Eastern Economic Journal
https://doi.org/10.1057/s41302-023-00256-z

ORIGINAL ARTICLE

Analysis of Hybrid Epidemiological-Economic Models of COVID-19 Mitigation Policies

Jessica Carrick-Hagenbarth[1] · **Eric Edlund**[1] · **Avanti Mukherjee**[1]

Abstract
Some studies have examined the tension between public health benefits and economic costs of interventions in the United States for the COVID-19 pandemic by merging epidemiological and economic models. We extract from these studies a data set consisting of lives saved and the associated costs for each intervention. Our analysis calculates cost–benefit ratios that allow the effectiveness of intervention types to be compared against each other and against the value of the statistical life and the value of the statistical life year. This work identifies patterns in the cost–benefit ratios that illuminate the strengths and weaknesses of different intervention strategies and thereby enhances the practical application of complex theoretical analyses.

Keywords COVID-19 · Health policy interventions · Value of a statistical life · Models

JEL Classification I18

Introduction

The COVID-19 pandemic initially caused policymakers and the public to question whether the costs of disease mitigation strategies in terms of foregone national income outweighed the benefits of intervention in terms of lives saved (Hilsenrath and Armour 2020; Qin 2020). This question was complicated by a large degree of uncertainty regarding disease parameters, the health effects on individuals, and the consequences of restricting parts of the economy. Economists were invited to conduct research bridging the disciplines of economics and epidemiology (Avery et al. 2020; Murray 2020), but warned against naively adopting epidemiological methods (Murray 2020). From January 2020 to August 2021, around 15% of economics working papers in some of the major series, such as the National Bureau of

✉ Jessica Carrick-Hagenbarth
jessica.carrickhage@cortland.edu

[1] SUNY Cortland, P.O. Box 2000, Cortland, NY 13045, USA

Published online: 17 October 2023

Economic Research or Centre for the European Policy Research, addressed COVID-related themes (Bürgi and Wohlrabe 2022). While the discussion has become increasingly nuanced, the tradeoffs between saving lives and saving the economy remain a relevant public policy question and will likely continue to be so with new variants of COVID-19 and possibly altogether new pandemics.

Some economists sought to inform these discussions by searching for optimal policies through analysis of the tradeoffs between lives saved and economic losses (Bairoliya and İmrohoroğlu 2020; Favero et al. 2020; Glover et al. 2023; Gollier 2020; Hall et al. 2020; Kaplan et al. 2020; Rampini 2020; Thunström et al. 2020; Acemoglu et al. 2021; Alvarez et al. 2021; Brotherhood et al. 2021; Eichenbaum et al. 2021; Farboodi et al. 2021; Flaschel et al. 2021; Jones et al. 2021). A common approach of these studies was to combine economic and epidemiological models, which we refer to as *hybrid models*, such that optimal strategies could be identified. Many of these studies are not particularly accessible to an audience outside of specific subfields of economics, even for those with specialized mathematical skills, on account of the highly theoretical nature of the models and discipline-specific language. Furthermore, the range of intervention methods examined in these studies is large and, where there is overlap between studies, they do not always agree.

A common approach in the hybrid modeling literature for assessing the success of various intervention strategies is to compare scenarios with and without government intervention (Bairoliya and İmrohoroğlu 2020; Rampini 2020; Thunström et al. 2020; Alvarez et al. 2021; Brotherhood et al. 2021; Eichenbaum et al. 2021; Farboodi et al. 2021). This literature finds a broad range of deaths and economic losses across scenarios, making it difficult to discern patterns. In this article, we ask: Did the analyses that emerged from combining economic and epidemiological models present consistent outcomes or trends that could be useful for informing policy? We answer this question with two main contributions. The first is a literature review of US-relevant hybrid modeling analyses from which we can extract lives saved and economic losses. The second contribution is a new metric for interpreting the results of such models—a cost–benefit ratio for each intervention analysis that allows for a direct comparison of the economic cost of policy interventions to the value of lives saved, thereby drawing meaningful patterns for discussion. We argue that the most effective interpretation of this data requires that the outcomes with intervention be compared to the outcomes without intervention (the baseline, i.e., a "do-nothing" or "laissez-faire" strategy). By subtracting the deaths and economic losses of a model's specific policy scenario from those for the model's baseline, we calculate the economic cost of intervention and the lives saved by intervention. We construct a ratio of these measures (a cost–benefit ratio) that allows scenarios to be compared on equal footing across the studies and intervention types.

Our paper contributes to the field by interpreting salient points of a highly specialized subject for a wider audience including policy advocates and policymakers. By doing so, we make an original contribution to the field by providing a new metric, a cost–benefit ratio, for interpreting the results of hybrid studies and allowing a direct comparison of economic cost of policy interventions to the values of lives lost. The use of such a cost–benefit ratio appears to have been hitherto overlooked in this literature. The simplicity of this method of interpretation is a virtue that serves the dual purposes of providing a mechanism for commensurate comparison of the costs and benefits of intervention strategies while also increasing the accessibility of these studies to a broader audience.

Central to these hybrid modeling studies is a comparison of lives and economic losses. A quantitative analysis of the tradeoffs of different interventions requires a method of valuing life. As is common in this literature, we reference two widely used standards to assign a value to lives saved: the value of statistical life (VSL), which assigns a universal dollar value for each statistical life, and the value of a statistical life year (VSLY), which assigns a dollar value to every year of life saved or lost and has been used to account for age-dependent mortality variations (Schelling 1968; Viscusi 2018). The VSL is the larger of these values, roughly around $10 million per life (Viscusi 2018) compared to about $3 million for the VSLY in the context of the COVID-19 pandemic (Hall et al. 2020; Alvarez et al. 2021; Farboodi et al. 2021). Our cost–benefit ratio assigns a value to each modeling data point in units of dollars per life, which allows for a direct comparison to the VSL and VSLY. There are many policy scenarios where the cost–benefit ratio is below the VSLY, meaning that the costs incurred to save a life are below the measures of statistical life used by the studies. Our analysis also demonstrates that social distancing mandates, widespread testing, and age/health targeted policies have very low cost–benefit ratios compared to stay-at-home (or lockdown) policies.

Our analysis does not assess the validity of the specific epidemiological and macroeconomic models or the precision of point estimates. Instead, by developing and employing cost–benefit ratios that can be applied across studies, we demonstrate which intervention strategies appear to be the most effective at saving lives with the least economic costs. An important point in the larger discussion of responses to crises is the disproportionate effects experienced by the income-poor, people of color, and women. While assessing these impacts is outside the scope of our paper, any discussion of optimal strategies has significant bearing for those who have borne the greater burdens in terms of both lives lost and repercussions such as lost jobs and/or greater work burden.

The rest of the paper is organized as follows. The following section provides an overview of the hybrid modeling, explaining the basic epidemiological model, the use of the macroeconomic models, an introduction to and discussion of critiques of the value of a statistical life, and how all of this is combined via a valuation of lives saved versus costs incurred. The third section presents our data sources, with analysis methods and a presentation of our cost–benefit ratio in the fourth section. A broader discussion of these results with possible lessons for policy-making is presented in the fifth section, followed by concluding remarks.

Overview of the Epidemiological-Economic Modeling

Economists have modeled the potential economic costs of COVID-19 and of policy interventions on the economy by combining epidemiological models with macroeconomic models (Bairoliya and İmrohoroğlu 2020; Favero et al. 2020; Gollier

2020; Hall et al. 2020; Kaplan et al. 2020; Rampini 2020; Thunström et al. 2020; Acemoglu et al. 2021; Alvarez et al. 2021; Brotherhood et al. 2021; Eichenbaum et al. 2021; Farboodi et al. 2021; Flaschel et al. 2021; Jones et al. 2021; Glover et al. 2023). These models, which we call *hybrid models*, allow for the estimation of deaths and economic losses of government interventions aimed at disease mitigation. Epidemiological models are useful for studying the spread of a disease, the associated mortality, and exploring how we can exert control over these outcomes through behavioral change or various policy interventions. This type of modeling is an important tool for designing public policies aimed at disease mitigation. Macroeconomic modeling can be used to assess the economic impact of different disease mitigation policies, thereby allowing the comparison of the costs of various possible policies with their disease mitigation efficacy. We discuss some of the main features of such modeling below, including how they have been combined, to help provide insight into the strengths and weaknesses of different policy options.

Epidemiological Models

A class of epidemiological models, called SIR (Susceptible/Infectious/Recovered) models (Kermack and McKendrick 1927), and their contemporary extensions have played an important role in modeling the COVID-19 pandemic (Avery et al. 2020). Epidemiologists and others have estimated the number of lives lost to COVID-19 under a wide range of scenarios, including those without government intervention where the disease vanishes only by an unmitigated progression to herd immunity. In the simplest SIR models, the population moves through three states, Susceptible → Infectious → Recovered. More sophisticated models, in an effort to achieve a more accurate representation of the propagation of the disease and the extent of its impacts, may subdivide each of these states by age or location, include additional illness categories such as critical or asymptomatic, or account for high-density "super-seeder" events (Avery et al. 2020). If individuals cannot be reinfected, the recovered population increases as the susceptible population decreases. Herd immunity is reached when a sufficiently large fraction of the population attains immunity, either through exposure or vaccination, such that infections begin to decrease.[1] While individuals can suffer multiple infections of COVID-19, such cases are relatively rare in the short-run and it is far more common that individuals develop significant natural immunity that persists for the time-period of up to approximately 20 months (Nordström et al. 2022), indicating that over the timescale of one year the concepts of immunity and herd immunity are approximately true.

These models depend on several important parameters, including the case fatality ratio, the basic reproduction number (R_0), and the effective reproduction

[1] At the beginning of the COVID-19 pandemic many epidemiological models modeled a single variant and assumed the end of the pandemic would come through reaching herd immunity. However, there are significant challenges to reaching herd immunity in the case of COVID-19 because of vaccine hesitancy, the lack of widespread global access to COVID-19 vaccines, breakthrough infections, short-term immunity, and new variants (Aschwanden 2021; D'Souza and Dowdy 2021). More recently, it has become clear that COVID-19 will be endemic, but perhaps less virulent, and the challenge will be learning how to live with it, rather than eradicating it.

number (R_t), among others. The range of values for these critical parameters underlies the wide range of possible deaths calculated by the models. The case fatality ratio (CFR), also described as the case fatality rate, is a measure of the number of deaths among all persons with the disease. The CFR can be especially uncertain in the presence of asymptomatic cases. The basic reproduction number is an estimate of disease contagion. It is the average number of people infected by one infected person at the beginning of an epidemic when approximately the whole population is susceptible, people have not yet modified their behavior, and no mitigation measures have been put in place (Shaw and Kennedy 2021). The basic reproduction number determines the maximum scale of the pandemic in the case of no coordinated mitigation efforts or behavioral change. The effective reproduction number reflects disease contagion given changes over time due to a declining susceptible population, behavior modification, or the implementation of mitigation measures (Shaw and Kennedy 2021).

For a disease to spread, the effective reproduction number must be greater than one. In the absence of government intervention and behavioral change, the overall infection rate and size of the recovered population increase with R_0, whereas the time to peak infection and remaining susceptible population decrease with R_0. Severe diseases with large values of R_0 (roughly for values greater than 2), like COVID-19, carry the risk of overwhelming the health care system as infection rates exceed its ability to manage cases. Most of the models in the literature surveyed here used an R_0 of around 2.4 (Greenstone and Nigam 2020; Rampini 2020; Thunström et al. 2020; Alvarez et al. 2021; Brotherhood et al. 2021), based on the influential and early COVID-19 epidemiological model created by Ferguson et al. (2020), although some also used values between 1.23 and 3.1 (Bairoliya and İmrohoroğlu 2020; Farboodi et al. 2021). Reproduction numbers greater than about 2 result in a rapid spread of COVID through the population, typically reaching herd immunity in less than one year. We found this result to be robust to alternative parameter specifications in SIR models. Reproduction numbers close to one, which can be effected through intervention measures, allow the pandemic to be slowed sufficiently for other policy interventions, such as the development of a vaccine, to be implemented.

Macroeconomic Models

The macroeconomic models that are coupled to epidemiological models predict potential economic losses from disease mitigation strategies. Such models estimate the costs of policy interventions that combat the progression of the disease by restricting the economy, such as the closure of nonessential businesses, schools, and travel restrictions that necessarily incur economic costs. This should be differentiated from direct costs that may be related to a specific type of intervention, such as the cost of medical testing kits. The models examined here have generally limited themselves to including the cost of economic restrictions and have not included the direct costs of the intervention. There was large variation in the economic impacts felt by individuals and industries. Most of the models examined in this study treated the population as an aggregate, except for some cases which allowed for income, age, and health-related impacts. Additionally, most of the models do not consider how the costs of the interventions could be offset by economic stimulus.

The studies included in our paper present several different types of macroeconomic models and estimations of economic loss. One group of models build from the microeconomic foundations of a representative agent optimizing a discounted lifetime utility function (Bairoliya and İmrohoroğlu 2020; Brotherhood et al. 2021; Eichenbaum et al. 2021; Farboodi et al. 2021). That is, an individual whose goal is to maximize the well-being they receive from their consumption and work time over their lifetime. Such utility functions assume that an agent values their future consumption and work time at a discounted rate in the present, i.e., they value present consumption more than future consumption. Representative agents weigh the increased probability of infection against utility derived from consumption and wages from work, often with the result that they both consume and work less or work from home at an assumed lower productivity. These types of models allow for the possibility of behavioral change as representative agents respond to high rates of infection. These models can also provide a representative firm that employs workers and sees a loss in output due to behavioral change, infection, death, or mitigation measures that can cause or require agents to work less (Bairoliya and İmrohoroğlu 2020; Brotherhood et al. 2021; Eichenbaum et al. 2021). This output loss is reported in many of the papers as a decline in GDP. We use their estimated GDP loss for all cases as close to the one-year mark as possible.

The second group of macroeconomic models employs a more diverse variety of methods to estimate economic losses. They vary from the use of an aggregate output function, wherein output losses occur when workers are unable to work (Rampini 2020), a social planner function minimizing the present value of the trade-off between economic losses and lives lost (Alvarez et al. 2021), a net present value function based on Goldman Sach's GDP projections (Thunström et al. 2020), and a utility function measuring lost utility due to decreased social interaction because of economic shutdown or worry (Farboodi et al. 2021).

In all these models, except Farboodi et al. (2021), economic losses are reported as output losses due to lost production or consumption. In contrast, the Farboodi et al. (2021) model economic losses as utility losses due to forgone social activity measured in utils and translated into an equivalent GDP loss using the VSL. We extrapolate the GDP value for the models at one year to make them comparable, with the exception of Alvarez et al. (2021), whose model was only run to 300 days. The variations in losses among models are discussed in greater detail in the section titled Economic Losses.

Hybrid Models

Hybrid models, the combination of epidemiological and macroeconomic models, vary along several lines meaningful to our discussion: behavioral change, hospital overloading (HO), and heterogeneous agents. Table 1 describes each of the papers included in this analysis by which of these elements are included in their model. There are certainly other important effects that could be included in these models.

Table 1 Model variations

Author	Behavioral change	Hospital overload	Heterogeneous agents	Values of lives saved	Value of GDP
Bairoliya and İmrohoroğlu	N	N	Y (age, income, health)	Y	Y
Rampini	N	N	Y (age)	Y	Y
Alvarez et al.	N	N	N	Y	Y
Ferguson et al.	N	N	N	Y	N
Thunstrom et al.	N	Y	N	Y	Y
Greenstone and Nigam	N	Y	N	Y	N
Farboodi et al.	Y	N	N	Y	Y
Brotherhood et al. 2021	Y	Y	Y (age)	Y	Y
Eichenbaum et al.	Y	Y	N	Y	Y

The table is organized first by behavioral change and then by the hospital overloading effect. Y is yes it does include this effect. N is no it does not include this effect.

One such effect that is known to have had an impact on public health is the delay in treatment of other medical procedures, either preventative health measures or critical health interventions, such as treatment of cardiac disease, for which the delay of treatment can have significant consequences. None of the studies surveyed included such effects. However, all hybrid modeling studies examined here include extensive discussions of their calibration/benchmarking methods for epidemiological and economic parameters. These methods typically involve matching the time-scale and severity of disease propagation in the SIR models to disease behavior described in the epidemiological literature.

The first model element, behavioral change, describes how/whether a model's agents respond to population-wide levels of infection or death. The underlying idea is that people will voluntarily change their consumption and work patterns when confronted with news of higher rates of deaths and infection. The epidemiological effect of such behavioral change results in reduced exposure to the disease, thereby reducing propagation and the infection rate. Some of the hybrid models examined in this analysis do not incorporate behavioral change so that the counterfactual experience of no government intervention results in zero or negligible economic loss. In reality, people tend to adapt their behavior when faced with a serious pandemic like COVID-19. For example, Farboodi, Jarosch, and Shimer found that during the first week of the pandemic (3/13/20–3/20/20), google data measuring social activity showed declines of 25–33% in the use of public transit hubs and in workplace, retail, and recreation activities, along with a corresponding increase in residential activity of 12%, all prior to the implementation of lockdowns (2021, p. 7). Juranek et al. studying labor markets in the Nordic countries during the first wave of the COVID-19 pandemic, found considerable unemployment increases due to people voluntarily reducing work activity to prevent infection even in the absence of economic shutdown (although they also found that stay-at-home policies significantly exacerbated unemployment) (2021). As such, models including behavioral change are more realistic than those that do not include behavioral change. The economic consequence of behavioral change is that when included in the no-intervention scenarios they will show substantial economic losses, thereby making the actual cost of intervention less severe.

A second line of variation between models is whether the epidemiological models include the possibility of hospital overload (HO) (see Table 1). When hospitals reach capacity and there are not enough ICU beds available for critical patients, the care required to treat the disease is not available, which results in a greater number of deaths. This effect is often expressed mathematically as an increased mortality rate conditional on a threshold that compares the infected population to the total number of ICU beds or can be expressed as a mortality rate that increases continuously with patient load.

The third line of variation in the models is whether they assume homogenous or heterogeneous agents. Several of the models in the given papers include agents which vary by age or health. Age is particularly important as the elderly have died at significantly higher rates with COVID-19. Heterogeneous agent models allow for age or health-targeted policies. When combined with behavioral change, hospital overloading and heterogeneous agents can produce a much richer range of economic outcomes.

Measuring the Value of a Statistical Life

The hybrid models reviewed here are engaged in a cost–benefit analysis of COVID-19 interventions aimed at saving lives at a potential economic loss. Such an analysis requires commensurate units and thus needs a monetary valuation of life. While considered a necessary condition for cost–benefit analysis, the monetization of the value of life is often controversial (Banzhaf 2014). Despite the challenges of accurately measuring the value of a statistical life and the controversies associated therewith, described subsequently, such measures appear throughout the hybrid modeling literature because they provide a simple method for comparing economic losses and deaths. It is important to note that this approach has developed significant credibility through its application by numerous federal agencies.

Prior to the emergence of the value of a statistical life approach, explained below, the main method for ascribing value to life was based on the human capital approach (Hood 2017), which uses labor market productivity accounting for the discounted present value of future earnings either to the individual or to others (Mishan 1971). This approach has been criticized in that it values people only in as much they contribute to gross national product and places no inherent value on life and little to no value on those without earnings, such as the elderly, children, and homemakers (Hood 2017).

Some economists found this approach deficient and created new valuations based on the concept of willingness to accept (WTA) or willingness to pay (WTP) some amount for a small increase or reduction in risk, which is more consistent with welfare economics in which individual preferences, utility maximization, and Pareto improvements are central (Schelling 1968; Mishan 1971; Jones-Lee 1976). This new approach became the value of a statistical life (VSL). Banzaf (2014) traces the concept of the value of a statistical life to 1948 when the US Air Force tasked the RAND Corporation with maximizing bombing damage subject to the constraints of a limited budget. Their initial analysis placed no weight on the lives of airplane crews and advised a strategy that would have met the bombing objective but with a high casualty rate. These calculations were criticized by the Air Force because they failed to value the lives of the crewmen (Banzaf 2014).

Schelling's 1968 paper, "The life you save may be your own," responded to these criticisms by proposing the concept of a statistical life based on measuring an individual's willingness to pay to prevent small increases in the risk of death that could then be aggregated over many individuals to derive the value of a statistical life. Schelling (1968) emphasized the importance of distinguishing an individual or identified life (an ex post known person who dies) from a statistical life (an unknown ex ante possible loss of life) because we tend to place far greater value on the former and the latter is the domain of policy that can be applied to a population. The VSL can be measured by using stated preferences or revealed preferences (Friedman 2020). Stated preferences are often based on surveys asking individuals about their willingness to pay to make small reductions in their risk of death or their willingness to accept increased wages for small increases in their chance of death on the job. Revealed preferences rely on the observation of labor market and consumption decisions of individuals and study individuals' willingness to accept or willingness to pay for small increases or decreases in risk. The approach of using labor market data to extract revealed preferences, from which the WTA and can be calculated, has become one of the primary means for determining the VSL in the United States (Viscusi 1992).

Given its ubiquity and acceptance as a standard for valuing lives, the VSL is a logical choice for evaluating the tradeoffs presented by the COVID-19 pandemic. However, considering that the COVID-19 pandemic disproportionately affects the elderly, many have argued that the more appropriate measure is years of life lost, or the value of a statistical life year (VSLY) (Bairoliya and İmrohoroğlu 2020; Colmer 2020; Alvarez et al. 2021; Farboodi et al. 2021). This measure adjusts the statistical value of life by applying an annual discount over the remaining years of life, which can then be summed over the population (Aldy and Viscusi 2007). Since COVID-19 mortality rates asymmetrically affect older people, the VSLY tends to decrease the population-averaged value of life when considering the deceased. There is ongoing debate about how exactly this accounting should be performed. Evidence indicates that people may assess the value of risk and life differently depending on age; yet, how these effects impact the calculation of the VSL and VSLY is uncertain (Evans and Smith 2006; Robinson et al. 2021). While Aldy and Viscusi (2007) find the relationship between age and WTP for risk reduction to be an inverted U, Evans and Smith (2006) find that for older adults in good health this relationship may have a positive correlation.

The WTA/WTP method to calculating VSL/VSLY and the resulting values have been the subject of substantial critique in the literature. Some economists have rejected the idea completely, arguing that human life is irreducible to monetary valuation (Broome 1985; Ackerman and Heinzerling 2004; Aldred 2022). Others have argued that it is the name—the value of a statistical life—that is the problem. For example, Cameron asserts that the name has led to widespread misinterpretation of the concept, wherein people assume it is the value we assign to a whole life rather than the willingness to pay for small reductions in risk applied to many different people (2010). Cameron and others argue that it would be better to not aggregate up to a human life (Cameron 2010; Simon et al. 2019). However, the practice has evolved such that regulatory agencies often use the VSL to value whole lives (Ackerman and Heinzerling 2001), just as the hybrid studies examined here did. Another critique is that both the WTP and the WTA approaches assume that people make informed choices (often in situations where it is difficult even for experts to determine the risk level), as well as failing to consider how differences in income may affect WTP and WTA (Ashenfelter 2006; Sunstein 2014; Friedman 2020). Furthermore, it is often assumed in revealed preference analysis using labor market data that higher risk is associated with higher wages. The approach of calculating the VSL through the WTA has also been criticized because it incorporates assumptions that do not necessarily hold; for example, that people have the ability to negotiate salaries to account for increased risk or that they can easily change jobs in the face of higher risk (Anderson 1993; Dorman 2009; Friedman 2020). Additionally, a VSL based on labor market data, which tends to represent risk of death from industrial accidents, may be considerably lower than people's risk aversion due to the possibility of death from disease (Cameron and DeShazo 2013).

There are several identified weaknesses of the stated preferences approach, which has often been derived through surveys regarding individuals' willingness to pay to reduce risk or through analyzing investments of time or money. Surveys are subject to selection bias, sensitivity to question framing, and the ability of people to interpret and/or extrapolate the very small probabilities used in these surveys (Friedman 2020). The VSL thus appears to be sensitive to specification (Ashenfelter 2006), which leads to substantial variation in estimates of the VSL. Banzhaf presents a meta-analysis of seven meta-analyses that provide mean VSLs ranging from $4.3 to $13.8 million (2022).

Finally, there are several criticisms specific to its use in the context of the COVID-19 crisis (Colmer 2020). First, the mortality risk of COVID increases with age and so a VSL based on the median person may underestimate the social benefits of policy intervention (Colmer 2020). Second, the elderly are often not represented in WTA calculations in proportion to their population as they do not reflect a large portion of the working population on which the VSL is generally calculated, and as such the value they assign to risk may be inaccurately measured (Colmer 2020).

Despite the varied and serious critiques of the VSL concept, it has developed significant credibility because it has been used broadly by various federal agencies, including the OMB, OSHA, FDA, EPA, and DOT (Viscusi 2018). More recently, Viscusi calculated the VSL to be between 9 and 11 million dollars (Viscusi 2018, p. 6; Viscusi 2021). The VSLY calculated using the COVID-19 pandemic mortality rates places a population-averaged value of a life in the range of 2.6 to 3.9 million dollars (Hall et al. 2020; Alvarez et al. 2021; Farboodi et al. 2021). All of the hybrid models cited here use the VSL/VSLY in their calculations; we follow this pattern and use the VSL and VSLY as reference values in our cost–benefit analysis.

Data and Methods

Throughout this and the following sections, we will use the following general terminology. We use "study" to mean an individual paper, "model" to mean the numerical algorithms used in a particular study to generate a set of outcomes, "intervention" or "intervention strategy" or just "strategy" to mean a policy focused on a single method of intervention to curb the spread of the pandemic, and "case" or "scenario" to refer to a particular instance of an intervention run under specific model parameters. This section describes the studies examined (i.e., the data), the categorical distinctions that define our interpretation of the data, and concludes with a definition of our method for interpreting these studies in terms of a cost–benefit ratio that allows for comparison across studies and across intervention strategies.

Selection of Data Sources

We analyzed only those papers that focused on the United States and from which we could extract monetary valuations of economic costs and lives saved around the one-year mark. At the beginning of the pandemic there was a great deal of research by economists on COVID-19, but since the publication timeline in economics tends to be long, the quickest way to release research was via working papers. Sources used to find papers were Google Scholar and the following working paper series: National Bureau of Economic Research, Centre of Economic and Policy Research Covid Economics, and Social Science Research Network. Studies were found using combinations of the search terms "macroeconomic" and "epidemiology" and "COVID" and "United States" and/or "pandemic" and/or "policy" and/or "intervention" and/or "SIR." Additionally, we followed any relevant references cited by the papers found via the above searches. Each of the seven studies examined here uses a hybrid epidemiological-macroeconomic model to explore tradeoffs between different policy decisions, augmented by two epidemiological models (one of which was highly cited in the literature), which together provide 27 data points. The list of studies examined can be found in "Appendix."

Selection of Intervention Categories

Our analysis uses a subset of all cases presented in the literature. This down-selection resulted from the requirements that the scenarios be (a) single intervention and (b) reasonably unique compared to other scenarios within the same study. Examples of the first condition are the elimination of scenarios that examined combined testing and age-targeted policies or others that included the effects of a vaccine. While all papers examined single intervention types, some papers examined the effects of multiple, simultaneous intervention strategies as well (Alvarez et al. 2021; Brotherhood et al. 2021; Eichenbaum et al. 2021). Because it is important to first understand the effectiveness of each intervention strategy on its own, we only examined cases with single interventions. We categorize intervention strategies into five main groups: stay-at-home policies, social distancing, testing, health-targeted policies, and age-targeted policies. This classification of intervention types presents the most common types of interventions examined in the hybrid modeling studies.

The second condition excluded cases that are nearly redundant with others in the same study, such as similar variations on moderate stay-at-home orders that do not have a significant difference in either deaths or economic losses. Where studies included more than one variation of a numerical model (with different effects included), we chose the more realistic model. For example, of the Eichenbaum et al. (2021) models with and without a hospital overloading effect, we have included only those with this effect. Table 2 presents the policy intervention types by paper.

Stay-at-home policies

We use the terminology of "stay-at-home" policies, as defined by the Centers for Disease Control and Prevention (CDC) (2021), to refer to the collection of other policies often interchangeably referred to as lockdowns, shelter-in-place, stay-at-home, and safer-at-home. The primary purpose of stay-at-home policies is to reduce disease spread by limiting activities outside the home that tend to bring people into closer and more frequent contact with each other. They may recommend or require that people not leave their homes except for daily outdoor exercise, grocery shopping, essential activities, and work, especially for essential workers. In terms of the mathematics of pandemic modeling, such policies have the effect of reducing the R_t.

We divide stay-at-home policies into three categories defined by a parameter $Q = T \times f$, where T is the duration of the stay-at-home policy measured in weeks and f is the average fraction of the population under such orders. We define *mild stay-at-home* policies by the range $0 \leq Q < 5$, *moderate stay-at-home* policies by $5 \leq Q < 10$, and *severe stay-at-home* policies by $Q \geq 10$. For example, a *mild* policy could be achieved by having 100% of the population stay at home for up to 4 weeks or 40% of the population at home for up to 10 weeks. Similarly, a *moderate* policy could be achieved by having 50% of the population stay at home for a period of time between 10 and 20 weeks.

The studies that evaluated the effects of stay-at-home policies were Alvarez et al. (2021), Bairoliya and Imrohorglu (2020), Brotherhood et al. (2021), Eichenbaum et al. (2021), and Rampini (2020). Stay-at-home policies were modeled in a variety of ways. Bairoliya and İmrohoroğlu (2020) present a model with heterogeneous agents that differ by age, income, and health status (good, fair, bad) and intervene with different stay-at-home policies. In their general case, they model a scenario where 50% of the population is unproductive for a period of 3 months and scale this to a period of one year, resulting in an effective stay-at-home policy affecting 12.5% of the population for one year (Bairoliya and İmrohoroğlu 2020, p. 2). Brotherhood et al. (2021) present a variety of stay-at-home policies involving combinations of 4-week and 26-week periods in combination with 25%, 75%, and age-targeted

Table 2 Policy intervention types by paper

Policy intervention type	Paper	Representative shape in figures
Stay-at-home	Alvarez et al. (2021), Bairoliya and Imrohoroglu (2020), Brotherhood et al. (2021), Eichenbaum et al. (2021), Ferguson et al. (2020), Greenstone and Nigam (2020), Rampini (2020)	Square
Testing	Brotherhood et al. (2021)	Star
Social distancing	Farboodi et al. (2021), Thunstrom et al. (2020)	Diamond
Age-targeted	Bairoliya and Imrohoroglu (2020), Brotherhood et al. (2021), Rampini (2020)	Triangle
Health-targeted	Bairoliya and Imrohoroglu (2020)	Cross

fractions.[2] Of the two Eichenbaum et al. (2021) cases included here, stay-at-home policies were uniformly applied across the population. These included policies that ramped the time spent at home up or down in response to the infection rate. The policy interventions lasted from around 100–120 weeks. Alvarez et al. (2021) also examined a ramped stay-at-home policy. Rampini (2020) evaluates three models of government lockdowns: early lifting, late lifting, and sequential lifting (age-targeted) of stay-at-home orders.

Social Distancing

Social distancing policies that targeted social and consumption-related activities typically aim to limit the density or the number of people in public and private spaces, including such things as social gatherings, public transport, gyms, tourist attractions, and schools (Thunström et al. 2020; Farboodi et al. 2021). We treat social distancing as distinct from stay-at-home policies and use social distancing to refer to mandate personal space boundaries and limits on gathering sizes. The two studies that evaluated social distancing policies are Farboodi et al. (2021) and Thunström et al. (2020).

Testing

Testing refers to policies that are primarily focused on large-scale and frequent testing to allow for the tracking of infections and to reduce propagation through quarantining the infected. It is through the mechanism of quarantine that testing earns its effectiveness at mitigating the spread of a disease, so any such quarantine policy is lumped under the umbrella of testing.

The models that examined testing were all conducted by Brotherhood et al. (2021), which considered the testing of individuals presenting symptoms of a fever, for which they included a fevered state in their epidemiological model. This study presented a variety of possible models, including some that include only testing and some that include testing and quarantine of those who test positive. We present two out of their six testing models, which examined various mixes of testing all ages, testing young and testing old, and testing plus quarantine.[3]

[2] "All ages shelter at home for an extra 75% of the time for 35 weeks" is a longer duration with lesser restrictions on mobility outside home compared to "all ages shelter at home for an extra 90% of the time for 26 weeks" which is a shorter duration but with greater restrictions on mobility outside home (Brotherhood et al. 2021).

[3] We restricted ourselves to these two models for the following reasons. First, the models found little effectiveness of just testing the old as they make up a small percentage of the population. Second, testing plus quarantine was more effective than solely testing, and it is commonsense that quarantines would be imposed on those testing positive. Third age-based testing and quarantine that only targeted the young had very similar outcomes in terms of deaths and costs to those targeting all ages.

Health-Targeted Policies

The intention behind health-targeted policies is to isolate and protect individuals who are at greater risk of death due to the disease. Bairoliya and İmrohoroğlu examined a policy where individuals were selectively forced to shelter based on their health status (2020). While this scenario also involved a condition on age, the primary objective was to identify individuals based on overall risk. Their condition was that individuals at high risk of any age category and elderly individuals at fair health (or worse) were forced to stay at home.

Age-Targeted Policies

Age-targeted policies here refer to policies that target groups based on their age alone. In principle, age-targeted policies could be either directed toward the young or the elderly. Some of the hybrid modeling examined a range of age-targeted policies, with some scenarios focusing on the young and others focusing on the elderly. While these studies tended to conclude that policies targeting the elderly were more effective, it is worth noting that it is not obvious which is better since transmission and mortality differ significantly across age groups. For example, some recent studies (not part of this analysis) that examined vaccine distribution strategies have shown that the optimal strategy (young first or elderly first) depends on the quantity of vaccines and vaccine effectiveness (Bubar et al. 2021; Matrajt et al. 2021).

Hybrid models that examined age-targeted policies focusing on the young (i.e., up to retirement age) are costly because the young tend to be the major consumers and producers and have comparatively low mortality rates due to COVID-19. In contrast, policies targeted at the elderly tend to be fairly effective at saving lives and have comparatively lower economic impacts because the elderly comprise a smaller fraction of the population, tend to consume and produce less (Lee et al. 2014), and have higher COVID-19 mortality rates. Given that policies targeting the young seem to be exceptionally costly and not very effective at saving lives, our analysis focused only on those age-targeted policies that sought to shelter the elderly via stay-at-home policies. For the remainder of this discussion, age-targeted will mean only those policies that are targeted specifically toward the elderly.

Baseline Scenarios

The baseline scenarios estimate deaths and the costs to the economy assuming no government intervention. Figure 1 presents all the baseline scenarios in the two-dimensional space of deaths and economic loss compared to a world without the pandemic. Table 1 lists each of these baseline models together with the model's main features. Except for Rampini (2020), all the papers we analyzed include at least one baseline model, which estimated the death rate and the economic loss in the case of no government intervention in the economy. In the case of the Rampini (2020) scenarios, we use the Greenstone and Nigam (2020) data for baseline deaths, which is similar to the widely cited Ferguson (2020) study. The error bars on the Rampini (2020) baseline

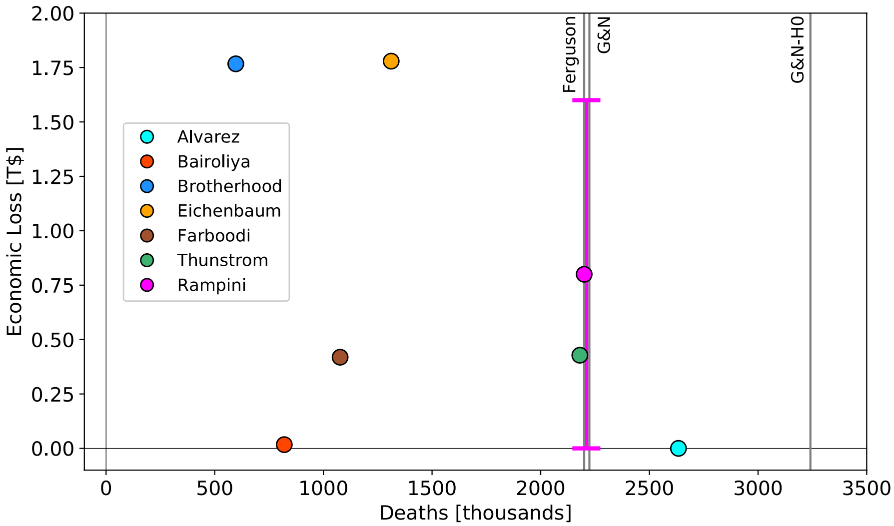

Fig. 1 Baseline (no-intervention) cases showing calculated deaths due to COVID-19 and, for the circles, the associated economic loss in trillions of US dollars. The vertical lines represent the calculations of the purely epidemiological models of Ferguson et al. and Greenstone and Nigam (G&N and G&N-HO), the latter including a hospital overloading (HO) effect in the model that increases the probability of death when ICU capacity is exceeded. Importantly, all of the cases with deaths greater than 2000 thousand do not include behavioral change, whereas all cases below that do include behavioral change or are scaled so as to produce an effect similar to behavioral change (as done by Bairoliya and İmrohoroğlu 2020).

represent the standard deviation of the economic losses in the other hybrid baseline cases.

Deaths

The baseline models cluster around either one million or two million deaths. The difference between these two clusters of deaths depends on two factors: whether the model allows for behavioral change and whether hospital capacity is reached (the hospital overloading (HO) effect, discussed in previously in the section Hybrid Models). The cluster of scenarios near 2 million deaths is those that do not incorporate behavioral change. In contrast, models that incorporate behavioral change allow agents to adjust their behavior in response to the number of infections or deaths in the population. Such behavioral change will decrease the predicted death rate as people react to higher infection numbers by decreasing their productive, consumptive, or social activities, all of which result in a lower R_t. Most of these models cluster around 1 million deaths. The Brotherhood et al. (2021) paper illustrates the effect of behavioral change with heterogeneity. The baseline model incorporated age-specific behavior that resulted in 1.184 million deaths without government intervention.

During a pandemic, the need for care in a hospital can easily outpace supply, thereby leading to an increased mortality rate. In the Eichenbaum et al. (2021) HO model, the mortality probability depends on the infected population beyond a threshold, which, once surpassed, creates an additional 426 thousand deaths. The Greenstone and Nigam (2020) baseline model that incorporates hospital overload but not behavioral change finds that deaths increase from 2.224 million without HO to 3.241 million with HO.

Economic Losses

The baseline models (no government intervention) show economic losses of anywhere between 17 million dollars (Bairoliya and İmrohoroğlu 2020) to over one trillion dollars (Brotherhood et al. 2021). This is presented in Figure 1. In the models that allow for behavioral change, individuals voluntarily reduce consumption, work, or leisure activities in response to higher infection rates to avoid risk, leading to decreases in income, utility, or workforce participation (Brotherhood et al. 2021; Eichenbaum et al. 2021; Farboodi et al. 2021). The error bars on the Rampini (2020) scenarios indicate that its baseline case was discussed in terms of deaths, referencing Ferguson et al. (2020), but did not discuss explicitly the economic losses in the baseline. However, the error bars also serve a more general purpose. There is, of course, uncertainty in all of these estimates, and the error bars on the Rampini (2020) results give a sense of scale of the uncertainty in the results. Importantly, the range of the economic costs of the interventions shown in Figure 2 is much larger than this uncertainty, meaning that the results are fairly robust with regard to such errors.

Time Frames

The time frames for the studies examined here span six months to six years. The purely epidemiological study of Greenstone and Nigam (2020) concluded their analysis at six months, but predicted that the pandemic had fully spread through the population by that point, reaching herd immunity, so that the number of deaths had effectively saturated. The models of Alvarez et al. (2021) were run for a period of 300 days, though it is important to note that the pandemic reached a steady-state well before this in all of their cases, so the number of deaths at one year could be extrapolated from this model with strong confidence. Ferguson et al. (2020) and Farboodi et al. (2021) ran their models for a period of 1 year. The remaining papers ran their models for lengths of time greater than one year: Brotherhood et al. (2021) concluded at 100 weeks, Rampini (2020) concluded at two years, Eichenbaum et al. (2021) concluded at 150 weeks, Thunström et al. (2020) concluded at five years, and Bairoliya and İmrohoroğlu (2020) concluded at six years. For those studies that provided data up to or beyond one year, we extract data close to the one-year mark. We do this for several reasons. The hybrid models were composed near the beginning of the pandemic when there was a great deal that was unknown both epidemiologically and economically. For example, there had not yet been the creation of a vaccine and discussions regarding federal stimulus were in their infancy. It was difficult

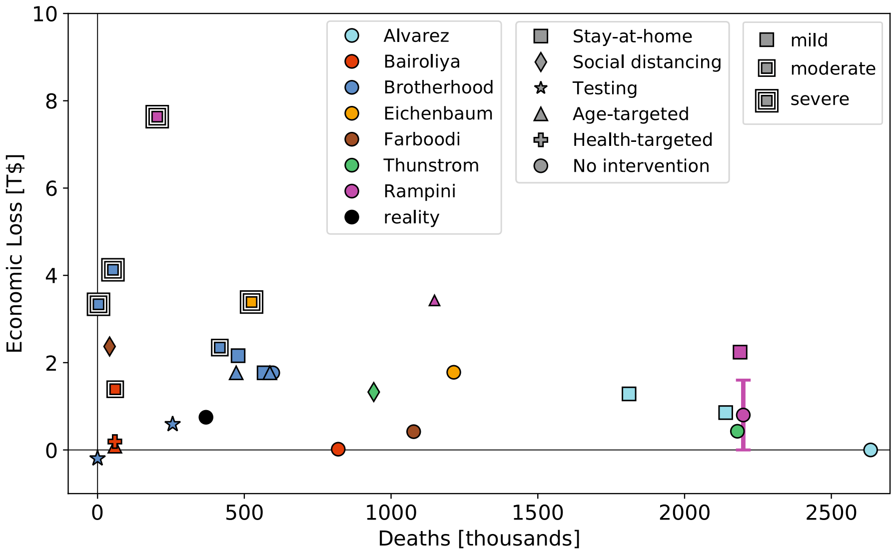

Fig. 2 Summary of findings from seven hybrid studies, showing the calculated economic loss against the total number of deaths. The circles indicate the baseline (no intervention) cases. The other shapes represent different types of interventions. The outlines around the squares represent different intensities of stay-at-home policies. The black circle represents the estimated loss of GDP and deaths realized in the United States as of December 31st, 2020.

to predict out reliably to one year, much less farther. Second, most of the models provided the time period of one year as a reference point in their analysis and, lastly, because COVID-19 vaccines were developed in about a year.

VSL and VSLY Ranges

All of the studies analyzed for this work used some form of VSL or VSLY as a reference for measuring the cost-effectiveness of strategies or scenarios or for scaling their agents' utility discount functions, with fairly broad ranges of these values across studies. We present the full range of VSLs and VSLYs used in these studies, which acknowledges that there is no consensus on the exact valuation that should be used for statistical lives or statistical life-years, thereby increasing the papers' comparability. The common value used by the US government is around 9.9 million (Greenstone and Nigam 2020). Brotherhood et al. (2021) and Eichenbaum et al. (2021) set their VSLs at 9.3 million, Thunström et al. (2021) use a VSL of 10 million, and Greenstone and Nigam (2020), adjusting for income growth and inflation, set theirs at 11.5 million. Of the authors using a VSLY, Alvarez et al. (2021) used a value of around 2.6 million per life, and Farboodi et al. (2021) used a value of 3.915 million to account for the fact that mortality rates were much higher for the elderly population than the young population. Bairoliya and İmrohoroğlu (2020) adopt

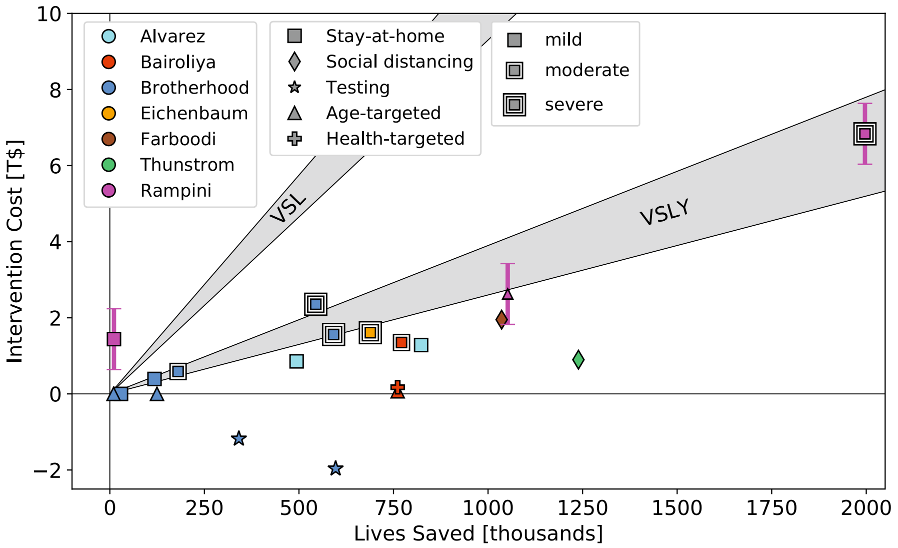

Fig. 3 Summary of findings from seven hybrid studies, showing the lives saved (baseline models minus the intervention models) against intervention costs (intervention model economic loss minus the baseline model economic loss). The vertical dashed lines represent the lives saved as calculated from the purely epidemiological models of Ferguson et al. (2020) and Greenstone and Nigam (2020) when interventions are applied. The other shapes represent different types of interventions. The outlines around the squares represent different intensities of stay-at-home policies. The VSL, varying from 9.3 to 11.5 million dollars, represents the statistical value of life, and the VSLY, varying from 2.6 to 3.9 million dollars per life, accounts for the average number of years left to the typical COVID-19 victim. Since Rampini (2020) does not report a baseline model, we use the Greenstone and Nigam (2020) calculations because this is the only model that has baseline deaths greater than the calculated deaths in Rampini's (2020) models. In the absence of a baseline value of economic loss, we use a value of 0.8 T$, which represents the approximate mean of the baseline losses from the other studies. The error bars on the Rampini points represent an uncertainty of +/− 0.8 T$ on this baseline.

age-specific VSLs from Aldy and Viscusi (2008) and assume those 65 or older have a VSL of $1 million.

Cost–Benefit Ratio

The data from the studies examined here are presented in two different formats. Figure 2 presents the data in the parameter space of deaths and economic losses. Figure 3 presents a similar parameter space, but subtracts the baseline economic loss and baseline deaths from each of the cases separately for each study. By subtracting the baseline scenario (the no-intervention strategy) from each case, the primary measures of efficacy become the number of lives saved through intervention (the horizontal axis) and the economic cost of intervention (the vertical axis). We define the ratio of these values, the vertical divided by the horizontal, as the cost–benefit ratio and employ this as a

Table 3 Ranked list of intervention by cost–benefit ratio

Paper	Intervention type	Cost–benefit ratio [$M/life]
Brotherhood et al.	**Testing**	**− 3.45**
Brotherhood et al.	**Testing**	**− 3.29**
Bairoliya and İmrohoroğlu	Age-targeted	0.09
Bairoliya and İmrohoroğlu	Health-targeted	0.23
Brotherhood et al.	**Age-targeted**	**0**
Brotherhood et al.	**Age-targeted**	**0**
Brotherhood et al.	**Stay-at-home (mild)**	**0**
Thunstrom et al.	Social distancing	0.73
Alvarez et al.	Stay-at-home (mild)	1.56
Alvarez et al.	Stay-at-home (mild)	1.73
Bairoliya and İmrohoroğlu	Stay-at-home (moderate)	1.81
Farboodi et al.	**Social distancing**	**1.88**
Eichenbaum et al.	**Stay-at-home (severe)**	**2.34**
Rampini	Age-targeted	2.4
Lower limit of VSLY		2.6
Brotherhood et al.	**Stay-at-home (severe)**	**2.65**
Brotherhood et al.	**Stay-at-home (moderate)**	**3.26**
Brotherhood et al.	**Stay-at-home (mild)**	**3.33**
Rampini	Stay-at-home (severe)	3.42
Upper limit of VSLY		3.9
Brotherhood et al.	**Stay-at-home (severe)**	**4.33**
Lower limit of VSL		9.3
Upper limit of VSL		11.5
Rampini	Stay-at-home (mild)	132.06

*The rows highlighted in bold indicate studies that included behavioral change.

measure of relative efficacy that can be compared across studies and across intervention strategies. Viewed thusly, low cost–benefit ratios represent low cost per life saved. The reason for representing the costs and benefits in this way is that the ratio of these values (cost divided by benefit) has the same units as the VSL and VSLY. In this sense, we can interpret the VSL and VSLY as a kind of cost–benefit ratio, that is, the cost we should pay to save a life or, equivalently, a valuation of the loss incurred with each death.

Analysis

We return to the question of whether the hybrid economic-epidemiological models created consistent outcomes that are useful for informing policy? Our answer is presented in the form of two graphic representations of the data, in Figures 2 and 3, and in the ranking of cost–benefit ratios in Table 3. While the figures present the data in a similar format, the subtle differences between them have important consequences for the insight we can derive from them. In short, a raw presentation of the data like that in Figure 2 reveals no discernable patterns. In contrast, subtraction of the baseline effect from each intervention scenario, as shown in Figure 3, allows trends to emerge.

Figure 2 represents the model data as taken directly from each study. Each scenario is presented in the parameter space represented by the number of deaths calculated from the epidemiological model and the associated economic losses. No clear patterns emerge from the data presented in Figure 2, though it suggests that the severe stay-at-home policies are more economically costly than other interventions, yet they also result in comparably few deaths. In comparison, the other intervention methods tend to have much lower economic costs. Little else can be stated with confidence.

The implicit assumption in the representation of the data as presented in Figure 2 is that we are comparing against a world without a pandemic. The better lens for assessing possible interventions is to compare each scenario against its associated "do-nothing" policy, represented by the baseline models. Such an interpretation of the data is achieved by transforming the model data through subtraction of the baseline values, which reveals the cost of intervention and the number of lives saved through intervention.

Figure 3 presents the data thus transformed, which provides a clearer sense of the underlying patterns. The ratio of these values (the cost of intervention divided by the number of lives saved) should be interpreted as the cost–benefit ratio for each strategy. That is, points in this space that exist at a higher cost for a given number of lives saved represent a higher cost–benefit ratio (i.e., a more expensive policy, in the normalized sense). Importantly, the VSL and VSLY, which are themselves forms of a cost–benefit ratio, can be represented in this parameter space as straight lines. The swaths for each of the VSL and VSLY represent the range of values encountered in the literature. While there is substantial disagreement about whether the VSL or VSLY is the more appropriate measure to be used, three things are clear with regard to their use. First, given that the VSL is always greater than the VSLY, any intervention that appears in the parameter space of Figure 3 above the VSL swath must be considered expensive on all accounts. Second, similar to the first point, any intervention appearing in the parameter space beneath the VSLY swath must be considered as cost effective. Lastly, those interventions falling between the VSL and VSLY are ambiguous and must be carefully considered. Table 3 lists the cost–benefit ratio for all studies shown in Figure 3, ranked from least cost per life saved to highest cost per life saved, in comparison with the VSL and VSLY.

Discussion

Across studies, we see patterns in the cost–benefit ratios of the interventions despite substantial variation between studies. In analyzing the data, we argue that we should pay attention to two important parameters, the cost–benefit ratio for each intervention and the total number of lives saved by that intervention. The starkest result to emerge from the data presented in Figure 3 and Table 3 is that all models have a cost-benefit ratio that is within or below the VSLY range with the exception of two cases. The high cost-benefit ratio Rampini (2020) case should be interpreted with caution because it did not include behavioral change, a factor shown in other studies to greatly reduce deaths. If combinations of policies can be used to achieve greater efficiency, then it must be that these optimal interventions will be even further below the VSLY band. This is an important result because this shows that a broad range of intervention strategies must generally be interpreted as highly effective and, therefore, also socially responsible.

As we compare strategies against each other, it becomes clear that the stay-at-home policies have cost–benefit ratios that are larger (more expensive per life saved), on average, than all other types of interventions. The comparatively high cost of stay-at-home policies should be a point of general interest and concern, as many government interventions have focused on various forms of stay-at-home orders. This tendency toward higher cost–benefit ratios as a group also shows that our analysis does not hinge on the definitions of the *mild-moderate-severe* classification of the stay-at-home orders. Although, this additional categorization may be useful for eliciting more information about what kinds of stay-at-home policies are more or less effective than others.

In comparison with stay-at-home policies, the negative cost–benefit ratios of the testing policies show that they are highly favorable. Negative cost–benefit ratios imply that the intervention reduces losses compared with the do-nothing cases. In models that include behavioral change, an unmitigated pandemic comes with significant economic costs as people reduce risk through curtailment of spending and production activities (Brotherhood et al. 2021; Eichenbaum et al. 2021; Farboodi et al. 2021). Without additional intervention efforts by the government, the pandemic has widespread effects, even with agent behavioral change. It follows that intervention strategies that increase agents' confidence through individualized response, without a coarse population-level mandate, allow agents to maintain higher levels of consumption and productivity. Thus, widespread testing emerges as a clear winner in terms of effectiveness at reducing death and preserving the economy. The more successful Brotherhood testing model assumed all individuals presenting symptoms were tested, and all people testing positive were quarantined. They found that such an undertaking would require approximately 15 million tests per week at the height of the pandemic (Brotherhood et al. 2021, p. 30). While this is a large number of tests, it is not out of the realm of possibility. For example, data from Johns Hopkins University show that the US was conducting about 13 million tests per week in January 2021 near the period of peak infection (Johns Hopkins University and Medicine 2021). However, in the United States, it was only well into the pandemic that such testing capacity was established.

Social distancing policies tend to have higher cost–benefit ratios than testing, but still occur below the VSLY range. Similar to testing, social distancing policies entail relatively fewer restrictions, thereby allowing greater continuity of economic activities with safer interaction. Moreover, empirical analyses of social distancing suggest that such policies effectively slowed the spread of the disease (Price and Holm 2021). In a study of 37 OECD member countries, stay-at-home orders showed low effectiveness at saving lives for durations of less than 70 days, and beyond this were less effective than closing schools and public transportation, which they found to have a persistently significant impact (Clyde et al. 2021).

Some economists have argued that achieving lower-cost and higher-benefit outcomes can be accomplished through health-targeted and/or age-targeted policies involving stay-at-home policies for specific groups (Bairoliya and İmrohoroğlu 2020; Acemoglu et al. 2021). This holds to some extent for the Bairoliya and İmrohoroğlu (2020) and the Rampini (2020) models because they assumed that the disease passes through the young population (reaching herd immunity) before the restrictions on the elderly are lifted. Brotherhood et al. (2021) found age-targeted policies to be low-cost, though saving relatively few lives. This is because they assume that the old and young interact sufficiently to counteract some of the gains due to isolation and that the policy is lifted before herd immunity is reached. In reality, age-targeted policies were relatively rare in pandemic responses, and it may be that age-targeted policies were not politically or logistically feasible. Requiring only one age group or risk group to stay at home could be perceived as discrimination against that group. Besides, multigenerational households or workplaces make such policies very difficult to implement.

In general, the decision to use a VSL or a VSLY as a point of reference for these types of analyses can drastically change the conclusions as to which policy actions are considered feasible. The COVID-19 pandemic has had higher mortality for older age groups; if the VSLY is used, rather than the VSL, the value of lives saved is diminished. If instead of causing a disproportionate effect on the elderly, the COVID-19 pandemic had disproportionately affected the young, as the 1918 Spanish flu did (Gagnon et al. 2013), then the population-weighted VSLY would shift upward, toward the VSL. For the most part, the papers here did not consider the ethics of their choice, but the use of the VSLY by some papers appeared to represent some discomfort with the financial onus of using the VSL. We do not attempt to resolve the ethical question of which value is the "right" value to use (or if life is priceless). Instead, we have provided the information of lives saved versus economic losses for the full range of possible valuations of life given by the papers.

Conclusion

The policymaker confronts a difficult task; ambiguities abound in the epidemiological data, the economic modeling, and the public response to government intervention. Early in the pandemic, epidemiologists were uncertain of disease characteristics resulting in large differences in estimates of infection and mortality. Initial empirical studies on masking, social distancing and behavioral change found that there was room to have saved more lives with decisive mask and social distancing mandates.[4] Conversations around disease mitigation efforts have, at times, considered a tradeoff

[4] For instance, Lyu and Wehby (2020) suggest that early implementation of mask mandates would have averted more than 200,000 COVID-19 cases by May 22, 2021. Chernozhukov, Kasahara, and Schrimpf (2021) argue that nationally mandating face masks for employees early in the pandemic could have reduced the weekly growth rate of cases and deaths by more than 10 percentage points in late April and between 19 and 47 thousand saved lives by the end of May.

between deaths and economic losses, but there has never been consensus on how, or even whether, we should attempt to compare lives and dollars. The valuation of life used in the studies examined here varied significantly from 2.6 million to 11.5 million per life. However, if one accepts that at a societal level there must be some accounting of lives and economy and that a calculation where each life is of infinite value is not tenable, then an approach along the lines of using a VSL or VSLY may have some merit. In addition, economic losses not only affect the economy but also negatively impact physical and mental health. In an article studying the impact of job loss on mental health, the authors report the incidence of depression in the U.S. population increased from 5% in 2019 to almost 20% in 2021 (Huato and Chavez 2021). Lastly, government stimulus, a factor mostly not included in the models studied, largely mitigated many of the potential GDP losses estimated by the models, but we have yet to see its long-term impacts.

Despite the uncertainties inherent in any single study, each of which uses imperfect data and models that are gross simplifications of reality, the existing analyses are helpful, especially when viewed collectively, because they allow the efficacy of various intervention strategies to be assessed. While such studies have been performed, all of them fall somewhat short in clarity, transferability, and representation in practical terms that are useful for policymakers. If a purpose of these studies is to guide policymakers, what would be valuable for future studies is to use language that is more approachable, to model output in real units (dollars or fraction of GDP), to report results at time periods of six months and one year, and to explicitly include the baseline models as a point of reference.

One of our main contributions to this effort is the baseline subtracted cost–benefit method that allows for interventions to be compared across studies, thereby permitting the underlying patterns in the models to emerge. This method enables us to view the two hypotheticals, a “do-nothing” approach with a variety of possible policies. We extrapolate from this approach that some of the least severe intervention measures, like testing and social distancing, are also highly effective at saving lives and should be considered one of the most effective weapons in our fight against pandemics. While stay-at-home policies are generally effective at saving lives, they also tend to be more expensive. However, stay-at-home policies may be especially important at the beginning of a pandemic, prior to the development of accurate and widespread testing or widespread vaccines. Additionally, the role of age-targeted and health-targeted policies is interesting to this discussion. Their very low, near-zero, cost–benefit ratio is compelling, and it may be a significant advantage for the entire population for policymakers to examine whether/how such policies could be implemented, in whole or in part. Viewing the data in this way is the manifestation of the idea that we do not get to choose a world without a pandemic, only how we respond to it.

Appendix

Table 4 lists the papers used in this analysis.

Table 4 Articles included

Author	Title	Date	Published
Alvarez, Fernando, David Argente, and Francesco Lippi	A simple planning problem for COVID-19 lockdown, testing and tracing	September, 2021	American economic review: insights
Bairoliya, Neha, and Ayse İmrohoroğlu	Macroeconomic consequences of stay-at-home policies during the COVID-19 pandemic	May, 2020	Centre for economic policy research, COVID economics
Brotherhood, Luiz, Philipp Kircher, Cezar Santos, and Michèle Tertilt	An economic model of the Covid-19 pandemic with young and old agents: Behavior, testing and policies	April, 2021	Previous version circulated as CEPR DP 14695 with the title "An economic model of the Covid-19 epidemic: the importance of testing and age-specific policies"
Eichenbaum, Martin S, Sergio Rebelo, and Mathias Trabandt	The macroeconomics of epidemics	October, 2021	The review of financial studies
Farboodi, Maryam, Gregor Jarosch, and Robert Shimer	Internal and external effects of social distancing in a pandemic	June, 2021	Journal of economic theory
Ferguson, Neil N., Daniel Laydon, Gemma Nedjati-Gilani, Natsuko Imai, Kylie Ainslie, Marc Baguelin, Sangeeta Bhatia, et al.	Report 9 - impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand	March, 2020	Imperial College London
Greenstone, Michael, and Vishan Nigam	Does social distancing matter?	March, 2020	Centre for economic policy research, COVID economics
Rampini, Adriano A	Sequential lifting of COVID-19 interventions with population heterogeneity	April, 2020	NBER working paper
Thunstrom, Linda, Stephen Newbold, David Finnoff, Madison Ashworth, and Jason F. Shogren	The benefits and costs of using social distancing to flatten the curve for COVID-19	September, 2021	Journal of benefit-cost analysis

Acknowledgements The authors would like to thank Michael Ash, João Paulo de Souza, Lisi Krall, Dan Moyer and three anonymous referees for their helpful comments. Any errors are our own.

References

Acemoglu, Daron, Victor Chernozhukov, Iván. Werning, and Michael D. Whinston. 2021. Optimal targeted lockdowns in a multigroup SIR model. *American Economic Review: Insights* 3(4): 487–502.

Ackerman, Frank, and Lisa Heinzerling. 2004. *Priceless: On knowing the price of everything and the value of nothing*, Reprint, New York: The New Press.

Ackerman, Frank, and Lisa Heinzerling. 2001. *If it exists, It's getting bigger: Revising the value of a statistical life.* Global Development and Environment Institute Working Paper No. 01-06.

Aldred, Jonathan. 2022. Guiding Covid policy: Cost-benefit analysis and beyond. *Cambridge Journal of Economics* 46(3): 589–608.

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*.

Aldy, Joseph E., and W. Kip Viscusi. 2008. Adjusting the value of a statistical life for age and cohort effects. *Review of Economics & Statistics* 90(3): 573–581. https://doi.org/10.1162/rest.90.3.573.

Alvarez, Fernando, David Argente, and Francesco Lippi. 2021. A simple planning problem for COVID-19 lock-down, testing, and tracing. *American Economic Review: Insights* 3(3): 367–382.

Anderson, Elizabeth. 1993. *Value in ethics and economics*. Harvard University Press.

Aschwanden, Christie. 2021. Five reasons why COVID herd immunity is probably impossible. *Nature* 591(7851): 520–522. https://doi.org/10.1038/d41586-021-00728-2.

Ashenfelter, Orley. 2006. Measuring the value of a statistical life: problems and prospects. *The Economic Journal* 116(510): C10–C23.

Avery, Christopher, William Bossert, Adam Clark, Glenn Ellison, and Sara Fisher Ellison. 2020. An economist's guide to epidemiology models of infectious disease. *Journal of Economic Perspectives* 34(4): 79–104.

Bairoliya, Neha, and Ay.şe İmrohoroğlu. 2020. Macroeconomic consequences of stay-at-home policies during the COVID-19 pandemic. *Centre for Economic Policy Research, COVID Economics* 13: 68–87.

Banzhaf, H. Spencer. 2014. Retrospectives: The cold-war origins of the value of statistical life.". *Journal of Economic Perspectives* 28(4): 213–226.

Banzhaf, H. Spencer. 2022. The value of statistical life: A meta-analysis of meta-analyses. *Journal of Benefit-Cost Analysis* 13(2): 182–197.

Broome, John. 1985. The economic value of life. *Economica* 52(207): 281–294.

Brotherhood, Luiz, Philipp Kircher, Cezar Santos, and Michèle Tertilt. 2021. An economic model of the Covid-10 pandemic with young and old agents: Behavior, testing and policies (April). Available on author's website: https://sites.google.com/site/luizbrotherhood/research. Previous version: *An economic model of the Covid-19 epidemic: The importance of testing and age-specific policies*. CEPR Discussion Paper No. DP1495 (May 2020).

Bubar, Kate M., Kyle Reinholt, Stephen M. Kissler, Marc Lipsitch, Sarah Cobey, Yonatan H. Grad, and Daniel B. Larremore. 2021. Model-informed COVID-19 vaccine prioritization strategies by age and serostatus. *Science* 371(6532): 916–921.

Bürgi, Constantin, and Klaus Wohlrabe. 2022. The influence of Covid-19 on publications in economics: bibliometric evidence from five working paper series. *Scientometrics* 127(9): 5175–5189.

Cameron, Trudy Ann. 2010. Euthanizing the value of a statistical life. *Review of Environmental Economics and Policy* 4: 2.

Cameron, Trudy Ann, and J.R. DeShazo. 2013. Demand for health risk reductions. *Journal of Environmental Economics and Management* 65(1): 87–109.

Centers for Disease Control and Prevention. 2021. *US State, Territorial, and County Stay-At-Home Orders: March 15-May 5 by County by Day Data.* Centers for Disease Control and Prevention (February 10). https://data.cdc.gov/Policy-Surveillance/U-S-State-Territorial-and-County-Stay-At-Home-Orde/qz3x-mf9n.

Chernozhukov, Victor, Hiroyuki Kasahara, and Paul Schrimpf. 2021. Causal impact of masks, policies, behavior on early covid-19 pandemic in the US. *Journal of Econometrics* 220(1): 23–62.

Clyde, William, Andreas Kakolyris, and Georgios Koimisis. 2021. A study of the effectiveness of governmental strategies for managing mortality from COVID-19. *Eastern Economic Journal* 47(4): 487–505.

Colmer, Jonathan. 2020. What is the meaning of (statistical) life? Benefit–cost analysis in the time of COVID-19. *Oxford Review of Economic Policy* 36(Supplement_1): S56–S63.

Dorman, Peter. 2009. *Markets and mortality: Economics, dangerous work, and the value of human life*. Cambridge: Cambridge University Press.

D'Souza, Gypsyamber, and David Dowdy. 2021. *What is Herd Immunity and How Can We Achieve It With COVID-19?* Johns Hopkins Bloomberg School of Public Health (April 6). https://www.jhsph.edu/covid-19/articles/achieving-herd-immunity-with-covid19.html

Eichenbaum, Martin S., Sergio Rebelo, and Mathias Trabandt. 2021. The macroeconomics of epidemics. *The Review of Financial Studies* 34(11): 5149–5187.

Evans, Mary F., and V. Kerry Smith. 2006. Do we really understand the age–VSL relationship? *Resource and Energy Economics* 28(3): 242–261.

Farboodi, Maryam, Gregor Jarosch, and Robert Shimer. 2021. Internal and external effects of social distancing in a pandemic. *Journal of Economic Theory* 196: 105293.

Favero, Carlo A., Andrea Ichino, and Aldo Rustichini. (2020). Restarting the economy while saving lives under Covid-19. Social Science Research Network (November 8th). https://ssrn.com/abstract=3580626

Ferguson, Neil, Daniel Laydon, Gemma Nedjati Gilani, Natsuko Imai, Kylie Ainslie, Marc Baguelin, Sangeeta Bhatia et al. 2020. *Report 9—Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand*. Imperial College London, (March 16). http://www.imperial.ac.uk/medicine/departments/school-public-health/infectious-disease-epidemiology/mrc-global-infectious-disease-analysis/covid-19/report-9-impact-of-npis-on-covid-19/

Flaschel, Peter, Giorgos Galanis, Daniele Tavani, and Roberto Veneziani. 2021. *Pandemics and aggregate demand: A framework for policy analysis* FMM Working paper No. 62, Düsseldorf: Hans-Böckler-Stiftung, Macroeconomic Policy Institute (IMK), Forum for Macroeconomics and Macroeconomic Policies.

Friedman, H.S. 2020. *Ultimate price: The value we place on life*. Oakland: University of California Press.

Gagnon, Alain, Matthew S. Miller, Stacey A. Hallman, D. Robert Bourbeau, Ann Herring, David JD. Earn, and Joaquin Madrenas. 2013. Age-specific mortality during the 1918 influenza pandemic: unravelling the mystery of high young adult mortality. *PLoS ONE* 8(8): e69586.

Glover, Andrew, Jonathan Heathcote, Dirk Krueger, and José-Víctor Ríos-Rull. 2023. Health versus wealth: On the distributional effects of controlling a pandemic. *Journal of Monetary Economics*.

Greenstone, Michael, and Vishan Nigam. 2020. Does Social Distancing Matter? Becker Friedman Institute for Economics Working Paper No. 2020-26, University of Chicago.

Gollier, Christian. 2020. Cost–benefit analysis of age-specific deconfinement strategies. *Journal of Public Economic Theory* 22(6): 1746–1771.

Hall, Robert E., Charles I. Jones, and Peter J. Klenow. 2020. *Trading off consumption and covid-19 deaths*. NBER Working Paper No. 27340, Cambridge: National Bureau of Economic Research.

Hilsenrath, Jon and Stephanie Armour. 2020. As economic toll mounts, Nation Ponders Trade-Offs. *Wall Street Journal*, (March 23). https://www.wsj.com/articles/as-economic-toll-mounts-nation-ponders-the-trade-offs-11584970165

Hood, Katherine. 2017. The science of value: Economic expertise and the valuation of human life in US federal regulatory agencies. *Social Studies of Science* 47(4): 441–465.

Huato, Julio, and Aida Chavez. 2021. Household income, pandemic-related income loss, and the probability of anxiety and depression. *Eastern Economic Journal* 47(4): 546–570.

Johns Hopkins University and Medicine. 2021. *COVID-19 Map*. Johns Hopkins Coronavirus Resource Center. https://coronavirus.jhu.edu/map.html

Jones, Callum, Thomas Philippon, and Venky Venkateswaran. 2021. Optimal mitigation policies in a pandemic: Social distancing and working from home. *The Review of Financial Studies* 34(11): 5188–5223.

Jones-Lee, Michael W. 1976. *The value of life*. Chicago: The University of Chicago Press.

Juranek, Steffen, Jörg. Paetzold, Hannes Winner, and Floris Zoutman. 2021. Labor market effects of COVID-19 in Sweden and its neighbors: Evidence from administrative data. *Kyklos* 74(4): 512–526.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante. 2020. *The great lockdown and the big stimulus: Tracing the pandemic possibility frontier for the US*. NBER Working Paper No. 27794, Cambridge: National Bureau of Economic Research.

Kermack, William Ogilvy, and Anderson G. McKendrick. 1927. A Contribution to the Mathematical Theory of Epidemics. *Proceedings of the Royal Society of London. Series a, Containing Papers of a Mathematical and Physical Character* 115(772): 700–721.

Seonglim, Lee, Sang-Hee. Sohn, Eunyoung Rhee, and Yoon G. Lee. 2014. Consumption patterns and economic status of older households in the United States. *Monthly Labor Review* 137: 1.

Lyu, Wei, and George L. Wehby. 2020. Community use of face masks and COVID-19: Evidence from a natural experiment of state mandates in the US. *Health Affairs* 39(8): 1419–1425.

Matrajt, Laura, Julia Eaton, Tiffany Leung, and Elizabeth R. Brown. 2021. Vaccine optimization for COVID-19: Who to vaccinate first? *Science Advances* 7(6): eabf1374.

Mishan, Ezra J. 1971. Evaluation of life and limb: a theoretical approach. *Journal of Political Economy* 79(4): 687–705.

Murray, Eleanor J. 2020. Epidemiology's time of need: COVID-19 calls for epidemic-related economics. *Journal of Economic Perspectives* 34(4): 105–120.

Nordström, Peter, Marcel Ballin, and Anna Nordström. 2022. Risk of SARS-CoV-2 reinfection and COVID-19 hospitalisation in individuals with natural and hybrid immunity: A retrospective, total population cohort study in Sweden. *The Lancet Infectious Diseases* 22(6): 781–790.

Price, Gregory, and Eric van Holm. 2021. The effect of social distancing on the early spread of the novel coronavirus. *Social Science Quarterly* 102(5): 2331–2340.

Qin, Amy. 2020. China may be beating the coronavirus, at a painful cost. *The New York Times* (March 7). https://www.nytimes.com/2020/03/07/world/asia/china-coronavirus-cost.html

Rampini, Adriano A. 2020. *Sequential lifting of COVID-19 interventions with population heterogeneity* NBER Working Paper No. 27063, Cambridge: National Bureau of Economic Research.

Robinson, Lisa A., Ryan Sullivan, and Jason F. Shogren. 2021. Do the benefits of COVID-19 policies exceed the costs? Exploring uncertainties in the age–VSL relationship. *Risk Analysis* 41(5): 761–770.

Schelling, Thomas C. 1968. The Life You Save May be Your Own. *Problems in Public Expenditure Analysis: Papers Presented at a Conference of Experts, Washington, DC : Brookings Institution, September 15 - 16, 1966,* 127–162.

Shaw, Clara L., and David A. Kennedy. 2021. What the reproductive number R0 can and cannot tell us about COVID-19 dynamics. *Theoretical Population Biology* 137: 2–9.

Simon, Nathalie B., Chris Dockins, Kelly B. Maguire, Stephen C. Newbold, Alan J. Krupnick, and Laura O. Taylor. 2019. Policy brief—What's in a name? A search for alternatives to "VSL." *Review of Environmental Economics and Policy* 13(1): 155–161.

Sunstein, C.R. 2014. *Valuing life: Humanizing the regulatory state*. Chicago: University of Chicago Press.

Thunström, Linda, Stephen C. Newbold, David Finnoff, Madison Ashworth, and Jason F. Shogren. 2020. The benefits and costs of using social distancing to flatten the curve for COVID-19. *Journal of Benefit-Cost Analysis* 11(2): 179–195.

Viscusi, W. Kip. 1992. *Fatal tradeoffs: Public and private responsibilities for risk*. New York: Oxford University Press.

Viscusi, W. Kip. 2018. *Pricing lives: Guideposts for a safer society*. Princeton: Princeton University Press.

Viscusi, W. Kip. 2021. Extending the domain of the value of a statistical life. *Journal of Benefit-Cost Analysis* 12(1): 1–23.