*Top*
20	0/0	0/0	0/0	0/0	0/0
30	0/1	0/1	1/1	2/2	4/3
40	1/2	2/3	2/4	6/4	9/7
50	12/17	16/22	19/28	39/33	53/45
60	45/58	52/68	58/75	80/80	80/80
70	61/76	67/80	71/80	80/80	80/80
80	63/75	69/80	73/80	80/80	80/80

Older cohorts prefer a larger GDP reduction than younger ones. Within each age band, richer individuals prefer a larger GDP reduction than younger ones. This latter effect is dampened by shifting from regressive to proportional incidence, but not eliminated.

The choice of social distancing policy is a matter of balancing the (starkly) opposing interests of the various cohorts. VSL-CBA reaches a very different point of equipoise in this balancing, as compared with utilitarianism and prioritarianism.

Table 8 gives the optimal degree of reduction in GDP, according to the three methodologies—in the base case; under alternative assumptions (more optimistic or less optimistic) regarding the efficacy of social distancing policy in eliminating COVID-19 risk; and in the convex case. [44]

<div align="center">Table 8: Optimal reduction in GDP</div>

	VSL-CBA	Utilitarianism	Prioritarianism
Base Case- Regressive Incidence	27%	11 %	12 %
Base Case – Proportional Incidence	27%	13%	14%
Optimistic Case— Regressive Incidence	40%	19%	19%
Optimistic Case – Proportional Incidence	40%	21%	23%
Pessimistic Case— Regressive Incidence	7%	4%	4%
Pessimistic Case – Proportional Incidence	7%	4%	4%
Convex Case— Regressive Incidence	80%	0%	0%
Convex Case— Proportional Incidence	80%	0%	0%

---

[44] For utilitarianism and prioritarianism, the optimal reduction was calculated exactly rather than using the marginal approximations stated in equations 2 and 4 above.

Let's first compare VSL-CBA to utilitarianism. In the base case, utilitarianism recommends a 11% reduction with regressive incidence and a 13% reduction with proportional incidence. The utilitarian social cost to a given reduction in income is lower if that loss is borne by higher-income groups (see Table 5b) and so shifting from regressive to proportional incidence increases the utilitarian optimum. VSL-CBA is insensitive to cost distribution (see Equation 3 and Table 5a), and so the optimum is a 27% reduction regardless of whether incidence is regressive or proportional. In either event, VSL-CBA's optimal reduction of GDP is *much* larger (27% versus 11%/13%) than the utilitarian optimum. Shifting to the optimistic and pessimistic cases does not change the pattern. VSL-CBA continues to prefer a much larger reduction than utilitarianism (40% versus 19%/21%; 7% versus 4%/4%).

Finally, in the convex case, the disparity between VSL-CBA and utilitarianism is quite extreme. The two methodologies choose corner solutions, but different ones: utilitarianism (with either regressive or proportional cost incidence) recommends no reduction in GDP at all, while VSL-CBA recommends wholly eliminating the risk at the cost of an 80% GDP reduction.

Why is VSL-CBA more stringent with respect to social distancing than utilitarianism?[45] There are two reasons, the first somewhat subtle. Let π denote the degree of GDP reduction, and let $\Delta p_c(\pi)$ denote the reduction in cohort c's fatality risk as a function of π; $\Delta y_c(\pi)$ the reduction in cohort c's income as a function of π; and $\Delta U_c(\pi)$ the increase in cohort c's expected utility. As per Equation 2, $\Delta U_c(\pi) \approx \frac{\partial U_c}{\partial p_c} \Delta p_c(\pi) - \frac{\partial U_c}{\partial y_c} \Delta y_c(\pi)$. As per Equations 1 and 3, the VSL-based monetary value of a π reduction in GDP is as follows:

$MV_c(\pi) = \Delta p_c(\pi)VSL_c - \Delta y_c(\pi) = \Delta p_c(\pi)\left(\frac{\partial U_c/\partial p_c}{\partial U_c/\partial y_c}\right) - \Delta y_c(\pi)$. Multiplying both sides of this

equation by $(\partial U_c/\partial y_c)$, we have that $MV_c(\pi)\frac{\partial U_c}{\partial y_c} = \Delta p_c(\pi)\frac{\partial U_c}{\partial p_c} - \Delta y_c(\pi)\frac{\partial U_c}{\partial y_c}$, or that

$MV_c(\pi) \approx \Delta U_c(\pi)k_c$, with $k_c = 1/\left(\frac{\partial U_c}{\partial y_c}\right)$. In short, a cohort's VSL-based monetary valuation of

a policy is approximately its change in expected utility, times a weighting factor equaling the inverse of the marginal utility of income.

The optimal value of π, according to utilitarianism, is the value that maximizes the sum of ΔU_c across cohorts; while its optimal value, according to VSL-CBA, is the value that maximizes the sum of MV_c across cohorts. The k_c weighting factor increases with income. Thus VSL-CBA in this optimization gives greater weight to utility impacts on higher-income cohorts, and less weight to utility impacts on lower-income cohorts, as compared to utilitarianism.

---

[45] For a related discussion of why CBA and utilitarianism can diverge with respect to risk regulation, see Armantier and Treich (2004).

Note now that, with regressive and even with proportional incidence, richer cohorts tend to prefer a more restrictive social distancing policy than poorer cohorts. (See Table 7.)[46] Thus the effect of this skew towards the interests of higher-income cohorts is that VSL-CBA ends up with a more restrictive policy than utilitarianism.

The second reason that VSL-CBA is more stringent with respect to social distancing than utilitarianism is that VSL-CBA overstates the benefits of "large" policies, which reduce risks at the cost of significant changes in individual incomes.[47]

Turning now to prioritarianism: although in many contexts utilitarianism and prioritarianism can differ, here they converge. In all cases (base, optimistic, pessimistic, and convex, with proportional or regressive incidence), the optimum reduction in GDP as per

---

[46] Why do richer cohorts prefer a more restrictive social distancing policy than poorer cohorts? Shutdown policy is (approximately) a uniform risk reduction across income groups. Thus, on the risk side, each increment in π produces a larger increase in the expected utility of a richer group as compared to a poorer group. (See Table 4b.) On the cost side, whether each increment in π yields a smaller or larger loss in the expected utility of richer groups depends upon cost incidence. With a logarithmic utility function and proportional incidence, a given increment in π yields the same utility cost for different income groups within each age band. With regressive incidence, a given increment yields a larger utility cost for lower income groups. In either event, the *net* utility impact of an increment in π (utility benefit from risk reduction minus utility cost from income loss) is always larger for richer than smaller groups. Thus richer groups tend to prefer a larger π than poorer groups, as illustrated in Table 7.

[47] Consider the additive utility function used in the simulation; see Appendix B. Let A_c denote the number of the current period for cohort c. As in the text, let p_c denote the cohort's current survival probability (shorthand for $p_c^{A_c}$) and let y_c denote the cohort's current income ($y_c^{A_c}$). So as to avoid using the symbol π to mean two different things, let $\delta_c(t; A_c) = \prod_{s=A_c}^{t} p_c^s$ here denote the probability of a cohort member surviving to the end of period t.

$$\Delta U_c(\pi) = \left[\Delta p_c(\pi) u(y_c - \Delta y_c(\pi)) + \Delta p_c(\pi) \sum_{t=A_c+1}^{T} \frac{\delta_c(t; A_c)}{p_c} u(y_c^t) \right] - p_c\left(u(y_c) - u(y_c - \Delta y_c(\pi)) \right).$$ Note that, with this additive model, $\frac{\partial U_c}{\partial p_c} = u(y_c) + \sum_{t=A_c+1}^{T} \frac{\delta_c(t; A_c)}{p_c} u(y_c^t)$ and $\frac{\partial U_c}{\partial y_c} = p_c u'(y_c)$. Assume that $u(\cdot)$ is strictly increasing and also strictly concave (diminishing marginal income utility). Observe that the term in brackets in the formula for $\Delta U_c(\pi)$ is less than $\frac{\partial U_c}{\partial p_c} \Delta p_c(\pi)$ and that $p_c(\pi)\left(u(y_c) - u(y_c - \Delta y_c(\pi)) \right)$ is greater than $\frac{\partial U_c}{\partial y_c} \Delta y_c(\pi)$ by the strict concavity of $u(\cdot)$.

Thus, $\frac{\partial U_c}{\partial p_c} \Delta p_c(\pi) - \frac{\partial U_c}{\partial y_c} \Delta y_c(\pi) = \Delta U_c(\pi) + E$, with E a positive error term the magnitude of which increases with $\Delta y_c(\pi)$. In turn, $MV_c(\pi) = (\Delta U_c(\pi) + E)k_c$. In short, MV_c is skewed upwards relative to ΔU_c, and increasingly so as Δy_c becomes large.

prioritarianism is the same as, or quite close to, the utilitarian optimum—and thus much lower than VSL-CBA's recommendation.[48]

In short, relative to the utilitarian benchmark and to a prioritarian benchmark, VSL-CBA is significantly more restrictive.[49]

Can CBA's recommendations with respect to shutdown policy be defended on alternative grounds—by appeal to Kaldor-Hicks efficiency? I calculated cohorts' actual willingness to pay/accept amounts at the CBA optimum in the base case (27% reduction). At the optimum, all the cohorts aged 40 and younger are worse off than the status quo, as are the three age-50 cohorts at low, moderate and middle income. All the cohorts aged 60 and above are better off than the status quo, as are the two higher-income age-50 cohorts.[50] Total willingness-to-pay of the better-off cohorts does in fact exceed the total willingness-to-accept of the worse-off cohorts. So the optimum is Kaldor-Hicks efficient relative to the no-intervention baseline.

However, many of the non-optimal policies that CBA scores as an improvement relative to the status quo are *not* Kaldor-Hicks efficient. In the model here, *every* reduction in GDP (from 1% to 80%) is assigned a positive score (positive sum of MVs), relative to the baseline of 0% reduction. In other words, every reduction is seen by CBA as an improvement over the status quo. However, reductions at or above 37% are not Kaldor-Hicks efficient. In particular, CBA prefers to wholly eliminate the COVID-19 risk, at the cost of 80% of GDP, as compared to the no-intervention status quo; but at 80% reduction, only 6 of the 35 cohorts are better off than

---

[48]Prioritarianism is more concerned than utilitarianism with reducing risks among the young. (Compare Tables 4b and 4c.) Baseline COVID-19 risk increases with age, and thus the effect of any COVID-19 policy (in the model here) is to produce a smaller risk reduction for younger than older cohorts. These effects roughly cancel, and so prioritarianism ends up with roughly the same social distancing recommendation as utilitarianism.

As mentioned, the prioritarian results in the social-distancing analysis are for "ex post" prioritarianism. As a sensitivity analysis, I computed optimal GDP reduction for "ex ante" prioritarianism in the base case. Ex ante prioritarianism prefers a 9% reduction given regressive incidence and an 11% reduction given proportional incidence.

[49] A third possible source of the divergence between VSL-CBA and utilitarianism with respect to social distancing policy is the fact that the two methodologies, in the model here, have different patterns of valuation of risk reduction with age. Utilitarian values decrease monotonically, while VSL-CBA values do not. (Compares Tables 4a and 4b.) To test whether this divergence helps explain the divergent social-distancing recommendations, I undertook an alternative analysis in which the time path of income within each quintile is constant. At every age, individuals receive the very same income, depending on quintile: $21,961 (low); $30,118 (moderate); $41,349 (middle); $57,538 (high); and $134,840 (top). In this case, I find that utilitarianism with regressive incidence recommends a 12% GDP reduction, and with proportional incidence a 14% GDP reduction, while VSL-CBA recommends a 29% reduction. Thus the differing patterns of risk valuation by age do not appear to contribute to the VSL-CBA/utilitarianism divergence with respect to social distancing.

[50] These willingness-to-pay/accept amounts are the individuals' compensating variations (equilibrating changes to the incomes that individuals would have with the policy). It is the sum of compensating variations which tracks whether a policy is Kaldor-Hicks efficient relative to the status quo: if the sum is positive, there is a scheme of transfers which, if costlessly implemented together with the policy, would make everyone better off than in status quo.

Compensating variations were calculated assuming regressive incidence. With progressive incidence, a different array of groups might be better off than the status quo (compare the numbers above and below the slash in each cell in Table 7), but the sum total of compensating variations would be the same.

the status quo (high and top-income individuals aged 60, 70 and 80) and their total willingness to pay for 80% shutdown is insufficient to compensate the 29 cohorts who are worse off.

IV.   Population-Average VSL

The previous two parts focused on textbook VSL-based CBA ("VSL-CBA"). VSL varies among individuals (Equation 1) and these individual-specific VSLs are used to compute MVs.

As already mentioned, governmental agencies in the U.S. employ a single population-average VSL. I'll refer to CBA with a population-average VSL as "VSL^{avg}-CBA." It uses this single conversion factor to calculate MVs.

While MV$_i$ calculated with individual i's VSL is a good approximation to her compensating variation and equivalent variation for a small change in survival probability and other attributes, relative to the status quo, this is not true of MV$_i$ calculated with a population-average VSL. The point is well illustrated with the simulation model presented in Part III, which I'll continue to use in this Part. The cohort-specific VSLs are given in Table 3. The population-average VSL (weighting by the proportion of the various age groups in the population) is $8,635,355. The cohort-specific VSLs are generally significantly different (larger or smaller) than this average. For example, a 60-year old middle-income individual has a VSL of $3,943,065. That individual's equivalent variation and compensating variation for a 1-in-1 million reduction in current fatality risk is approximately $3.94, and for a 1-in-100,000 reduction is approximately $39.43. But his MV according to VSL^{avg}-CBA is $8.64 in the first case and $86.35 in the second.

We saw earlier that VSL-CBA need not conform to Kaldor-Hicks efficiency. The divergence between VSL-CBA and Kaldor-Hicks efficiency occurs as MV$_i$ diverges from i's equivalent and compensating variations. VSL^{avg}-CBA also need not conform to Kaldor-Hicks efficiency, and can diverge even for small changes from the status quo (since MV$_i$ calculated with population-average VSL is not a good approximation to the equivalent or compensating variation even for a small change). For example, imagine that there are 1 million individuals in the 40-year old, low-income cohort (with a cohort VSL of $3,254,671), and a policy will reduce each individual's risk by 1 in 100,000, at a total cost of $50 million. VSL^{avg}-CBA approves the policy, since the population-average VSL of $8,635,355 multiplied by 1-in-100,000 summed over the million individuals equals $86.35 million, which exceeds the cost of $50 million. But individuals in the cohort are actually willing to pay, in total, only $32.55 million for the policy. Whoever the cost bearers might be, $32.55 million would be insufficient to fully compensate them.

Table 4d is the analogue to tables 4a, 4b, and 4c. It gives the social value of an increment in current-year fatality risk according to VSL^{avg}-CBA, with 1 indicating the value of a risk increment to the member of the 80-year-old, low-income cohort.

Table 4d: Social value of risk increment according to VSL^{avg}-CBA
(relative to social value of risk increment for 80-year-old, low-income cohort)

	Income: Low	*Moderate*	*Middle*	*High*	*Top*
Age: 20	1	1	1	1	1
30	1	1	1	1	1
40	1	1	1	1	1
50	1	1	1	1	1
60	1	1	1	1	1
70	1	1	1	1	1
80	1	1	1	1	1

The entry in every cell in table 4d is the same (1), because a risk increment Δp to any cohort is assigned the very same value, namely Δp multiplied by $8,635,355.

On the income side, VSL^{avg}-CBA values an income increment the same way as VSL-CBA—a Δy increment to anyone's income changes her MV by Δy—and so the relative value of income increments is given by the same table for both methodologies, namely table 5a.

The risk-increment and income-increment tables show that VSL^{avg}-CBA diverges radically from utilitarianism and prioritarianism. Nor, as just explained, can it be supported with reference to Kaldor-Hicks efficiency.

Considerations of political feasibility may well push governments towards VSL^{avg}-CBA. In the U.S., a tentative move by the EPA in 2003 in the direction of adjusting VSL by age prompted a harsh political reaction by interest groups representing older individuals.[51] But such political constraints of course don't constitute a normative justification for VSL^{avg}-CBA.

In some respects, the methodology's schedule of relative risk valuations (table 4d—every group has the same value of risk reduction) is intuitively attractive. Surveys eliciting citizens' views regarding the allocation of health care show no support for preferring the rich.[52] In this respect, Table 4d is more intuitive than Table 4a (VSL-CBA) or Table 4b (utilitarianism). Moreover, Table 4d conforms to the intuition that every life has equal value—that there should be no differentiation whatsoever in deciding whose life to save or risk to reduce. Yet surveys of citizens also find support for a conflicting intuition, namely that the young should receive priority over the old.[53] Is it really the case that society should be indifferent between reducing an 80-year-old's risk of dying by Δp and reducing a 20-year-old's risk of dying by the same amount? So intuitive support for Table 4d can hardly be said to be rock-solid.

---

[51] This episode is discussed in Viscusi (2018, p. 93).

[52] See Dolan et al. (2005).

[53] See *id*.

On the cost side, VSL^{avg}-CBA, like VSL-CBA, is neutral to the distribution of the costs of risk reduction (table 5a). This is not especially intuitive.

Turning to the model social-distancing policy: VSL^{avg}-CBA is dramatically *more* restrictive than VSL-CBA, let alone utilitarianism and prioritarianism. In the base case, it recommends an 80% reduction, while VSL-CBA recommends 27%. In the pessimistic case, it recommends a 25% reduction, while VSL-CBA recommends 7%. (See Table 10.) VSL^{avg}-CBA advises a substantially greater reduction in risk and loss of income than VSL-CBA because the benefits of social distancing are skewed towards older cohorts, whose cohort-specific VSLs are generally less than the population average. (VSL^{avg}-CBA reaches the same result as VSL-CBA in the optimistic and convex cases.)

Recall that VSL-CBA's recommendation (a 27% reduction) is Kaldor-Hicks efficient relative to the status quo. VSL^{avg}-CBA's recommendation (an 80% reduction) is not.

V.   VSLY

The value of statistical life year (VSLY) is VSL divided by the life expectancy gained from a risk reduction. Life expectancy gained might be discounted at the market interest rate or the individual rate of time preference, or undiscounted. The conceptual points here apply to both discounted and undiscounted VSLY.[54] The specific illustrations, drawn, from the simulation model, involved undiscounted VSLY.

At the individual level, VSL with respect to a particular time period can be expressed as such, or as VSLY. Consider a policy that increases individual i's period s survival probability by $\Delta p_i{}^s$. Recall that individual i's VSL for that period, VSL_i^s, is equal to her marginal rate of substitution between survival probability in period s and current income (Equation 1). MV_i in this case equals $(\Delta p_i^s)VSL_i^s$, which is a good approximation to her equivalent variation and compensating variation for the policy if $\Delta p_i{}^s$ is small.

Let LE_i^s be the difference between (a) i's current life expectancy if her probability of surviving period s conditional on being alive at the beginning of s is 1, and (b) her current life expectancy if her probability of surviving period s conditional on being alive at the beginning of s is 0. Note that $(\Delta p_i^s)LE_i^s$ is the gain to current life expectancy from a $\Delta p_i{}^s$ increase in period s survival probability. Let's now define $VSLY_i^s$ as VSL_i^s / LE_i^s. Then, by construction, $MV_i = (\Delta p_i^s)VSL_i^s = (\Delta p_i^s)LE_i^s \times VSLY_i^s$. Individual i's MV for a period s risk reduction can be expressed *either* as her VSL for that period multiplied by the risk reduction, *or* as the gain to life expectancy from the risk reduction, multiplied by her VSLY for that period. This period-and-individual specific VSLY is well-defined, but has no advantages (or disadvantages!) over the

---

[54] The conceptual points here about the relation between VSL and VSLY are not novel; see Hammitt (2007) and Jones-Lee et al. (2015).

period-and-individual specific VSL, since monetary valuations calculated either way are identical.

It might be thought that we could calculate a period-*invariant* individual VSLY (let's call it VSLY$_i$), such that an individual's VSL with respect to the current period or any future period is just equal to her life expectancy gain from surviving that period, multiplied by VSLY$_i$. That is: for any period s, $VSL_i^s = VSLY_i \times LE_i^s$. In other words, an individual's willingness to pay for a small risk reduction in the current period or any future period is the gain to life expectancy, multiplied by a constant individual VSLY. However, there is nothing in the theory of VSL to ensure that this is true. For example, imagine that an individual consumes income when she receives it, and that the time path of future income is variable. Between period s^* and s^{**}, her income will be higher than after s^{**}. If so, the individual will be willing to pay *more* in current income for a given increase in life expectancy secured through an increase in period s^* survival probability, than for the same increase in life expectancy secured through an increase in period s^{**} survival probability. There is *no* period invariant VSLY$_i$ such that $VSL_i^{s^*} = VSLY_i \times LE_i^{s^*}$ and $VSL_i^{s^{**}} = VSLY_i \times LE_i^{s^{**}}$.[55]

In short, at the individual level VSLY is either equivalent to VSL, or undefined.

At the population level, however, we can define a population-average VSLY that is well defined and that gives rise to a form of CBA—for short, "VSLY^{avg}-CBA"—that differs from both CBA with individual-specific VSLs (VSL-CBA) and CBA with a population-average VSL (VSL^{avg}-CBA). VSLY^{avg} for a given period is just the average of individual VSLYs for that period. VSLY^{avg}-CBA assigns a monetary valuation to an individual's risk reduction in period s equaling the increase in life expectancy multiplied by VSLY^{avg}.

We can use the simulation model to illustrate. Table 9 gives the life-expectancy gains from preventing the current death of an individual in each of the 35 cohorts (that is, LE_i^s with s the current period).

---

[55] In the case at hand, $VSL_i^{s^*} / LE_i^{s^*} > VSL_i^{s^{**}} / LE_i^{s^{**}}$. On a related point, Hammitt (2007, p. 236) observes: "If [as is often observed] an individual's VSL first rises then falls with age, then her VSLY cannot be constant over her lifespan. Life expectancy typically decreases with age."

Table 9: Life expectancy gains (years) from saving a cohort member in the current year

	Income: Low	Moderate	Middle	High	Top
Age: 20	54.3	56.8	58.9	60.1	62.1
30	45.1	47.5	49.5	50.6	52.5
40	36.1	38.4	40.2	41.3	43.1
50	27.4	29.5	31.2	32.2	33.9
60	19.6	21.4	23.0	23.8	25.3
70	12.8	14.2	15.4	16.2	17.4
80	7.1	8.2	9.1	9.6	10.5

For each cohort, we divide the VSL (Table 3) by the life expectancy gain in Table 9 to arrive at a cohort VSLY, and then average these (weighting for the proportion of the different age groups in the population) to arrive at the population-average VSLY, which is $240,676.

An individual's monetary valuation calculated using the population average VSLY need not approximate her equivalent or compensating variation, even for small changes. This is because the individual's VSLY for the specific period (here, the current period) need not be the same as the population-average VSLY for that period. For example, VSL for a member of the 20 year old, high-income cohort is $5,240,842. Her VSLY is $5,240,842/60.1 = $87,202. Her equivalent variation and compensating variation for a 1-in-100,000 increase in survival probability is approximately $52, which is her VSL times 1-in-100,000 or, equivalently, her VSLY times 1-in-100,000 times the life expectancy gain from preventing her death in the current period (60.1 years). However, the population-average VSLY ($240,676) multiplied by 1-in-100,000 multiplied by 60.1 years gives a monetary value of $145.

Thus VSLY^{avg}-CBA need not conform to Kaldor-Hicks efficiency, even for small risk changes (like VSL^{avg}-CBA, but by contrast with VSL-CBA).

The relative social values of risk reduction implied by VSLY^{avg}-CBA are presented in Table 4e. As in the previous such tables, the values are normalized so that 1 indicates the value of risk reduction to a member of the 80-year-old, low-income cohort. These relative valuations are just the relative life expectancy gains in Table 9.

Social value of risk increment according to VSLY^{avg}-CBA
(relative to social value of risk increment for 80-year-old, low-income cohort)

	Income: Low	Moderate	Middle	High	Top
Age: 20	7.6	8.0	8.3	8.4	8.7
30	6.3	6.7	6.9	7.1	7.4
40	5.1	5.4	5.6	5.8	6.0
50	3.8	4.1	4.4	4.5	4.8
60	2.8	3.0	3.2	3.3	3.6
70	1.8	2.0	2.2	2.3	2.4
80	1.0	1.1	1.3	1.3	1.5

These relative valuations are quite different from those in Tables 4b (utilitarianism) and 4c (prioritarianism); and on the income side VSLY^{avg}-CBA has a constant valuation (Table 5a), so is quite different from utilitarian and prioritarian valuations of income increments (Tables 5b and 5c). In short, VSLY^{avg}-CBA is grounded neither in utilitarianism, nor prioritarianism, nor (as we've just seen) Kaldor-Hicks efficiency. Why adopt it?

VSLY^{avg} –CBA conforms to the intuition that the social value of risk reduction should be lower for older individuals, since they have fewer life years remaining. (The valuations in Table 4e decrease according to life expectancy remaining.) If one has this intuition, VSLY^{avg}-CBA will be seen to be an improvement over VSL^{avg}-CBA. Further, it is dramatically less biased to the rich than VSL-CBA (Table 4a), and significantly less so than utilitarianism (Table 4b). Values increase with income in each age band in Table 4e only because higher income is associated with a more favorable survival curve and thus higher life expectancy—and not for the utilitarian reason that a year of life at higher income is better for well-being.

Note, however, that prioritarianism can avoid *any* bias towards the rich (see Table 4c) and in this sense has intuitive advantages over VSLY^{avg}-CBA.

With respect to social distancing policy, VSLY^{avg}-CBA is somewhat less restrictive than VSL-CBA in the base case. It recommends a 23% reduction rather than 27%. However, this recommendation is significantly more restrictive than utilitarianism and prioritarianism, which (recall) recommend at most 14% depending on incidence. In the pessimistic case, VSLY^{avg}-CBA is again somewhat less restrictive than VSL-CBA (6% versus 7%). In the optimistic and convex cases the two concur.

As regards Kaldor-Hicks efficiency, the recommendation of VSLY^{avg}-CBA (23% reduction) is Kaldor-Hicks efficient relative to the status quo. However, VSLY^{avg}-CBA scores every reduction as an improvement over the status quo. Recall that reductions at or above 37% are not Kaldor-Hick efficient.

Table 10 summarizes the recommendations of all the methodologies in all the cases, and Table 11 summarizes the results with respect to Kaldor-Hicks efficiency.

Table 10. Social Distancing Policy: Optimal Degree of GDP Reduction

	VSL-CBA	VSL^{avg}-CBA	VSLY^{avg}-CBA	Utilitarianism	Prioritarianism
Base Case-Regressive Incidence	27%	80%	23%	11 %	12 %
Base Case – Proportional Incidence	27%	80%	23%	13%	14%
Optimistic Case—Regressive Incidence	40%	40%	40%	19%	19%
Optimistic Case – Proportional Incidence	40%	40%	40%	21%	23%
Pessimistic Case—Regressive Incidence	7%	25%	6%	4%	4%
Pessimistic Case – Proportional Incidence	7%	25%	6%	4%	4%
Convex Case—Regressive Incidence	80%	80%	80%	0%	0%
Convex Case—Proportional Incidence	80%	80%	80%	0%	0%

Table 11. Kaldor-Hicks Efficiency (Base Case)

	VSL-CBA	VSL^{avg}-CBA	VSLY^{avg}-CBA
Optimal degree of GDP reduction	27%	80%	23%
Is the optimum Kaldor-Hicks efficient relative to the status quo?	Yes	No	Yes
Does the methodology assign a positive score (indicating that the GDP reduction is an improvement over the status quo) to GDP reductions that are not Kaldor-Hicks efficient relative to the status quo?	Yes (all reductions at or above 37%)	Yes (all reductions at or above 37%)	Yes (all reductions at or above 37%)

VI.   Conclusion

This Article has critically examined the use of VSL-based CBA as a methodology for setting fatality risk reduction policy. I have done so via close consideration of a policy problem that, sadly, is all too timely: trading off the costs and benefits of social distancing to mitigate COVID-19. I consider three versions of VSL: textbook, population-average, and VSLY.

Strikingly, I find that all three recommend policies that *deviate* from Kaldor-Hicks efficiency (potential Pareto superiority), the traditional touchstone of CBA. All three versions of VSL conclude that a range of social distancing policies would be better than the status quo (in my simulation, policies at or above a 37% reduction of GDP), even though these policies are not in fact Kaldor-Hicks efficient relative to the status quo. CBA with population-average VSL *optimizes* at a policy that is not Kaldor-Hicks efficient relative to the status quo.

Thus these VSL-based methodologies *lack* grounding in the standard normative justification offered for CBA. Nor can they find normative support elsewhere. Utilitarianism, the oldest example of a systematic welfarist approach to normative reasoning, recommends quite different social-distancing policies—as does prioritarianism, a newer variant of welfarism. As for intuition: textbook VSL is quite counterintuitive, because it places a dramatically higher value on risk reductions for richer individuals; population-average VSL fails to differentiate with respect to age; and all three versions on the cost side are completely insensitive to the incidence of the costs of social distancing policy.

My own view is that social distancing policy, and risk regulation more generally, should be set with reference to a utilitarian or prioritarian social welfare function.[56] To be sure, this position implicates the long-running debate about the role of distributional considerations in non-tax policies.[57] The purist view, here, is that *all* non-tax policies should be designed to maximize the size of the "pie"—to be precise, the sum total of compensating variations relative to the status quo—and that the tax-and-transfer system should be used to share the "pie," to everyone's benefit.

There are various difficulties with the purist view,[58] above all this: the changes to existing tax-and-transfer laws required to render a non-tax policy Pareto-superior to the status quo— *universally* beneficial—are often merely hypothetical, given the actual political economy of the tax system. In the case of social distancing policy, a pie-maximizing-and-sharing combination of policies would mean quite stringent and prolonged social distancing requirements combined with significant taxes on older individuals and substantial payments to younger ones. This doesn't seem to be in the offing. Pie-maximization without pie-sharing may well leave some groups much worse off than the status quo.

---

[56] See Adler (2019, ch. 5); Adler (2017); Adler, Hammitt and Treich (2014): Adler, Ferranna, Hammitt and Treich (2019).
[57] See Kaplow (1996, 2004, 2008).
[58] See Adler (2019, pp. 225-33); Adler (2017).

In any event, a key takeaway from this Article should be that the debate about the role of distributional considerations in non-tax policy is *orthogonal* to VSL. The proponent of pie-maximization-and-sharing needs a methodology that approves policies if and only if they are Kaldor-Hicks efficient relative to the status quo, and that optimizes at the largest sum of compensating variations relative to the status quo. Such a methodology is not VSL-CBA, nor VSL^{avg}-CBA, nor VSLY^{avg}-CBA.

References

Acemoglu, Daron, et al. 2020. "A Multi-Risk SIR Model with Optimally Targeted Lockdown." NBER Working Paper No. 27102 (dated May 2020). Note: This paper was originally available at http://www.nber.org/papers /w27102, but has now been replaced at that link by Acemoglu et al. (2020b) immediately below. I have retained an electronic copy of Acemoglu et al. (2020) and can make it available on request.

Acemoglu, Daron, et al. 2020b. "Optimal Targeted Lockdowns in a Multi-Group SIR Model." NBER Working Paper No. 27102 (described on NBER site as a June 2020 revision, although dated May 2020 on the title page). Available at http://www.nber.org/papers /w27102.

Adler, Matthew D. 2012. *Well-Being and Fair Distribution*: *Beyond Cost-Benefit Analysis*. New York: Oxford University Press.

Adler, Matthew D. 2016a. "Benefit-Cost Analysis and Distributional Weights: An Overview." *Review of Environmental Economics and Policy* 10: 264-85.

Adler, Matthew D. 2016b. "Extended Preferences." In Matthew D. Adler and Marc Fleurbaey, eds., *The Oxford Handbook of Well-Being and Public Policy*, pp. 476-517. New York: Oxford University Press.

Adler, Matthew D. 2017. "A Better Calculus for Regulators: From Cost-Benefit Analysis to the Social Welfare Function." Duke Law School Public Law and Legal Theory Working Paper No. 2017-19. Available at https://papers.ssrn.com/sol3/papers.cm?abstract_id=2923829.

Adler, Matthew D. 2019. *Measuring Social Welfare*: *An Introduction*. New York: Oxford University Press.

Adler, Matthew D., Maddalena Ferranna, James K. Hammitt, and Nicolas Treich. 2019. "Fair Innings? The Utilitarian and Prioritarian Value of Risk Reduction over a Whole Lifetime." Working paper. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3493976.

Adler, Matthew D., James K. Hammitt, and Nicolas Treich. 2014. "The Social Value of Mortality Risk Reduction: VSL versus the Social Welfare Function Approach." *Journal of Health Economics* 35: 82-93.

Adler, Matthew D., and Eric A. Posner. 2006. *New Foundations of Cost-Benefit Analysis*. Cambridge, MA: Harvard University Press.

Aldy, Joseph E., and W. Kip Viscusi. 2007. "Age Differences in the Value of Statistical Life: Revealed Preference Evidence." *Review of Environmental Economics and Policy* 1: 241-60.

Alvarez, Fernando E., David Argente, and Francesco Lippi. 2020. "A Simple Planning Problem for Covid-19 Lockdown." NBER Working Paper No. 26981. Available at http://www.nber.org/papers/w26981.

Armantier, Olivier, and Nicolas Treich. 2004. "Social Willingness to Pay, Mortality Risks, and Contingent Valuation." *Journal of Risk and Uncertainty* 29: 7-19.

Bairoliya, Neha, and Ayşe Imrohoroğlu. 2020. "Macroeconomic Consequences of Stay-at-Home Policies During the COVID-19 Pandemic." Working paper. Available at https://www.marshall.usc.edu/sites/default/files/bairoliy/intellcont/macroeffects_of_corona-1.pdf.

Barnett-Howell, Zachary, and Ahmed Mushfiq Mobarak. 2020. "The Benefits and Costs of Social Distancing in Rich and Poor Countries." Working paper. Available at https://arxiv.org/abs/2004.04867.

Béland, Louis-Phillippe, Abel Brodeur, and Taylor Wright. 2020. "Covid-19, Stay-at-Home Orders and Employment: Evidence from CPS Data." IZA Discussion Paper No. 13282. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3608531.

Bethune, Zachary A., and Anton Korinek. 2020. "Covid-19 Infection Externalities: Trading Off Lives vs. Livelihoods." NBER Working Paper No. 27009. Available at http://www.nber.org/papers/w27009.

Boadway, Robin. 2016. "Cost-Benefit Analysis." In Matthew D. Adler and Marc Fleurbaey, eds., *The Oxford Handbook of Well-Being and Public Policy*, pp. 47-81. New York: Oxford University Press.

Broome, John. 1978. "Trying to Value a Life." *Journal of Public Economics* 9: 91-100.

Bronsteen, John, Christopher Buccafusco, and Jonathan Masur. 2013. "Well-Being Analysis vs. Cost-Benefit Analysis." *Duke Law Journal* 62: 1603-89.

Broughel, James. 2020. "Rethinking the Value of Life: A Critical Appraisal of the Value of a Statistical Life." Center for Growth and Opportunity, Policy Paper. Available at https://www.thecgo.org/research/rethinking-the-value-of-life/.

Chetty, Raj, et al. 2016. "The Association between Income and Life Expectancy in the United States, 2001-2014." *JAMA* 315: 1750-66.

Corcoran, Terence. May 15, 2020. "The Price of Life: Lockdown Costs are Real, but are the Benefits?" National Post (Canada).

Cropper, Maureen, James K. Hammitt, and Lisa A. Robinson. 2011. "Valuing Mortality Risk Reductions: Progress and Challenges." *Annual Review of Resource Economics* 3: 313-36.

Dolan, Paul, et al. 2005. "QALY Maximisation and People's Preferences: A Methodological Review of the Literature." *Health Economics* 14: 197-208.

Dolan, Paul, et al. 2008. "Valuing Lives and Life Years: Anomalies, Implications, and an Alternative." *Health Economics, Policy, and Law* 3: 277-300.

Eeckhoudt, Louis R., and James K. Hammitt. 2001. "Background Risks and the Value of Statistical Life." *Journal of Risk and Uncertainty* 23: 261-79.

Evans, Mary F., and V. Kerry Smith. 2010. "Measuring How Risk Tradeoffs Adjust with Income." *Journal of Risk and Uncertainty* 40: 33-55.

Ferguson, Neil M., et al. 2020. "Impact of Non-Pharmaceutical Interventions (NPIs) to Reduce Covid-19 Mortality and Healthcare Demand." Imperial College, London, working paper. Available at https://doi.org/10.25561/77482.

Freeman, A. Myrick. 2003. *The Measurement of Environmental and Resource Values*, 2d ed. Washington, D.C: Resources for the Future.

Frick, Johann. 2015. "Contractualism and Social Risk." *Philosophy and Public Affairs* 43: 175-223.

Greenstone, Michael, and Visham Nigam. 2020. "Does Social Distancing Matter?" Working paper. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3561244.

Gros, Claudius, et al. 2020. "Containment Efficiency and Control Strategies for the Corona Pandemic Costs." Working paper. Available at https://arxiv.org/pdf/2004.00493.pdf.

Grüne-Yanoff, Till. 2009. "Mismeasuring the Value of Statistical Life." *Journal of Economic Methodology* 16: 109-23.

Hall, Robert E., Charles I. Jones, and Peter J. Klenow. 2020. "Trading Off Consumption and COVID-19 Deaths." Working paper. Available at https://web.stanford.edu/~chadj/Consumption_v_Covid.pdf.

Hammitt, James K. 2000. "Valuing Mortality Risk: Theory and Practice." *Environmental Science and Technology* 34: 1396-1400.

Hammitt, James K. 2007. "Valuing Changes in Mortality Risk: Lives Saved versus Life Years Saved." *Review of Environmental Economics and Policy* 1: 228-40.

Hammitt, James K. 2020. "Valuing Mortality Risk in the Time of Covid-19." Working paper. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3615314.

Hammitt, James K., et al. 2020. "Premature Deaths, Statistical Lives, and Years of Life Lost: Identification, Quantification, and Valuation of Mortality Risks." *Risk Analysis* 40: 674-95.

Henderson, David R., and Jonathan Lipow. June 16, 2020. "The Data are In: It's Time for Major Reopening." *Wall Street Journal*.

Hilsenrath, Jon, and Stephanie Armour. March 23, 2020. "As Economic Toll Mounts, Nation Ponders Trade-Offs." *Wall Street Journal*.

Holtug, Nils. 2010. *Persons, Interests, and Justice*. Oxford: Oxford University Press.

Holtug, Nils. 2017. "Prioritarianism." *Oxford Research Encyclopedia of Politics*, http://dx.doi.org/10.1093/acrefore/9780190228637.013.232.

Ingraham, Christopher. March 30, 2020. "Moderate Social Distancing Yields $8 Trillion in Economic Benefits, Study Finds." *Washington Post*.

Jenkins, Holman W., Jr. March 27, 2020. "Yes, We Need Cost-Benefit Analysis." *Wall Street Journal*.

Johansson, Per-Olov. 2002. "On the Definition and Age-Dependency of the Value of Statistical Life." *Journal of Risk and Uncertainty* 25: 251-63.

Jones-Lee, Michael, et al. 2015. "Valuing Gains in Life Expectancy: Clarifying Some Ambiguities." *Journal of Risk and Uncertainty* 51: 1-21.

Kaplow, Louis. 1996. "The Optimal Supply of Public Goods and the Distortionary Cost of Taxation." *National Tax Journal* 49: 513-33.

Kaplow, Louis. 2004. "On the (Ir)Relevance of Distribution and Labor Supply Distortion to Government Policy." *Journal of Economic Perspectives* 18: 159-75.

Kaplow, Louis. 2008. *The Theory of Taxation and Public Economics*. Princeton: Princeton University Press.

Kniesner, Thomas J., and W. Kip Viscusi. 2019. "The Value of a Statistical Life." Working paper. Available at http://ssrn.com/abstract_id=3379967. Forthcoming in *Oxford Research Encyclopedia of Economics and Finance.*

Krupnick, Alan. 2007. "Mortality-risk Valuation and Age: Stated Preference Evidence." *Review of Environmental Economics and Policy* 1: 261-82.

Masur, Jonathan, and Eric Posner. April 20, 2020. "Cost-Benefit Analysis Supports Continuing the National Shutdown." *The Regulatory Review.*

OECD. 2012. *Mortality Risk Valuation in Environment, Health and Transport Policies*. Paris: OECD Publishing.

Office of Management and Budget. 2017. "2017 Report to Congress on the Benefits and Costs of Federal Regulations and Agency Compliance with the Unfunded Mandates Reform Act." Available at https://www.whitehouse.gov/omb/information-regulatory-affairs/reports/#ORC.

Parfit, Derek. 2000. "Equality or Priority?" In Matthew Clayton and Andrew Williams, eds., *The Ideal of Equality*, pp. 81-125. Houndmills: Palgrave.

Pindyck, Robert S. 2020. "Covid-19 and the Welfare Effects of Reducing Contagion." NBER Working Paper 27121. Available at https://www.nber.org/papers/w27121.pdf.

Pretnar, Nick. 2020. "The Intergenerational Welfare Implications of Disease Contagion." Working paper. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3624857.

Robinson, Lisa A. 2007. "How U.S. Government Agencies Value Mortality Risk Reduction." *Review of Environmental Economics and Policy* 1: 283-99.

Robinson, Lisa A., James K. Hammitt, and Richard J. Zeckhauser. 2016. "Attention to Distribution in U.S. Regulatory Analyses." *Review of Environmental Economics and Policy* 10: 308-28.

Robinson, Lisa A., Ryan Sullivan, and Jason F. Shogren. 2020. "Do the Benefits of COVID-19 Policies Exceed the Costs? Exploring Uncertainties in the Age-VSL Relationship." Working paper. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3626710.

Scherbina, Anna. 2020. "Determining the Optimal Duration of the Covid-19 Suppression Policy: A Cost-Benefit Analysis." Working paper. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3562053.

Sunstein, Cass R. March 26, 2020. "This Time the Numbers Show We Can't Be Too Careful." *Bloomberg Opinion*.

Thurnstrom, Linda, et al. 2020. "The Benefits and Costs of Using Social Distancing to Flatten the Curve for Covid-19." Working paper. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3561934.

Ugarov, Alexander. 2020. "Inclusive Costs of NPI Measures for Covid-19 Pandemic." Working paper. Available at https://www.medrxiv.org/content/10.1101/2020.03.26.20044552v1.

U.S. Department of Health and Human Services. 2016. *Guidelines for Regulatory Impact Analysis*. Available at https://aspe.hhs.gov/pdf-report/guidelines-regulatory-impact-analysis.

U.S. Department of Transportation. 2016. "Revised Departmental Guidance on Valuation of a Statistical Life in Economic Analysis." Available at https://www.transportation.gov/regulations/economic-values-used-in-analysis.

U.S. Environmental Protection Agency. 2010. (Updated 2014). *Guidelines for Preparing Economic Analyses*. Available at https://www.epa.gov/environmental-economics/guidelines-preparing-economic-analyses.

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

Viscusi, W. Kip, and Joseph E. Aldy. 2003. "The Value of a Statistical Life: A Critical Review of Market Estimates throughout the World." *Journal of Risk and Uncertainty* 27: 5-76.

Wilson, Linus. 2020. "Estimating the Life Expectancy and Value of Statistical Life (VSL) Losses from Covid-19 Infections in the United States." Working paper. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3580414.

Appendix

A.   *Simulated Population*

        The population is divided into 35 cohorts: seven age groups (individuals age 20, 30, 40, 50, 60, 70, and 80), each divided into five income quintiles ("Low," "Moderate," "Middle," "High" and "Top"). For the Middle income quintile, the survival probability in each year of life (the probability of surviving to the end of that year, conditional on being alive at the beginning) is set equal to the survival probability at each age in the most recent available U.S. population survival curve.[59] Survival probabilities for the other groups are adjusted so as to roughly match the estimates of life expectancy for different income groups in Chetty et al. (2016).[60] These survival probabilities for the five quintiles are then reduced to account for COVID-19 risk; see below, Appendix C.

        Incomes by quintiles are based on the most recent available official data on average after-tax-and-transfer household annual income. By quintile, these incomes are: $35,000; $48,000; $65,900; $91,700; and $214,900. The incomes were divided by the square root of average household size (2.54) to arrive at estimates of after-tax-and-transfer individual income: $21,961; $30,118; $41,349; $57,358; $134,840.[61]

        Data about the age distribution of income was used to estimate a time profile of income.[62] I estimated time factors for each year of life, and multiplied the quintile incomes above by the time factors to arrive at the income in that quintile in that year of life. The time factors are as follows (rounded to no decimal places): Ages 0 to 24: 38%. Ages 25 to 29: 84%. Ages 30 to

---

[59] U.S. Department of Health and Human Services, Centers for Disease Control and Prevention, *National Vital Statistics Reports* (vol. 68, No. 7, June 24, 2019), Table 1 (U.S. life table for 2017).

[60] The mortality risk at each age is 1 minus the survival probability. For income quintiles other than Middle income, mortality risks were taken from the U.S. survival curve, and then adjusted by a multiplicative factor in each year. The adjusted survival probabilities are, then, 1 minus the adjusted mortality risks. The multiplicative factors for the Low, Moderate, High and Top quintiles were, respectively 1.5, 1.2, 0.9, and 0.75. These multiplicative factors were chosen so that the ratio between life expectancy at age 40 for individuals in that quintile, and life expectancy at age 40 for 50th percentile individuals, was approximately equal to the ratio as estimated by Chetty et al. (2016).

[61] The most recent data about after-tax-and-transfer household income (for 2016) was taken from: Congressional Budget Office, *Projected Changes in the Distribution of Household Income*, *2016 to 2021* (December 2019; available at www.cbo.gov/publication/55941). See *id.* Appendix B, Table B-1. The recommendation to estimate individual income from household income by dividing by the square root of household size is taken from *id.*, Appendix A. The estimate of average household size at 2.54 was taken from the 2017 data (to match the survival curve) provided in U.S. Census Bureau, Current Population Survey, "Historical Households Tables," Table HH6 ("Average Population Per Household and Family"), available at https://www.census.gov/data/tables/time-series/demo/families/households.html.

[62] Mean income for various age bands (for 2016) was taken from U.S. Census Bureau, Current Population Survey, "Tables for Personal Income," Table PINC-01 ("Selected Characteristics of People 15 Years and Over, by Total Money Income, Work Experience, Race, Hispanic Origin, and Sex"), available at https://www.census.gov/data/tables/time-series/demo/income-poverty/cps-pinc/pinc-01.2016.html.
 The age bands were as follows: 15 to 24; 25 to 29; 30 to 34; 35 to 39; 40 to 44; 45 to 49; 50 to 54; 55 to 59; 60 to 64; 65 to 69; 70 to 74; 75 and over. The income in each band was divided by mean income for all individuals 15 and over ($46,550 as per this table) to arrive at the time factor for each year of life within the age band. The time factor for individuals under 15 was set equal to that for individuals in the age band 15 to 24.

34: 101%. Ages 35 to 39: 121%. Ages 40 to 44: 129%. Ages 45 to 49: 130%. Ages 50 to 54: 131%. Ages 55 to 59: 124%. Ages 60 to 64: 114%. Ages 65 to 69: 97%. Ages 70 to 74: 90%. Ages 75 and over: 67%.

For example, the lifetime income profile for the Low quintile is such that an individual at age 62 will receive income equaling 114% of the quintile income of $21,961. The lifetime income profile for the Top income quintile is such that an individual at age 38 will receive an income equaling 121% of the quintile income of $134,840. And so forth.

The percentages of the current population in the seven age groups were estimated to be as follows (rounding to no decimal places): age 20: 18%; age 30: 18%; age 40: 16%; age 50: 18%; age 60: 16%; age 70: 10%; age 80: 5%. These percentages were taken from estimates of the current age distribution of the U.S. population.[63]

B.   *Utility Function; Utilitarian and Prioritarian SWFs; VSL*

The formulas here for individual expected utility as a function of age, lifetime income profile, and lifetime profile of survival probabilities, and the formulas for the utilitarian and prioritarian SWFs and VSL, are those given in Adler, Ferranna, Hammitt and Treich (2019), except for not including a utility discount factor and changes in notation. This is in turn one version of the general model of utility and VSL presented in Part II. In this version, utility is temporally additive; non-income attributes are ignored; income is consumed when received (myopic consumption); and individuals have homogeneous preferences, captured in a common utility function.

There is a fixed population of N individuals. Each individual's life is divided into periods numbered 1, 2, ..., T, with T the maximum number of periods that any individual lives. An individual either dies immediately after the beginning of the period or does not die and survives, at least, until the end of the period.

The number of the current period, for individual i, is A_i. In the status quo, denoted with the subscript B ("baseline"), an individual i has a lifetime profile of survival probabilities: $(p_{i,B}^{A_i},...,p_{i,B}^{T})$. An individual also has a status quo lifetime income profile: $(y_{i,B}^{1},...,y_{i,B}^{A_i-1},y_{i,B}^{A_i},...,y_{i,B}^{T})$. A given policy P, denoted with the subscript P, endows each individual with a new lifetime profile of survival probabilities $(p_{i,P}^{A_i},...,p_{i,P}^{T})$ and income profile $(y_{i,P}^{1},...,y_{i,P}^{A_i-1},y_{i,P}^{A_i},...,y_{i,P}^{T})$.[64]

---

[63] U.S. Census Bureau, Current Population Survey, "Age and Sex Tables; Age and Sex Composition in the United States," Table 1 (2017) ("Population by Age and Sex"), available at https://www.census.gov/topics/population/age-and-sex/data/tables.2017.html. Each age group in the simulation was assigned the population percentage of the age bracket including that age from the table, and these percentages were then scaled up proportionally to sum to 100%.

[64] Policies don't change individuals' past incomes, so $y_{i,B}^{t} = y_{i,P}^{t}$ for $t < A_i$.

In what follows, the B and P subscripts will be removed if the same generic formula applies both to individuals' baseline profiles of survival probabilities and income amounts, and to individuals' policy profiles of survival probabilities and income amounts. $(\mathbf{p}_i, \mathbf{y}_i)$ is shorthand for an individual profile: $(\mathbf{p}_i, \mathbf{y}_i) = ((p_i^{A_i}, ..., p_i^T), (y_i^1, ..., y_i^T))$. $(\mathbf{p}_{i,B}, \mathbf{y}_{i,B})$ and $(\mathbf{p}_{i,P}, \mathbf{y}_{i,P})$ are, respectively, individual i's baseline profile and her profile with policy P.

The survival probabilities, denoted with p, are conditional probabilities: p_i^t is the probability that individual i survives to the end of period t, given that she is alive at the beginning. Related probabilities can be derived from these. Let $\pi_i(t; A_i)$, $t \geq A_i$, denote i's probability of surviving to the end of period t, given that she is currently alive at the beginning of A_i. Then $\pi_i(t; A_i) = \prod_{s=A_i}^{t} p_i^s$. Let $\mu_i(t; A_i)$, $t \geq A_i - 1$, denote i's current probability of surviving to the end of period t and then dying during the next period. (In other words, this is the current probability of living exactly t periods.) If $t = A_i - 1$, $\mu_i(t; A_i) = 1 - p_i^{A_i}$, If $t \geq A_i$, $\mu_i(t; A_i) = (1 - p_i^{t+1})\pi_i(t; A_i)$.

For a past period ($t < A_i$), y_i^t is the income that i earned in period t. For the current period or a future period ($t \geq A_i$), y_i^t is the income which i will earn in t if she survives to its end.

The common period utility function is denoted as $u(\cdot)$. Let $V_i^s(\mathbf{y}_i)$ denote the individual's realized lifetime utility if she lives exactly s periods. $V_i^s(\mathbf{y}_i) = \sum_{t=1}^{s} u(y_i^s)$. Individual i's current expected utility, denoted as $U_i(\cdot)$,[65] is as follows: $U_i(\mathbf{p}_i, \mathbf{y}_i) = (1 - p_i^{A_i})V_i^{A_i - 1}(\mathbf{y}_i) + \sum_{t=A_i}^{T} \mu_i(t; A_i)V_i^t(\mathbf{y}_i)$. It can also be shown that $U_i(\mathbf{p}_i, \mathbf{y}_i) = \sum_{t=1}^{A_i - 1} u(y_i^t) + \sum_{t=A_i}^{T} \pi_i(t; A_i)u(y_i^t)$.

The utilitarian social welfare function, denoted W^U, assigns a score to a policy or to the baseline as a function of individuals' profiles of survival probabilities and incomes, and ranks policies relative to each and the baseline as a function of these scores. The same is true for the ex post prioritarian social welfare function, denoted W^{EPP}.

$W^U((\mathbf{p}_1, \mathbf{y}_1), ..., (\mathbf{p}_N, \mathbf{y}_N)) = \sum_{i=1}^{N} U_i(\mathbf{p}_i, \mathbf{y}_i)$. W^{EPP} is, more precisely, a family of social welfare functions—each defined by some strictly increasing and strictly concave transformation function $g(\cdot)$. $W^{EPP}((\mathbf{p}_1, \mathbf{y}_1), ..., (\mathbf{p}_N, \mathbf{y}_N)) = \sum_{i=1}^{N} \left[(1 - p_i^{A_i})g(V_i^{A_i - 1}(\mathbf{y}_i)) + \sum_{t=A_i}^{T} \mu_i(t; A_i)g(V_i^t(\mathbf{y}_i)) \right]$

---

[65] The "i" subscripts on $U(\cdot)$ and $V(\cdot)$ could be dropped, since these are common utility functions.

In the simulation model, policies are combinations of changes to individuals' current survival probabilities and current incomes, relative to baseline. The formulas here for VSL and VSL-CBA are for this specific case. VSL_i, in this case, is the marginal rate of substitution between individual i's current survival probability and current income. $VSL_i = \dfrac{\partial U_i / \partial p_i^{A_i}}{\partial U_i / \partial y_i^{A_i}} \Big|_{(\mathbf{p}_{i,B};\mathbf{y}_{i,B})}$

. Using the formula for expected utility above, we have that $VSL_i = \dfrac{\sum\limits_{t=A_i}^{T} \dfrac{\pi_{i,B}(t;A_i)}{p_{i,B}^{A_i}} u(y_{i,B}^{t})}{p_{i,B}^{A_i} u'(y_{i,B}^{A_i})}$, with

$u'(\cdot)$ the first derivative.

Let $(\mathbf{p}_i,\mathbf{y}_i)$ differ from i's baseline profile only with respect to her current survival probability and/or income amount, if at all. $\Delta p_i^{A_i} = p_i^{A_i} - p_{i,B}^{A_i}$. $\Delta y_i^{A_i} = y_i^{A_i} - y_{i,B}^{A_i}$. Then the individual's VSL-based monetary valuation of this profile, MV_i, is as follows:
$MV_i(\mathbf{p}_i,\mathbf{y}_i) = \Delta p_i^{A_i} VSL_i + \Delta y_i^{A_i}$. VSL-CBA—like utilitarianism and prioritarianism—assigns a score to a policy or to the baseline as a function of individuals' profiles of survival probabilities and incomes, and ranks policies relative to each other and the baseline as a function of these scores. This score is the sum of MV amounts. $W^{VSL-CBA}((\mathbf{p}_1,\mathbf{y}_1),...,(\mathbf{p}_N,\mathbf{y}_N)) = \sum\limits_{i=1}^{N} MV_i(\mathbf{p}_i,\mathbf{y}_i)$.

Note that $MV_i(\mathbf{p}_{i,B},\mathbf{y}_{i,B}) = 0$ and so the score assigned by VSL-CBA to the baseline is 0. A policy is better than/worse than/equally good as baseline iff its score is positive/negative/zero.

In the simulation model, the periodization is annual and the period utility function is logarithmic. $u(y) = \log(y) - \log(y^{zero})$, with $y^{zero} = \$1000$. For an explanation of why a choice of y^{zero} is needed to specify the utility function, see Adler (2019, pp. 292-94); Adler (2017, pp. 71-72). y^{zero} can be thought of as subsistence income. Extending a life by one period with income y^{zero} leaves lifetime utility unchanged. The level of $1000 was chosen based roughly on the World Bank extreme poverty level of $1.90/day. See Adler (2017, pp. 71-72).

The Atkinson prioritarian SWF, which has attractive axiomatic properties, uses a power function. $g(V) = \dfrac{1}{1-\gamma} V^{1-\gamma}$, $\gamma > 0$. See Adler (2019, pp. 154-58, 274-75). As γ increases, priority for the worse off increase. I choose $\gamma = 1.5$, which has the effect of generally nullifying a utilitarian preference to reduce the risks of those at higher incomes. See Table 4c.

Baseline incomes for individuals in the 35 cohorts are as stated in Appendix A above; baseline survival probabilities are adjusted for the COVID-19 risk, as explained in Appendix C.

VSL^{avg}-CBA, as explained in the main text, uses the same formula as VSL-CBA except for using the population-average VSL to compute individuals' MV amounts.

Let L_i denote individual i's life expectancy with a given vector of survival probabilities.

$$L_i(\mathbf{p}_i) = (1 - p_i^{A_i})(A_i - 1) + \sum_{t=A_i}^{T} \mu_i(t; A_i)t = (A_i - 1) + \sum_{t=A_i}^{T} \pi_i(t; A_i).$$ Let $LE_i^{A_i}$ or, for short, LE_i denote

the difference between i's baseline life expectancy if she survives the current period and her

baseline life expectancy if she dies during the period. $LE_i = \sum_{t=A_i}^{T} \frac{\mu_{i,B}(t; A_i)}{p_{i,B}^{A_i}} t - (A_i - 1)$. As above,

let $(\mathbf{p}_i, \mathbf{y}_i)$ differ from i's baseline profile only with respect to her current survival probability

and/or income amount, if at all. $\Delta p_i^{A_i} = p_i^{A_i} - p_{i,B}^{A_i}$. $\Delta y_i^{A_i} = y_i^{A_i} - y_{i,B}^{A_i}$. Then the difference in i's

life expectancy between policy and baseline, $L_i(\mathbf{p}_i) - L_i(\mathbf{p}_{i,B}) = \Delta p_i^{A_i} LE_i$. $VSLY_i^{A_i}$ or, for short,

$VSLY_i$ equals VSL_i/LE_i.

Let $VSLY^{avg}$ be the population average VSLY. Then $VSLY^{avg}$-CBA sums monetary
valuations defined in terms of this population average.

$$MV_i^{VSLY^{avg}}(\mathbf{p}_i, \mathbf{y}_i) = \Delta p_i^{A_i}(LE_i)(VSLY^{avg}) + \Delta y_i^{A_i}$$

C.   *Social Distancing Policy*

Let $o_{i,B}$ denote individual i's baseline (no-governmental intervention) risk of dying during
the current year from COVID-19. $o_{i,B}$ for the various cohorts is given in Table 2 (based upon the
IFRs from Ferguson et al. (2020) and the estimate that, without intervention, 81% of the
population would be infected). $o_{i,P}$ is the individual's risk of dying during the current year from
COVID-19 after policy intervention P.

Let p_i^* denote individual i's current-year survival probability, but for the pandemic.
(These survival probabilities, for the various cohorts, are based upon the U.S. survival curve as
described in Appendix A, first paragraph.). Individual i's baseline current-year survival
probability with the COVID-19 risk, $p_{i,B}^{A_i}$, is set as follows: $p_{i,B}^{A_i} = p_i^*(1 - o_{i,B})$. Intuitively, i
survives the current year only if she is not killed by non-COVID-19 causes and is not killed by
COVID-19. Assuming these probabilities are independent—to be sure, a simplification—we
have the formula here. Similarly, $p_{i,P}^{A_i} = p_i^*(1 - o_{i,P})$.

I model COVID-19 reduction as a concave function of GDP reduction. If GDP is reduced
by π%, $o_{i,P} = o_{i,B}(1 - e\sqrt{\pi/80})$. In the base case ($e = 1$), 80% reduction wholly eliminates the
risk. In the optimistic case ($e = \sqrt{2}$), a 40% reduction wholly eliminates the risk. In the
pessimistic case ($e = 0.5$), an 80% reduction eliminates half the risk. In the convex case,

$$o_{i,P} = o_{i,B}(1 - (\frac{\pi}{80})^2).$$

On the cost side, regressive incidence is as follows. After a $\pi\%$ reduction in GDP: (1) if individual i is in the Low, Moderate, or Middle income quintiles, $y_{i,P}^{A_i} = y_{i,B}^{A_i} - (\frac{\pi}{80})(y_{i,B}^{A_i} - 1000)$; and (2) if individual i is in the High or Top income quintiles, $y_{i,P}^{A_i} = y_{i,B}^{A_i}(1 - (\frac{\pi}{80})q)$, with q chosen based upon cohort baseline incomes so that the reduction in total income is $\pi\%$. In the model here, $q = 0.7187$.

Proportional incidence is straightforward: After a $\pi\%$ reduction in GDP, $y_{i,P}^{A_i} = y_{i,B}^{A_i}(1 - (\frac{\pi}{100}))$ for individuals in all cohorts.

D. *Equivalent and Compensating Variations*

Using the utility model set forth in Appendix B., we can define equivalent variations (EV$_i$) and compensating variations (CV$_i$). $EV_i(\mathbf{p}_i, \mathbf{y}_i) = \Delta y^*$ such that:
$$U_i((p_i^{A_i},...,p_i^T),(y_i^1,...,y_i^T)) = U_i((p_{i,B}^{A_i},...,p_{i,B}^T),(y_{i,B}^1,...,y_{i,B}^{A_i-1}, y_{i,B}^{A_i}+\Delta y^*,...,y_{i,B}^T)).$$ $CV_i(\mathbf{p}_i, \mathbf{y}_i) = \Delta y^+$
such that: $U_i((p_{i,B}^{A_i},...,p_{i,B}^T),(y_{i,B}^1,...,y_{i,B}^T)) = U_i((p_i^{A_i},...,p_i^T),(y_i^1,...,y_i^{A_i-1}, y_i^{A_i}-\Delta y^+,...,y_i^T))$.

Assume, more specifically, that $(\mathbf{p}_i, \mathbf{y}_i)$ differs from baseline only with respect to current survival probability and/or current income, if at all. As a short hand, let U_i denote $U_i(\mathbf{p}_i, \mathbf{y}_i)$ and $U_{i,B}$ denote $U_i(\mathbf{p}_{i,B}, \mathbf{y}_{i,B})$. Then (1) $U_{i,B} = \sum_{t=1}^{A_i-1} u(y_{i,B}^t) + \pi_{i,B}(A_i;A_i)u(y_{i,B}^{A_i}) + \sum_{t=A_i+1}^{T} \pi_{i,B}(t;A_i)u(y_{i,B}^t)$.
Also, the compensating variation Δy^+ is s.t. (2) $U_{i,B} =$
$$\sum_{t=1}^{A_i-1} u(y_{i,B}^t) + \pi_i(A_i;A_i)u(y_i^{A_i}-\Delta y^+) + \sum_{t=A_i+1}^{T} \pi_i(t;A_i)u(y_{i,B}^t)$$

(3) $U_i = \sum_{t=1}^{A_i-1} u(y_{i,B}^t) + \pi_i(A_i;A_i)u(y_i^{A_i}) + \sum_{t=A_i+1}^{T} \pi_i(t;A_i)u(y_{i,B}^t)$. Also, the equivalent variation Δy^* is

s.t. (4) $U_i = \sum_{t=1}^{A_i-1} u(y_{i,B}^t) + \pi_{i,B}(A_i;A_i)u(y_{i,B}^{A_i}+\Delta y^*) + \sum_{t=A_i+1}^{T} \pi_{i,B}(t;A_i)u(y_{i,B}^t)$.

Subtracting the fourth equation from the first, we have that $U_{i,B} - U_i = \pi_{i,B}(A_i;A_i)\left(u(y_{i,B}^{A_i}) - u(y_{i,B}^{A_i}+\Delta y^*)\right)$. Subtracting the third from the second, we have that $U_{i,B} - U_i = \pi_i(A_i;A_i)\left(u(y_i^{A_i}-\Delta y^+) - u(y_i^{A_i})\right)$. Observe also that $\pi_{i,B}(A_i;A_i) = p_{i,B}^{A_i}$ and that $\pi_i(A_i;A_i) = p_i^{A_i}$.

With the logarithmic specification of $u(\cdot)$, $u(y) = \log(y) - \log(1000)$, we can use these last equations to solve explicitly for Δy^* and Δy^+. $\Delta y^* = y_{i,B}^{A_i}\left(\exp\left(\dfrac{U_i - U_{i,B}}{p_{i,B}^{A_i}}\right) - 1\right)$.

$$\Delta y^+ = y_i^{A_i}\left(1 - \exp\left(\dfrac{U_{i,B} - U_i}{p_i^{A_i}}\right)\right).$$