LAW AND ECONOMICS SEMINAR
Autumn 2026

Professor John J. Donohue III
Professor A. Mitchell Polinsky

Thursday, October 1, 2026
4:15 - 5:45 p.m.
Stanford Law School
Room 285

"Valuing Future Lives"

Daniel Hemel
New York University Law School

Note: It is expected that you will have reviewed the speaker's paper before the seminar. Because the paper is longer than average, Professor Hemel was asked to provide a reader's guide for readers who are only able to read about half of the paper. Here is his response:

***Reader's guide provided by the author*: "We wrote this paper because we thought that federal agencies -- under Democratic and Republican administrations -- were picking numbers for the income elasticity of the value of a statistical life without thinking rigorously about the justifications for those numbers or the implications of their choices. We wanted to explain -- in accessible terms -- why the income elasticity of the VSL is so important in regulatory cost-benefit analysis and how the income elasticity of the VSL relates to other parameters such as the discount rate. We encourage everyone to read the Introduction. Readers who are familiar with the VSL concept can skip Part I. Part II is the heart of the argument. If readers only have time to go a bit beyond that, we recommend focusing on Section III.A.2 (regarding the EPA), both because EPA is the agency for whom the income elasticity of the VSL matters most and because EPA is the agency that discounts future lives most dramatically."**

VALUING FUTURE LIVES

Daniel J. Hemel* & Jonathan S. Masur**

September 16, 2026

—DRAFT—

ABSTRACT

Federal regulation often involves a tradeoff between monetary costs in the present and life-saving benefits in the future. A central question in regulatory cost-benefit analysis is how to assign a present dollar value to future lives so that future lives and present dollars can be compared. For regulations that are projected to prevent deaths years or decades down the road, agencies make two key analytical moves. First, they adjust the value of a statistical life upward to reflect the fact that society's willingness to pay to save lives will rise as people become wealthier in the future. Second, they discount future benefits downward to reflect the fact that the wealthier people of the future will derive less utility from each dollar precisely because they are wealthier. In theory, these upward and downward adjustments could offset exactly, such that lives saved in the future carry the same weight in agency cost-benefit analyses as lives saved today. In practice, agencies adjust the value of a statistical life upward at a much slower rate than they discount future benefits. The consequence is that agencies assign much less weight to lives in the future than to lives in the present. For example, the Environmental Protection Agency considers a life saved fifty years from now to be worth 69% to 95% less than a life saved today. Other agencies likewise apply large haircuts to future lives. The result is a dramatic anti-regulatory bias, particularly with respect to regulations that will have large effects years in the future, such as measures to curb climate change.

This Article critically analyzes federal agencies' practice of devaluing future lives. Through step-by-step reconstructions of methodologies employed by four agencies—the Department of Transportation, the Environmental Protection Agency, the Department of Health and Human Services, and the Occupational Safety and Health Administration—the Article shows how current regulatory practice substantially discounts the prevention of premature deaths over the medium and long term. It then demonstrates—through case studies of recent air pollution and motor safety rulemakings—how the devaluation of future lives stacks the deck against regulatory action. It concludes by calling for agencies to adopt a presumption that lives in the future have the same value as lives saved today, and it explores the implications of this proposal for the broader enterprise of cost-benefit analysis.

INTRODUCTION

A central question facing agencies across the federal government is how to value lives saved in the future. Most high-cost federal regulations involve a tradeoff between dollars on one side of the ledger and lives on the other.[1] Federal agencies typically use a measure—known as the

* John S. R. Shad Professor of Law, New York University School of Law.
** John P. Wilson Professor of Law, University of Chicago Law School. We thank Matt Adler, James Hammitt, Aaron Jacobowitz, Marcel Kahan, David Kamin, Eric Posner, Cass Sunstein, David Weisbach, numerous editors

value of a statistical life (VSL)—to compare dollar costs to life-saving benefits. For example, suppose that an agency sets the VSL at $10 million. If a regulation (e.g., a clean air standard or a vehicle safety standard) will avert one premature death this year, the regulation will pass a cost-benefit test so long as it imposes less than $10 million in costs in the same year. By the mid-2020s, agencies across the executive branch had converged upon VSL figures in the range of $14 million.[2]

In many cases, the life-saving benefits of federal regulations will accrue far in the future rather than in the same year that costs are incurred. For example, a clean air standard adopted in 2026 might result in lower particulate matter emissions starting in 2028 or 2029—and thus in fewer deaths from asthma and chronic obstructive pulmonary disease many decades hence. A vehicle safety standard applied to model year 2026 vehicles might result in fewer traffic deaths into the 2030s and 2040s, as some model year 2026 vehicles will likely remain in use then. Yet these same regulations often involve costs that accrue immediately or in the near term. If a vehicle safety standard is applied to model year 2026 cars, consumers will have to pay for those safety features in 2026 when they purchase the cars. A clean air standard that makes electricity production more expensive may raise costs for consumers long before the public health benefits manifest. Federal agencies conducting cost-benefit analyses of regulations with later-in-time life-saving benefits and

---

with the Yale Journal on Regulation, and participants in the American Law and Economics Association Annual Meeting, the NYU School of Law Internal Law and Economics Workshop, the University of Michigan Law School Law and Economics Workshop, and the Society of Benefit Cost Analysis Annual Conference for helpful comments and suggestions, and Charles Tammons for superb research assistance. Masur thanks the Wachtell, Lipton, Rosen & Katz Program in Behavioral Law, Finance & Economics and the David & Celia Hilliard Fund for research support.

[1] According to one estimate, 70% of the total monetized benefits of major federal agency regulations "were directly attributable to the monetized value of reducing early mortality." Jonathan Colmer, *What Is the Meaning of (Statistical) Life? Benefit-Cost Analysis in the Time of COVID-19*, 36 OXFORD REV. ECON. POL'Y S56, S57 (2020); *see also* Daniel J. Hemel, *Regulation and Redistribution with Lives in the Balance*, 89 U. CHI. L. REV. 649, 666 (2022) (noting that out of 24 regulations promulgated by federal agencies from 2001 to 2018 with costs exceeding $1 billion, 20 were promulgated by the Environmental Protection Agency and/or the Department of Transportation, and the bulk of the benefits of those agencies' regulations come through reductions in mortality and morbidity).

[2] *See, e.g.*, Aaron Kearsley, *HHS Standard Values for Regulatory Analysis, 2026*, OFF. OF THE ASS'T SEC'Y FOR PLANNING & EVALUATION, U.S. DEP'T OF HEALTH & HUMAN SERVS. (Mar. 6, 2026), https://aspe.hhs.gov/reports/standard-ria-values (central estimate of $14.1 million in constant 2025 dollars for 2025); U.S. DEP'T OF TRANSP., DEPARTMENTAL GUIDANCE ON VALUATION OF A STATISTICAL LIFE IN ECONOMIC ANALYSIS (Mar. 20, 2026), https://www.transportation.gov/office-policy/transportation-policy/revised-departmental-guidance-on-valuation-of-a-statistical-life-in-economic-analysis ($14.2 million for 2025); U.S. Env'tl Prot. Agency, Economic Analysis for the Final Per- and Polyfluoroalkyl Substances National Primary Drinking Water Regulation Appendices, app. J., tbl.J-1 (Apr. 2024), https://www.epa.gov/system/files/documents/2024-04/pfas-npdwr_final-rule_ea_appendices.pdf (projected VSL for 2024 of $11,529,327 in 2022 dollars). Note that EPA's April 2024 estimate, which was stated in 2022 dollars, corresponds to approximately $14.4 million in 2026 dollars based on Consumer Price Index-measured inflation. *See* U.S. BUREAU OF LABOR STATISTICS, CPI INFLATION CALCULATOR, https://www.bls.gov/data/inflation_calculator.htm (last visited Apr. 11, 2026) (comparing March 2022 to March 2026).

near-in-time costs therefore must consider how to compare lives saved in the future with costs incurred in the present.

Two competing intuitions influence the valuation of future lives.[3] On the one hand, cost-benefit analysis generally relies on "willingness to pay" and "willingness to accept" concepts. For example, a VSL of $14 million reflects the assumption that an average American would be willing to pay up to $14 for a 1-in-1 million reduction in the risk of death—or alternatively, that an average American would be willing to accept a 1-in-1 million increase in the risk of death in exchange for a payment of $14 or more.[4] As income increases, the amount that individuals are willing to pay for a given reduction in the risk of death generally also increases (and, commensurately, the amount that individuals will demand in exchange for accepting a given increase in the risk of death rises). Since individuals in the future will probably be richer than individuals today, their VSLs will likely be higher too, even after adjusting for inflation.

On the other hand, cost-benefit analysis generally entails a discount for future dollars. One reason to discount future dollars is because, again, individuals in the future will probably be richer than individuals now. Due to the diminishing marginal utility of consumption, $1 worth of consumption will translate into less utility (i.e., less well-being) for individuals in the future than for individuals today—even after adjusting for inflation—because individuals in the future will have higher consumption (and thus lower marginal utility of consumption). Insofar as the goal of government policy is ultimately to maximize utility, policymakers should assign greater value to $1 of present consumption than to $1 of future consumption because $1 in the present is likely to bring more utility to whomever consumes it.

Cost-benefit analysis across federal agencies reflects both of these intuitions. Agencies typically use a concept known as the "income elasticity of the value of a statistical life" to reflect the percent change in the VSL for a percent change in overall income. The income elasticity of the VSL is conventionally denoted by the Greek letter theta (θ), sometimes with a subscript (θ_{VSL}). Agencies update their estimate of the VSL year-to-year by multiplying θ_{VSL} by the growth rate of real consumption or income per capita (g). Thus, the VSL used in regulatory cost-benefit analysis is higher—in real dollar terms—for future lives than for present lives. Agencies then discount future dollars at a rate conventionally denoted by the Greek letter delta (δ).

In theory, these two effects (the annual increase in the VSL at a rate of $\theta_{\mathrm{VSL}} \times g$ and the discounting of future dollars at a rate of δ) could offset precisely. In that case, a death averted in

---

[3] By "future lives," we refer to lives saved in the future. Many of these "future lives" are lives of people who are alive today and who would die sooner if the relevant regulation is not adopted.

[4] In theory, "willingness to pay" and "willingness to accept" may diverge. In the VSL context, however, Thomas Kniesner, Kip Viscusi, and James Ziliak find no statistically significant difference in VSL between workers who switch from more dangerous to less dangerous jobs (i.e., who are "paying"—through foregone wages—for greater safety) and workers who switch from less dangerous to more dangerous jobs (i.e., who are "accepting" a higher risk in exchange for higher wages). *See* Thomas J. Kniesner, W. Kip Viscusi & James P. Ziliak, *Willingness to Accept Equals Willingness to Pay for Labor Market Estimates of the Value of a Statistical Life*, 48 J. RISK & UNCERTAINTY 187, 190 (2014).

the future—after adjusting the VSL upward to reflect income growth and then discounting future dollars—would count the same in present value terms as a death averted today. In practice, however, federal agencies do not weigh present and future lives equally. Agencies have selected significantly lower values for the VSL growth rate ($\theta_{\text{VSL}} \times g$) than for the discount rate (δ), with the result that future lives count for much less than present lives.

To give some sense of the magnitude of the discrepancies: In 2013 guidance, the Department of Transportation estimated a value for θ_{VSL} equal to 1 and a value for g equal to 1.07%, while following a recommendation from the Office of Management and Budget to use alternative discount rates of 3% and 7%.[5] With those parameters, the value of a life saved 30 years from now would be—after adjusting upwards for income growth and then discounting at a 3% rate—approximately 43% less than the value of a life saved today. At a 7% discount rate, the present value of a life saved in 30 years would be 82% less than the present value of a life saved today.[6] At the Environmental Protection Agency, the haircut for future lives appears to be even more dramatic. Using the discount rate guidance currently in effect, the present value of a life saved in 30 years—according to the EPA—is 50% to 84% less than the present value of a life saved today.[7]

The approach to the valuation of future lives adopted by the EPA, the Department of Transportation, and other federal agencies has potentially profound anti-regulatory repercussions. Environmental, health, and safety regulations with delayed or durable effects run up against the obstacle that lives saved decades down the road don't count for much in cost-benefit analysis. That, in turn, results in weaker standards for air and water quality, vehicle and workplace safety, and other areas of policy where federal regulation has long-term life-or-death consequences.

---

[5] Memorandum from Polly Trottenberg, Under Sec'y for Policy & Robert S. Rivkin, Gen. Counsel, to Secretarial Officers & Modal Adm'rs, Guidance on Treatment of the Economic Value of a Statistical Life in U.S. Department of Transportation Analyses 2-3 (Mar. 1, 2013), https://www.transportation.gov/sites/dot.dev/files/docs/VSL%20Guidance%202013.pdf (hereinafter DOT 2013 Guidance); *see also infra* Section II.A.

[6] At a 3% discount rate: 1 – ((1 + 1 x 0.0107)/1.03)^30 ≈ 0.43. At a 7% discount rate: 1 – ((1 + 1 x 0.0107)/1.07)^30 ≈ 0.82.

[7] We discuss the components of the EPA's calculation in more detail in Section III.B, *infra*. The regulatory impact analysis supporting the agency's 2024 primary drinking water regulation indicated a VSL in 2024 of \$11,529,327 and a VSL in 2054 of \$13,907,201. *See* U.S. Env'tl Prot. Agency, *supra* note 2, at app. J tbl.J-1. Those figures imply a VSL growth rate of $\left(\frac{\$13{,}907{,}201}{\$11{,}529{,}327}\right)^{1/30} - 1 \approx 0.627\%$. The agency applied discount rates of 2%, 3%, and 7% in its analyses. *See* U.S. Env'tl Prot. Agency, Economic Analysis for the Final Per- and Polyfluoroalkyl Substances National Primary Drinking Water Regulation 2-3, 2-4 (Apr. 2024), https://www.epa.gov/system/files/documents/2024-04/pfas-npdwr_final-rule_ea.pdf. At a 2% discount rate: 1 – (\$13,907,201/(1.02^30))/\$11,529,327 ≈ 0.33. At a 3% discount rate: 1 – (\$13,907,201/(1.03^30))/\$11,529,327 ≈ 0.50. At a 7% discount rate: 1 – (\$13,907,201/(1.07^30))/\$11,529,327 ≈ 0.84. President Trump has since instructed agencies to use only the 3% and 7%—not 2%—discount rates. *See infra* Section III.A.5.

This Article presents a critical analysis of the federal government's approach to the valuation of future lives. After canvassing agency practice, we advance a case for urgent reform. Specifically, we argue that agencies should adopt a presumption that the growth rate of the VSL ($\theta_{\text{VSL}} \times g$) equals the discount rate (δ). Or put another way, agencies should assume—absent evidence strongly to the contrary—that the income elasticity of the VSL (θ_{VSL}) is equal to the ratio of the discount rate and the income growth rate (δ/g). When the VSL growth rate equals the discount rate, present lives and future lives will count the same in cost-benefit analysis. We show that economic theory and ethical intuitions both support the conclusion that the VSL growth rate is at least equal to the discount rate. We go on to demonstrate that setting the growth rate of the VSL equal to the discount rate would have large and potentially outcome-determinative consequences for cost-benefit analysis of major federal rules.

Our argument bears important implications for administrative law and policy. Executive Order 12866, promulgated by President Clinton in 1993, mandates that federal executive agencies quantify—"to the extent feasible"—the benefits and costs of significant regulatory actions based on the "best reasonably obtainable scientific, technical, economic, and other information."[8] Agencies are thus required to use the best estimates of the VSL growth rate and discount rate that are derivable from economic theory and data. And although Executive Order 12866 by its own terms cannot be enforced in court,[9] courts do consider the methodology, assumptions, and results of agency cost-benefit analyses for reasons other than enforcement of Executive Order 12866.[10] For instance, one ground on which agencies face challenges to their regulatory and deregulatory actions in federal court is with respect to whether the benefits justify the costs. If they do not, the agency's action may be "arbitrary" and "capricious."[11] We show that across both Democratic and Republican administrations, agencies have applied estimates for the value of future lives that are arbitrary, capricious, and—in some instances—riddled with arithmetic errors, leaving the resulting regulatory actions vulnerable to legal attack.

Beyond the implications for judicial review of administrative action, a coherent approach to the valuation of future lives will have potentially significant effects on the shape of the rules that agencies adopt in the first place. If agencies adopt the policies recommended by their own cost-benefit analyses, setting the VSL growth rate equal to the discount rate will have substantial impacts on the cars we drive, the air we breathe, and the risks we face in and beyond the workplace. Our argument also has important implications for equity across racial and ethnic groups. Because non-white and Latino groups constitute a minority of the U.S. population today but are projected

---

[8] Exec. Order 12866 of Sept. 30, 1993, §§ 1(b)(7), 6(a)(3)(B)(i)–(ii), 58 Fed. Reg. 51,735, 51,736, 51,741 (Oct. 4, 1993).
[9] *Id.* § 10, 58 Fed. Reg. at 51,744.
[10] *See* Caroline Cecot & W. Kip Viscusi, *Judicial Review of Agency Benefit-Cost Analysis*, 22 GEO. MASON L. REV. 575, 615–16 (2015). For example, the D.C. Circuit has held that "when an agency decides to rely on a cost-benefit analysis as part of its rulemaking, a serious flaw undermining that analysis can render the rule unreasonable." Nat'l Ass'n of Home Builders v. EPA, 682 F.3d 1032, 1040 (D.C. Cir. 2012).
[11] *See* 5 U.S.C. § 706(2)(A). For a compilation of 38 cases involving judicial consideration of agency cost-benefit analyses between 1985 and 2014, see Cecot & Viscusi, *supra* note 10, at 609–11 tbl.1.

to be a majority by mid-century,[12] current agency practice—which places greater weight on the present than the future—has the effect of valuing primarily white lives more than primarily non-white and Latino lives. Our argument for reform therefore implicates racial equity as well as intergenerational justice.

Our argument is especially relevant in the wake of the EPA's January 2026 announcement that it will no longer monetize the mortality benefits of reducing emissions of two major pollutants: fine particulate matter ($PM_{2.5}$) and ozone. EPA's stated rationale for dropping its VSL estimates for those two pollutants is "uncertainty," with the agency specifically highlighting uncertainty surrounding discount rates and income growth.[13] But if the EPA were to adopt our theoretically and empirically supported presumption—setting the growth rate of the VSL ($\theta_{VSL} \times g$) equal to the discount rate (δ)—then uncertainty surrounding discount rates and income growth would no longer be a barrier to monetizing mortality benefits. While discount rates and income growth are—concededly—difficult to estimate on their own, we show that these two values should move in tandem. As long as the income elasticity of the VSL (θ_{VSL}) is equal to the ratio of the discount rate and the income growth rate (δ/g), then EPA can plug in *any* value for δ and g, and the value of future lives will remain unchanged.

Our Article proceeds in three Parts. In Part I, we briefly describe how agencies value lives in the present. We explain the concept of the value of a statistical life (VSL), which agencies use when they must weigh the benefits of preventing premature deaths against the economic costs of regulations. Part II then shows how—as a matter of economic theory—agencies should apply the VSL when assessing the benefits of regulations that save lives in the future. We also review the empirical literature estimating key parameters relevant to the valuation of future lives, and we consider ethical arguments for and against our proposed equivalence presumption. Finally, Part III illustrates how four different federal agencies actually go about valuing lives that will be saved years into the future. As that Part makes clear, existing agency practice across the federal government deviates quite substantially from what the weight of economic theory and empirical evidence would recommend.

I. Valuing Present Lives

Federal agencies often must decide whether to adopt regulations that impose economic costs and yield life-saving benefits. The Office of Management and Budget's current guidance to federal agencies regarding regulatory analysis—Circular A-4—emphasizes the importance of

[12] *See* William H. Frey, *The US Will Become 'Minority White' in 2045, Census Projects: Youthful Minorities Are the Engine of Future Growth*, Brookings Inst.: The Avenue (Mar. 14, 2018), https://www.brookings.edu/blog/the-avenue/2018/03/14/the-us-will-become-minority-white-in-2045-census-projects.

[13] U.S. Envtl. Prot. Agency, Economic Impact Analysis for the New Source Performance Standards Review for Stationary Combustion Turbines: Final Rule 34 (Jan. 2026).

consistency in the monetization of mortality risk reduction.[14] For example, if the Department of Transportation adopts a safety standard (e.g., a new roof crush resistance requirement) that raises vehicle costs by $10 million per life saved, it should not reject—on cost-benefit grounds—a separate safety standard (e.g., a new side-impact protection requirement) that raises vehicle costs by $9 million per life saved. The amount that society is willing to pay to prevent a death in a motor vehicle should not depend on morally irrelevant factors such as whether the death would result from a crushed roof or a side collision.

Agencies seek to achieve consistency by selecting a standardized measure for the benefit of mortality reductions. The most commonly used measure in U.S. federal agency cost-benefit analyses is the "value of a statistical life" (VSL). The VSL is the value, expressed in dollars, of preventing one death. It is typically calculated on the basis of "local tradeoffs"—amounts that people are willing to pay or accept for small changes in their probability of death. As noted at the outset, federal agencies now use values in the range of $14 million for the VSL.[15]

Agencies generally base their VSL estimates on some combination of "revealed preference" and "stated preference" studies.[16] Revealed preference studies examine actual tradeoffs that individuals make; stated preference studies examine tradeoffs that participants in survey experiments say they would make. The most common type of revealed preference study in the VSL literature is a wage-risk study, which looks at wage differences across jobs with different mortality risk profiles. A VSL of $14 million derived from a wage-risk study implies that if there were two otherwise-identical jobs, one of which entailed a 1-in-1 million increase in the annual risk of death relative to the other, the riskier job would command $14 of extra annual pay.

The VSL is not the only measure that agencies use to monetize longevity benefits. Agencies sometimes employ a measure known as the "value of a statistical life *year*," or VSLY.[17] The agency estimates the number of additional years lived as a result of a regulatory intervention and then multiplies that by the VSLY. According to that approach, the value of an intervention that prevents the premature death of a person with 40 more years of remaining life expectancy is 40 times the value of an intervention that extends a life for one year. Agencies also occasionally use a measure known as a "quality adjusted life year," or QALY, which accounts for longevity gains and changes in the health-related quality of life.[18] A full exploration of VSLYs, QALYs, and related measures lies beyond the scope of this Article.[19] The key point for present purposes is that

---

[14] OFFICE OF MANAGEMENT AND BUDGET, CIRCULAR A-4, at 31 (Sept. 17, 2003), https://www.whitehouse.gov/wp-content/uploads/legacy_drupal_files/omb/circulars/A4/a-4.pdf; *see also* Exec. Order 14,192, 90 Fed. Reg. 9067 (Feb. 6, 2025) (reinstating the 2003 version of Circular A-4).
[15] *See supra* note 2 and accompanying text.
[16] *See, e.g.*, DOT 2013 Guidance, *supra* note 2, at 2-3.
[17] *See* Cass R. Sunstein, *Lives, Life-Years, and Willingness to Pay*, 104 COLUM. L. REV. 205, 252 app. (2004).
[18] *See id.*
[19] For further discussion of statistical life-years and QALYs, see John Bronsteen, Christopher Buccafusco & Jonathan Masur, *Well-Being Analysis vs. Cost-Benefit Analysis*, 62 DUKE L.J. 1603, 1676-82 (2013).

agencies have developed approaches that allow them to monetize the benefits of preventing premature deaths.

To this point, our discussion has concerned the methodology agencies use to value lives saved in the present. However, most significant regulations will not save lives immediately. Instead, they will prevent premature deaths many years into the future, when people who have breathed cleaner air or driven safer cars are not killed by heart disease or automobile accidents. Agencies must thus decide how much to value those future lives. That is the focus of our Article, which we take up in the next Part.

II. A Framework for Valuing Future Lives

Suppose the current value of a statistical life is $14 million, and a regulation will save a life 10 years from today. How much is that life worth in today's dollars? The reader might have the intuition that the answer should be $14 million. Perhaps lives should not be viewed as becoming cheaper or more valuable as time passes. But that is not how agencies have approached the question.

Agencies translate the value of present lives into a value of future lives through a two-step process. First, they estimate how much wealthier people will be in the future as the economy grows year by year. This greater wealth will translate into greater willingness to pay for life-saving precautions—or, in the language of wage-risk studies, greater demand for higher wages in exchange for risks to life and limb. If the people of 2027 are willing to pay $14 million to save a life in 2027, the people of 2037 will be willing to pay more than $14 million to save a life in 2037. Accordingly, agencies must increase the value of a statistical life by a small percentage each year to account for individuals' increasing wealth.

But the relevant question is how much the people of *2027* will be willing to pay to save a life in *2037*. The second step in valuing future lives is to translate the monetary value of a life in the future into a present-day monetary value. For a variety of reasons that we detail below, a dollar today is worth more than a dollar far in the future (even after accounting for inflation). Accordingly, when conducting cost-benefit analysis, agencies must *discount* future monetary values by a small percentage each year to identify their equivalent monetary values today.

The upshot is that in order to determine the value of a life saved in the future, agencies start with the current value of a statistical life, then increase that value by a small percentage each year to account for increasing wealth, and then decrease that value by a small percentage each year to account for the fact that a dollar in the future is less valuable than a dollar today. But as this capsule summary should indicate, this approach raises many more questions than it answers. In particular: by *how much* should agencies increase the value of life each year, and then by how much should they decrease the value of a dollar?

In this Part, we explain the theoretical and empirical framework for deciding how to value lives saved years in the future. Our focus is on the details of how agencies should select the rate at which they increase and then decrease the value of future lives. To preview our conclusion: there is a strong case, both theoretically and empirically, that the two rates—the rate at which VSL grows over time, and the rate at which dollars become less valuable over time—should be roughly equal, and thus a life saved in 2037 should be worth the same amount as a life saved in 2027. But as we will show in Part III, that is not at all what agencies have done.

The discussion that follows is sometimes technical, and it involves a number of mathematical symbols (including Greek letters) that are used to represent various economic concepts. We will introduce these terms and discuss the concepts they represent in detail below. In order to further simplify the discussion, we provide here a glossary of the symbols we will use. We invite the reader to refer back to this glossary whenever necessary.

Table 1. Glossary of Terms

Notation	Definition and Explanation
γ (gamma)	**Absolute value of the income elasticity of the marginal utility of consumption (or income).** γ reflects the change in the utility of each additional dollar as an individual gets richer. $\gamma = 1$ implies that when an individual's consumption doubles, the individual's marginal utility of consumption decreases by half. A higher (lower) value of γ implies that the marginal utility of consumption increases faster (slower) as an individual becomes richer.
θ (theta) or θ_{VSL}	**Income elasticity of the value of a statistical life.** θ_{VSL} reflects the change in the value of a statistical life as a function of income. $\theta_{\text{VSL}} = 1$ implies that a 1% increase in income translates to a 1% increase in the VSL. A higher (lower) value of θ_{VSL} implies that VSL rises faster (slower) as an individual becomes richer.
g	**Growth rate of real consumption or income per capita.** The adjective "real" means that g is adjusted for inflation.
ρ (rho)	**Rate of pure time preference.** ρ reflects the extent to which policymakers prefer to benefit people who are alive today over people who will be alive in the future purely because they care more about people who are alive today. $\rho = 0$ implies equal weight to present and future utility.
r	**Real risk-free rate of interest.** r is often estimated based on the inflation-adjusted yield on U.S. Treasury bonds.
δ (delta)	**Social discount rate.** δ reflects the extent to which policymakers value dollars today over dollars in the future. To a first approximation, $\delta = \gamma \times g + \rho$.

In Section II.A, we focus on γ, the absolute value of the income elasticity of the marginal utility of consumption or income. We start with γ because the elasticity of the marginal utility of consumption is essential to estimating both the income elasticity of the VSL (θ_{VSL}) and the social discount rate (δ). In Section II.B, we consider θ_{VSL}, the income elasticity of the value of a statistical life. Section II.C addresses g, the growth rate of real consumption or income per capita. In Section II.D, we analyze the construction of δ, the social discount rate. Section II.E compares our two central quantities of interest: δ, the social discount rate, and $\theta_{\text{VSL}} \times g$, the rate at which VSL increases. The central issue in valuing future lives is the race between the social discount rate (δ) and the VSL growth rate ($\theta_{\text{VSL}} \times g$), as we will explain in that Part. Finally, in Section II.F we discuss how our analysis can bring context and clarity to larger debates about discounting and—more generally—the cost-benefit analysis enterprise.

On to the main event.

A. Diminishing Marginal Utility of Consumption

At the center of the controversy over valuing future lives is the familiar concept of diminishing marginal utility of consumption.[20] Economists—and most non-economists who have thought about the question[21]—generally agree that the marginal utility of consumption is declining. In other words, the amount of happiness[22]—or the number of utils—that an individual derives from an extra dollar of consumption decreases over the number of dollars that the individual has to consume. For example, an extra dollar brings more utils to a minimum wage worker than to Elon Musk because Elon Musk already has more dollars than the minimum wage worker, so he derives fewer utils from each additional dollar.[23] The intuition behind this is that people will use their first dollars to buy the things they need most, their next dollars to buy the things they need next-most, and so forth. By the time they've gotten to their millionth dollar, whatever they're buying is no longer so essential. You can't get a lot of value out of that seventeenth streaming service.

In economics, the Greek letter gamma (γ) is often used to represent the relationship between the marginal utility of consumption and the level of consumption. To a first approximation, γ is the percent decrease in the marginal utility of consumption for a 1% increase in the level of consumption.[24] A value of zero for γ would imply (implausibly) that the marginal utility of

[20] In the absence of savings, the marginal utility of consumption is the same as the marginal utility of income. Savings upset the equivalence between those two concepts because a person who lives off savings can have high consumption with no income, while a very frugal person can have high income with low consumption. For our purposes, the consumption/income distinction will have little effect on the main analysis; we refer in text to consumption rather than income to align with the bulk of the relevant economic literature.

[21] *See* Joshua Greene & Jonathan Baron, *Intuitions About Declining Marginal Utility*, 14 J. BEHAVIORAL DECISIONMAKING 243 (2001) (documenting lay intuitions regarding declining marginal utility of consumption).

[22] Throughout this Article, we equate utility with happiness—a view defended in John Bronsteen, Christopher Buccafusco & Jonathan S. Masur, *Welfare as Happiness*, 98 GEO. L.J. 1583 (2010). However, nothing in our argument here depends on whether one views utility in terms of happiness or in terms of preference satisfaction.

[23] Everything we say about dollars could apply to euros, pounds, yen, renminbi, rupees, and so on.

[24] Economists often assume that the utility-of-consumption function takes an "isoelastic" or "power utility" form, whereby the utility of consumption ($u(c)$) can be expressed as:

consumption does not change with the level of consumption. The higher the value of γ, the more rapidly the marginal utility of consumption diminishes. A value of 1 for γ corresponds to a special case known as "log utility," where the utility of consumption is the natural logarithm of consumption ($ln(c)$) and the marginal utility of consumption is the inverse of consumption ($1/c$).

Estimates of γ come from several sources. One source is the body of research on the wage elasticity of labor supply. The wage elasticity of labor supply is the change in hours worked for a given change in the wage rate. The wage elasticity of the labor supply can be measured empirically by looking at changes in hours worked following changes in wages.[25]

When γ is high, the marginal utility of consumption declines very quickly. As wages rise, individuals rapidly become "sated with goods" (i.e., further increases in consumption bring them relatively little happiness).[26] In that case, we would expect the wage elasticity of labor supply to be low (or potentially negative): an increase in wages won't cause individuals to work much more (and might even cause them to work less). Low estimates of the wage elasticity of labor supply therefore suggest high values of γ (and, vice versa, high estimates of the wage elasticity of labor supply suggest low values of γ).[27]

$$u(c) = \begin{cases} \dfrac{c^{1-\gamma}-1}{1-\gamma}, & \gamma \geq 0, \gamma \neq 1 \\ \ln(c), & \gamma = 1 \end{cases}$$

The first derivative of the isoelastic utility function with respect to c is $c^{-\gamma}$, or $\frac{1}{c^{\gamma}}$, which is another way of expressing the marginal utility of consumption (i.e., the change in utility for a small change in the level of consumption).

For small percentage changes in the consumption level and values of γ in the low single digits, we can safely ignore the exponent formulation and simply think of γ as the absolute value of the percent change in the marginal utility of consumption for a percent change in the consumption level. For example, when γ equals 1, a 1% increase in the level of consumption corresponds to a decrease in the marginal utility of consumption of almost 1% (approximately 0.99%). When γ equals 2, a 1% increase in the level of consumption corresponds to a decrease in the marginal utility of consumption of almost 2% (approximately 1.97%). For higher values of γ and/or larger changes in consumption, the mental shortcut of thinking of γ as the percent decrease in the marginal utility of consumption for a percent increase in the consumption level no longer works as well, and γ is better understood with reference to the isoelastic utility formula.

[25] Specifically, the focus here is the *uncompensated* wage elasticity of labor supply, which reflects the income effect (the phenomenon that people choose to work less when their income rises because their marginal utility of consumption falls) and the substitution effect (the phenomenon that people choose to work more when their wage rises because they can buy more consumption for the same amount of labor effort). The compensated wage elasticity of labor supply, by contrast, ignores the income effect and isolates the substitution effect.

[26] Raj Chetty, *A New Method of Estimating Risk Aversion*, 96 AM. ECON. REV. 1821, 1821 (2006).

[27] Log utility ($\gamma = 1$) implies an uncompensated wage elasticity of zero. To see why, imagine that an individual's starting wage rate is w_0 and that she works h hours. Thus, her marginal utility of consumption is $1/(w_0 \times h)$. If she works an extra hour, she will earn w_0, allowing her to purchase additional consumption that brings her utility equal to her marginal utility of consumption multiplied by her additional earnings: $(1/(w_0 \times h)) \times w_0 = 1/h$. Because she will work only until the point that she is indifferent between an extra w_0 of consumption and an extra hour of leisure, her marginal utility of an extra hour of leisure when she works h hours is necessarily $1/h$ as well.

A second source for estimates of γ is the body of research on risk aversion.[28] Higher values of γ correspond to greater risk aversion. Intuitively, when γ is equal to zero (i.e., the marginal utility of consumption is constant), a rational agent should be risk neutral—all the agent is trying to do is to maximize the total amount of consumption across her entire life. When the marginal utility of consumption is declining, agents care more about consumption in future states in which they are poor than in future states in which they are rich. A rational, risk-averse agent may therefore turn down risky bets that the risk-neutral agent would accept.

To better understand the link between γ and risk aversion, think about whether you would prefer a future in which you are certain to receive $499,999 or one in which you toss a coin and have a 50% chance of receiving $1 million and a 50% chance of receiving nothing. Most people would prefer the former, because the downside in terms of your happiness or utility of going from $499,999 to zero is much, much greater than the upside of going from $499,999 to $1 million. The declining marginal utility of consumption is part of what generates this risk aversion. Someone with non-declining marginal utility of consumption would prefer the coin toss because the average outcome ($500,000) is higher than $499,999.

The relationship between the curvature of the utility function and risk aversion explains why γ is alternatively described as the "coefficient of relative risk aversion." One challenge in deriving an estimate of γ from research on risk aversion, though, is that individuals also tend to exhibit *loss aversion*: they reject bets with potential losses that, given plausible values of γ, expected utility theory suggests that they should accept.[29] Thus, estimating γ based on studies of risk-taking runs the risk—so to speak—that estimates of γ will be inflated because they reflect both risk aversion resulting from diminishing marginal utility of consumption and the separate phenomenon of loss aversion.

Now imagine that the individual receives a 1 percent raise (i.e., her wage rate rises to $1.01w_0$). If she continues to work h hours, her marginal utility of consumption will be $1/(1.01w_0 \times h)$. An extra hour of work will bring her consumption utility of $1/(1.01w_0 \times h) \times 1.01w_0 = 1/h$. Again, she will want to work until the marginal utility of leisure is $1/h$. Unless a change in consumption affects her marginal utility of leisure, log utility implies that an individual will continue to work the same number of hours when her wage rises (and likewise when her wage falls).

If γ is not equal to 1, then the uncompensated wage elasticity of labor supply will differ from zero. First, imagine that γ is less than 1. In that case, a wage raise will mean that the amount of extra consumption that an individual can afford will increase by more (proportionately) than her marginal utility of consumption will fall. Thus, the marginal utility of working an extra hour will rise, and the individual will choose to work more hours than before. Conversely, if γ is greater than 1, then a wage raise will mean that the amount of extra consumption that an individual can afford will increase by less, proportionately, than her marginal utility of consumption will fall. Thus, the wage raise will cause the individual to work less than she did before.

[28] *See* Ali Elminejad, Tomáš Havránek & Zuzana Irsova, *Relative Risk Aversion: A Meta-Analysis* (ZBW - Leibniz Information Centre for Economics, Working Paper, June 27, 2022), http://hdl.handle.net/10419/260586.

[29] *See* Matthew Rabin & Richard H. Thaler, *Anomalies: Risk Aversion*, 15 J. ECON. PERSP. 218 (2001).

Third, the body of research on the elasticity of intertemporal substitution in consumption can shed light on γ.[30] Given a choice between (option A) 1 unit of consumption at time 1 and (option B) $1 + r$ units of consumption at time 2 (where r, the real risk-free rate of interest, is greater than 0), an individual with constant marginal utility of consumption always will choose option B because $1 + r$ is greater than 1 and, by assumption, all dollars yield the same utility. As γ rises, shifting consumption from the present to the future in response to an increase in the interest rate becomes less attractive because our future (and richer) selves more rapidly become sated with goods.

The elasticity of intertemporal substitution in consumption is the percent change in consumption growth for a percent change in the real interest rate.[31] Lower values of γ correspond to higher elasticities of intertemporal substitution. For example, if γ is zero (i.e, the marginal utility of consumption is constant), then the elasticity of intertemporal substitution is infinity (i.e., any increase in the real interest rate above zero will cause an individual to save all of her income). Conversely, if γ is sufficiently high, then an increase in the real interest rate may cause an individual to save *less* of her income because she derives more utility from extra consumption today (when she is poorer) than in the future (when, as a result of the high interest rate, she will be richer).[32]

Fourth and finally, some researchers seek to estimate γ based on happiness survey data. When γ is zero (i.e., the marginal utility of consumption is constant), happiness increases proportionately with income. For higher values of γ, the effect of income on happiness is more muted, as individuals more quickly become sated with goods. Thus, the slope of self-reported happiness over income provides information on the value of γ, with a larger slope indicating a lower value of γ.

These various sources generate different estimates of γ. Based on a meta-analysis of labor supply studies, economist Raj Chetty derives a central estimate for γ of approximately 1, with an upper bound of approximately 2.[33] A recent meta-analysis of risk aversion studies reports a mean value for γ of 1 in the economics literature and 2 to 7 in the finance literature.[34] A meta-analysis of intertemporal substitution studies reports an implied γ of 3.[35] An analysis based on subjective

---

[30] For a meta-analysis, see Tomáš Havránek, *Measuring Intertemporal Substitution: The Importance of Method Choices and Selective Reporting*, 13 J. EUR. ECON. ASS'N 1180 (2015).
[31] With no pure time preference, the elasticity of intertemporal substitution is $1/\gamma$. We return to the question of pure time preferences in Section II.B.
[32] A potential challenge in estimating γ based on intertemporal substitution is that changes in the allocation of consumption may reflect changing time preferences as well as the interaction between interest rate changes and the diminishing marginal utility of consumption. For evidence that the rate of time preference remains relatively stable, at least in the aggregate, over the short to medium term, see Stephan Meier & Charles D. Sprenger, *Temporal Stability of Time Preferences*, 97 REV. ECON. & STAT. 273, 285–86 (2015).
[33] Chetty, *supra* note 26, at 1830.
[34] See Elminejad, Havránek & Irsova, supra note 28, at 5.
[35] *See* Havránek, *supra* note 30, at 1198.

wellbeing data from 50 countries across three decades yields values for γ ranging from 1.19 to 1.34.[36]

For our purposes, the value of γ matters less for its own sake than for what it tells us about two other parameters: the income elasticity of the VSL and the discount rate. It is worth emphasizing, though, that the various strategies for estimating γ tend to produce values of 1 or more.[37] Values of γ significantly below 1 are difficult to reconcile with real-world observations regarding labor supply, risk tolerance, savings behavior, and self-reported happiness.

B. The Income Elasticity of the Value of a Statistical Life

The diminishing marginal utility of consumption is the key force behind the income elasticity of the value of a statistical life. As noted at the outset, the income elasticity of the VSL (θ_{VSL}) is the percent change in the VSL for a small percent change in income.[38] In theory, the coefficient of relative risk aversion (γ) and the income elasticity of the VSL (θ_{VSL}) should be close to each other, though not necessarily identical.

To see the link between γ and θ_{VSL}, imagine that the value of preventing a premature death—in utility terms—does not vary with income or consumption. (For ease of exposition, we will assume here that income equals consumption.[39]) In other words, rich and poor people are both willing to sacrifice the same quantum of happiness (e.g., 1 util) to avoid a 1-in-1 million risk of premature death. However, due to the diminishing marginal utility of consumption, rich people derive less happiness from an additional dollar than poor people do. Rich people will therefore be willing to sacrifice more dollars to avoid a 1-in-1 million risk of death—not because they value their lives more, but because they value their money less.

If the assumptions in the previous paragraph apply, then γ and θ_{VSL} will be the same. Suppose that initially a person derives 1 util from \$1 of consumption and she is willing to sacrifice \$1 of consumption (or equivalently, 1 util) to avoid a 1-in-1 million risk of premature death. In that case, her VSL will be \$1 million. Suppose that γ is 1 and the person's income increases by 1%. If previously she derived 1 util from \$1 of consumption, now she derives 1 util from \$1.01 of consumption. That means that the utility cost to her of forgoing 1 unit of consumption has gone

---

[36] *See* R. Layard, G. Mayraz & S. Nickell, *The Marginal Utility of Income*, 92 J. PUB. ECON. 1846 (2008).

[37] A recent meta-analysis of estimates for γ in the United Kingdom yields a central estimate of 1.5. *See* Ben Groom & David Maddison Pr., *New Estimates of the Elasticity of Marginal Utility for the UK*, 72 ENVT'L & RES. ECON. 1155 (2019). Note also that strategies that generate estimates of γ significantly higher than 1 are vulnerable to methodological doubts about whether they are truly measuring γ or conflating γ with other phenomena (i.e., loss aversion, time preference, and the unitization of consumption in the utility function).

[38] To be more precise: With an isoelastic utility-of-consumption function and $\theta_{\text{VSL}} = \gamma$, the VSL should be proportionate to the inverse of the marginal utility of consumption. If the VSL is VSL_0 when consumption is c_0, then the VSL should be $VSL_0 * \frac{c_t}{c_o}^{\theta_{VSL}}$ when consumption is c_t. When $\theta_{\text{VSL}} = 1$ (as in the log utility case), the implication is that VSL is proportionate to consumption.

[39] *See supra* note 20.

down by 1%.[40] If spending on risk avoidance is 1% less costly (in forgone consumption utility terms), she will be willing to spend 1% more to avoid a 1-in-1 million risk of death. She will now be willing to pay up to $1.01 to avoid a 1-in-1 million risk of death (instead of $1). Thus her VSL will be $1.01 million—the 1% increase in her income corresponds to a 1% increase in her VSL. If γ is 2 and her income increases by 1%, the additional money is worth even less to her if used for consumption—it would take $1.02 to give her as much additional consumption utility as she previously derived from $1. Spending on risk avoidance is now approximately 2% less costly, meaning that she will be willing to pay $1.02 to avoid a 1-in-1 million risk of premature death. In other words, her VSL will be $1.02 million, with the 1% increase in her income corresponding to a 2% increase in her VSL.[41]

As Louis Kaplow notes, the relationship between happiness and income adds a further wrinkle to the relationship between γ and θ_{VSL}. If rich people are happier than poor people, rich people should be willing to sacrifice *more* utility to avoid a given risk of premature death than poor people. In that case, θ_{VSL} would be higher than γ.[42] For example, if γ is 1 and a person's income increases by 1%, then her VSL should rise by more than 1%. The first 1% increase in her VSL reflects the fact that her marginal utility of consumption has declined; the increase in the VSL on top of that reflects the fact that her utility from avoiding a given risk of premature death has increased because she is now happier with her life. While the magnitude of the latter effect may be small, Kaplow's logic suggests that γ sets a lower bound on θ_{VSL}.[43]

---

[40] *See* Louis Kaplow, *The Value of a Statistical Life and the Coefficient of Relative Risk Aversion*, 31 J. RISK & UNCERTAINTY 23, 24-25 (2005). Technically, when γ equals 1 and income rises by 1%, the utility cost of avoiding a 1-in-1 million risk of death will decline by $1 - 1/1.01 \approx 0.99\%$. The VSL, however, will increase by exactly 1%.

[41] *See id.* at 25. Technically, when γ equals 2 and income rises by 1%, the utility cost of avoiding a 1-in-1 million risk of death will decline by $1 - 1/1.01^2 \approx 1.97\%$, and the VSL will increase by $1.01^2 - 1 \approx 2.01\%$.

[42] *See id.*

[43] Mary Evans and Kerry Smith propose a reason why θ_{VSL} might be lower than γ. *See* Mary F. Evans & V. Kerry Smith, *Measuring How Risk Tradeoffs Adjust with Income*, 40 J. RISK & UNCERTAINTY 33 (2010). They posit that the risk of death (p) is an increasing function of the amount of time spent working (l). They define VSL as the hourly wage divided by the marginal increase in the risk of death associated with an additional hour of work: $VSL = w/p'(l)$. *See id.* at 40. They further posit that consumption and labor supply are complementary: "consumption makes work less painful." *Id*. Thus, when an individual's income exogenously increases, "[t]he compensation required for accepting the increase in risk associated with working an additional hour (VSL) is lower as a result." *Id.* at 41.

In our view, this argument conflates two quantities: (1) the compensation required for working an additional hour, and (2) the compensation required for accepting an increase in mortality risk. An exogenous increase in income plausibly reduces the disutility of working an additional hour because higher income workers can more easily pay for childcare, housekeeping, takeout meals, and other amenities that alleviate the burden of work. However, there is no obvious reason why an exogeneous increase in income would reduce the disutility from higher occupational mortality risk—why, for example, the pain and suffering associated with a workplace death would be negatively correlated with income. Thus, even if consumption and labor supply are complementary, consumption and the risk of death are not necessarily complementary, and an increase in income (which facilitates an increase in consumption) should not necessarily cause workers to be more accepting of on-the-job death.

Policymakers seeking to assign a numerical value to θ_{VSL} face two options. First, they can rely on theory to establish that θ_{VSL} equals or exceeds γ and then look to empirical estimates of γ. As noted above, most estimates place γ at 1 or higher, so this would suggest that θ_{VSL}, too, should be at least 1.

Second, policymakers can rely on a relatively small set of studies that estimate θ_{VSL} directly. A challenge in estimating θ_{VSL} directly is that wage-risk studies, which constitute the bulk of existing VSL estimates, rely on regression analysis with wages as the dependent variable and fatality risk among the independent variables. One cannot easily infer how VSL varies across income levels because VSL is itself estimated based on the relationship between wage income and fatality risk. If individuals in VSL studies also had significant sources of non-wage income, then wage-risk studies could potentially produce estimates of both VSL and θ_{VSL}. But for most of the working population, wage and income are highly collinear (i.e., most individuals receive the large bulk of their income from working rather than from savings).[44]

Researchers have therefore pursued several alternative strategies for estimating θ_{VSL}. One approach is longitudinal. Researchers can look at how VSL changes within the same country over time as income rises. Longitudinal studies are, in our view, the best source for deriving the value of θ_{VSL} that will be used for the valuation of future lives in cost-benefit analysis because we are ultimately interested in how VSL will evolve over time as per capita income rises.

The most comprehensive longitudinal analysis of the income elasticity of the VSL in the United States is a study by economists Dora Costa and Matthew Kahn, who combine U.S. Census microdata and Bureau of Labor Statistics fatality data to chart the evolution of VSL in the United States from 1940 to 1980.[45] Costa and Kahn estimate that the elasticity of the VSL with respect to per capita gross national product is approximately 1.5 to 1.7,[46] a range that they suggest should be interpreted as a "lower bound estimate" for several reasons.[47] For example, Costa and Kahn note that the share of jobs that are high-risk has declined over time; thus the marginal worker choosing between a high-risk and low-risk job is likely to be less risk-averse than before and is likely to demand a lower risk premium.[48] Costa and Kahn also note that workers' compensation has become more generous over time, which should likewise push the authors' θ_{VSL} estimate down.[49] (The

[44] *See* Hristos Doucouliagos, T.D. Stanley & W. Kip Viscusi, *Publication Selection and the Income Elasticity of the Value of a Statistical Life*, 33 J. HEALTH ECON. 67, 68 (2014).
[45] *See* Dora L. Costa & Matthew E. Kahn, *Changes in the Value of Life, 1940–1980*, 29 J. RISK & UNCERTAINTY 159 (2004).
[46] *See id.* at 172-73.
[47] *See id.* at 163.
[48] *See id*. To illustrate: Imagine that there are two types of jobs—high-risk and low-risk—and the only factors that affect a worker's choice of job are the wage and her risk aversion. When 50% of jobs are high-risk, the marginal worker will be at the 50th percentile of the risk-aversion distribution. When 10% of jobs are high-risk, the marginal worker will be at the 10th percentile of the risk-aversion distribution. Costa and Kahn write: "If the proportion of risk lovers is constant over time, then given that the share of jobs in risky occupations has shrunk, the marginal worker would require a lower risk premium." *Id.*
[49] *See id.*

reason is that individuals will be more willing to accept on-the-job hazards—and thus will demand a lower wage premium for occupational risk—when they and their family members are covered by more comprehensive insurance).[50] Costa and Kahn's estimate that θ_{VSL} lies substantially above 1 aligns with the one other comprehensive single-country longitudinal analysis of VSL—a study by James Hammitt, Jin-Tan Liu, and Jin-Long Liu using data from Taiwan for the 1982-1997 period and estimating θ_{VSL} to be in the range of 2 to 3.[51]

A second approach relies on a technique called "quantile regression," which examines the relationship between wages and fatality risks at different segments of the wage distribution using a single panel dataset. Thomas Kniesner, Kip Viscusi, and James Ziliak apply quantile regression methods to data from the 1993-2001 waves of the Panel Study of Income Dynamics. They arrive at an overall estimate for θ_{VSL} among U.S. workers of 1.44—a result roughly consistent with Costa and Kahn's. [52]

A third approach is to look at variation in VSL estimates across countries with different income levels. One challenge in estimating θ_{VSL} from cross-country studies is that factors other than income (e.g., cultural attitudes toward risk and death) may affect VSLs across the globe.

---

50 *Id.*

51 *See* James K. Hammitt, Jin-Tan Liu & Jin-Long Liu, *Is Survival a Luxury Good? Income Elasticity of the Value per Statistical Life*, 65 J. RISK & UNCERTAINTY 239 (2022). Hammitt and coauthors also conducted stated-preference risk-tradeoff studies in the Sichuan province of China in 2005 and 2016 and estimated VSL based on a restricted subsample of respondents whose answers exhibited consistency with theoretical conditions. They found a "dramatic increase" in average VSL between 2005 and 2016—consistent with an income elasticity of the VSL equal to 3.0. *See* James K. Hammitt, Fangli Geng, Xiaoqi Guo & Chris P. Nielsen, *Valuing Mortality Risk in China: Comparing Stated-Preference Estimates from 2005 and 2016*, 58 J. RISK & UNCERTAINTY 167, 169 (2019).

In their Taiwan study, Hammitt, Liu, and Liu also included estimates of the income elasticity of the VSL based on an instrumental variables (IV) strategy, which seeks to control for endogeneity between VSL and job choice. *See* Hammitt, Liu & Liu, *supra*, at 253. One requirement for IV estimation is that the instrument must be uncorrelated with any determinants of the dependent variable other than the endogenous regressor (the exclusion restriction). *See* JOSHUA D. ANGRIST & JÖRN-STEFFEN PISCHKE, MOSTLY HARMLESS ECONOMETRICS: AN EMPIRICIST'S COMPANION 117 (2009). However, Hammitt et al.'s instruments—including age, gender, marital status, and household size—are plausibly correlated with the dependent variable (VSL) through channels other than income. *See, e.g.*, Joseph E. Aldy & W. Kip Viscusi, *Age Differences in the Value of Statistical Life: Revealed Preference Evidence*, 1 REV. ENVT'L ECON. & POL'Y 241, 257 (2007) (noting inverted-U relationship between VSL and age). For example, gender may affect VSL both through its effect on income and its effect on longevity (i.e., males on average earn higher incomes, pushing VSL up, but also live shorter lives, potentially pushing VSL down). *See generally* Igor Fedotenkov & Pavel Derkachev, *Gender Longevity Gap and Socioeconomic Indicators in Developed Countries*, 45 INT'L J. SOC. ECON. 127, 128–29 (2020) (surveying potential causes of gender wage gaps and gender longevity gaps). Thus, it is not clear that the instruments chosen by Hammitt et al. can satisfy the exclusion requirement for IV estimation. Hammitt et al.'s IV-based estimates of the income elasticity of the VSL are substantially lower than their estimates based on changes in the VSL and income over time. *See* Hammitt, Liu & Liu, *supra*, at 228 (reporting IV-based estimates for the income elasticity of the VSL ranging from 0.65 to 0.91). We place less weight on these estimates in light of concerns regarding the validity of IV when the exclusion restriction cannot be met.

52 Thomas J. Kniesner, W. Kip Viscusi & James P. Ziliak, *Policy Relevant Heterogeneity in the Value of Statistical Life: New Evidence from Panel Data Quantile Regressions*, 40 J. RISK & UNCERTAINTY 15, 28 (2010).

Nonetheless, Ted Miller applies a cross-country approach and estimates a value of 1 for θ_{VSL}.[53] Kip Viscusi and Clayton Masterman estimate θ_{VSL} to be 1.1 based on a sample of wage-risk studies across different countries.[54] Other cross-country estimates range from as low as 0.2[55] to as high as 1.7 to 2.3.[56]

A fourth approach relies on stated preference VSL studies, which do not pose the same methodological hurdle as wage-risk studies (since VSL in stated preference studies is not estimated from regressions with wages as the dependent variable). This approach tends to generate lower estimates of θ_{VSL}: for example, James Hammitt and Kevin Haninger estimate elasticities in the range of 0.1 to 0.3.[57] One concern about estimates derived from these studies is that the stated preferences of survey experiment participants may not reflect their real-world behavior. A further concern is that the income reported by experiment participants may not accurately reflect consumption. For example, individuals who leave the labor force to attend graduate school or because of retirement may experience dramatic declines in income, but we would expect them to be able to smooth consumption (at least to some extent) through borrowing and saving, and thus we would not expect their VSL to change dramatically. A weak relationship between VSL and income across a single individual's lifecycle would not, therefore, imply a similarly weak relationship between VSL and population-wide per-capita income or consumption across time.

Finally, some researchers have collected VSL estimates from different studies and examined the relationship between VSL and the average income of subjects in each study. Several of these meta-analysis estimates of θ_{VSL} fall below 1.[58] One serious concern about this approach, though, is that the average income of subjects in a study may not be informative regarding the income to which the study's VSL estimate applies.

To illustrate: Imagine two VSL studies—Study A and Study B. Study A involves one nurse who earns $50,000 and has zero risk of on-the-job death plus one logger who earns $60,000 and

---

[53] *See* Ted R. Miller, *Variations Between Countries in Values of Statistical Life*, 34 J. TRANSP. ECON. & POL'Y 169, 182 (2000).
[54] *See* W. Kip Viscusi & Clayton J. Masterman, *Income Elasticities and Global Values of a Statistical Life*, 8 J. Benefit Cost Anal. 226, 242 tbl.5 (2017).
[55] *See* Chris Doucouliagos, T.D. Stanley & Margaret Giles, *Are Estimates of the Value of a Statistical Life Exaggerated?*, 31 J. HEALTH ECON. 197, 202 (2012).
[56] *See* Bradley J. Bowland & John C. Beghin, *Robust Estimates of Value of a Statistical Life for Developing Economies*, 23 J. POL'Y MODELING 385, 393 tbl.2 (2001).
[57] *See* James K. Hammitt & Kevin Haninger, *Valuing Fatal Risks to Children and Adults: Effects of Disease, Latency, and Risk Aversion*, 40 J. RISK & UNCERTAINTY 57 (2010); *see also* James K. Hammitt & Lisa A. Robinson, *The Income Elasticity of the Value per Statistical Life: Transferring Estimates Between High and Low Income Populations*, 2 J. BENEFIT-COST ANAL. no. 1, art. 1, at 10 (2011) (collecting estimates).
[58] *See* Clayton J. Masterman & W. Kip Viscusi, *The Income Elasticity of Global Values of a Statistical Life: Stated Preference Evidence*, 9 J. BENEFIT COST ANAL. 407, 410 tbl.1 (2018) (listing 12 meta-analysis studies that estimate θ_{VSL}, six of which result in ranges for θ_{VSL} below 1). A meta-analysis of meta-analyses by Chris Doucouliagos, T.D. Stanley, and Kip Viscusi arrives at an overall θ_{VSL} of approximately 0.6 after attempting to correct for publication bias. *See* Hristos Doucouliagos, T.D. Stanley & W. Kip Viscusi, *Publication Selection and the Income Elasticity of the Value of a Statistical Life*, 33 J. HEALTH ECON. 67, 68 (2014).

has a 1-in-1000 risk of on-the-job death. The VSL estimated from Study A would be $10 million: an increase in mortality risk of 1-in-1000 correlates to an increase in income of $10,000. The average income of subjects in Study A would be $55,000 (the average of the nurse's $50,000 and the logger's $60,000). Study B involves four nurses who earn $50,000 and have zero risk of on-the-job death plus one logger who earns $60,000 and, as before, has a 1-in-1000 risk of death. Again, the VSL would be $10 million, but the average income of subjects in Study B would be $52,000, simply because Study B includes more nurses. A meta-analysis of Study A and Study B would conclude that θ_{VSL} equals zero: an increase in average income from $52,000 to $55,000 doesn't change the VSL. But the conclusion would be spurious. Nothing in Study A or Study B tells us how VSL varies across the income spectrum. The apparent result that θ_{VSL} equals zero simply reflects the fact that the distribution of incomes across the two studies differs.

For the meta-analysis approach to generate reliable estimates of θ_{VSL}, we would need to make very strong assumptions about the consistency of the distribution of incomes across the disparate VSL studies included in the meta-analysis. As far as we know, none of the meta-analysis studies have sought to specify—much less substantiate—those assumptions. For this reason, we put little stock in estimates of θ_{VSL} that are based purely on meta-analyses of different VSL studies.

Moreover, our quantity of interest is growth in VSL over time. This is precisely what longitudinal studies measure: the change in VSL over time among a population of people. Of course, the fact that VSL increased by some percentage per year in the past does not mean that it will increase by the same percentage per year in the future. But it is an excellent guide to answering that question. By contrast, cross-sectional studies ask how VSL differs among different populations of people at a single moment in time. This can be a useful indication of how VSL will change over time—perhaps the people who are earning $65,000 per year a few decades from now will resemble the people who earn $65,000 per year today—but translating a cross-sectional estimate of θ_{VSL} into a projection of future VSL growth requires strong assumptions even apart from the fact that our cross-sectional estimates of θ_{VSL} are highly uncertain and variable. Accordingly, we think federal agencies, when updating VSL over time, should place greater weight on longitudinal studies, which directly estimate change over time and are less vulnerable to the methodological concerns limned above regarding cross-sectional studies.

Summing up: The empirical literature on θ_{VSL} yields wide-ranging estimates, but the most relevant evidence—particularly longitudinal evidence from the United States—suggests that VSL increases significantly faster than per capita income (i.e., $\theta_{\mathrm{VSL}} > 1$).[59] Some estimates based on other sources place θ_{VSL} below 1, but those approaches are also subject to more significant methodological doubts. Given the sound theoretical basis for concluding that θ_{VSL} equals or exceeds γ, along with the prevailing view that γ is at least 1, we think the best estimate of θ_{VSL} is also at least 1.

---

[59] *See* Costa & Kahn, *supra* note 46.

C. The Growth Rate

The second key input into the calculation of future VSLs is the real growth rate (g). Thus far we have focused on the marginal utility of *consumption*. Economic growth projections typically focus on per capita gross domestic product: the sum of consumption, investment, government expenditures, and net exports, all divided by population. If consumption as a percentage of GDP remained constant over time, then whether we use the growth rate of per capita consumption or the growth rate of per capita GDP for g in the future-VSL formula would not matter. In fact, personal consumption as a percentage of U.S. GDP has fluctuated in the postwar period, hitting a low of 58.6% in the first quarter of 1967 and reaching a high of 68.8% in the third quarter of 2009.[60] But presumably this trend cannot continue forever or else eventually there would be no government spending and no investment. The growth rate of real GDP per capita is thus likely to provide a serviceable guide to g, even if not a perfect match.

The Congressional Budget Office publishes projections for the growth rate of U.S. real GDP per capita for the next 30 years. CBO expects that the growth rate of real GDP per capita will be 1.3% over the 2025-2055 period. This is slightly below the growth rate of real GDP per capita over the previous 30 years (1.7%).[61] While CBO projections are not gospel, they generally provide a good approximation of the consensus view regarding future economic growth.[62]

D. The Social Discount Rate

The last ingredient in the formula for the valuation of future lives is the social discount rate (δ). There is widespread conceptual disagreement among scholars about the social discount rate, with different schools of thought arguing that the discount rate should be based on different underlying concepts or quantities. In this Part, we focus on the three major schools of thought. Despite the heterogeneity of existing views, we argue that there are strong theoretical and empirical reasons to set the social discount rate equal or very close to the product of γ and g—that is, equal to the rate of increase in the value of a statistical life.

We begin with classical utilitarian theory, which argues directly for a value of δ that equals the product of γ and g. The argument for setting δ equal to γg—which dates at least as far back as the work of early 20th century British polymath F.P. Ramsey[63]—arises from the diminishing marginal utility of consumption. That is, \$1 of consumption (or for Ramsey, £1) will bring less utility to people in the future because the marginal utility of consumption is declining and people

---

[60] *Shares of Gross Domestic Product: Personal Consumption Expenditures*, FEDERAL RESERVE BANK OF ST. LOUIS: FRED ECONOMIC DATA, https://fred.stlouisfed.org/series/DPCERE1Q156NBEA (last updated Feb. 16, 2026). The personal consumption share hit 68.8% again in the first quarter of 2011. *See id.*

[61] *See* CONGRESSIONAL BUDGET OFFICE, THE 2025 LONG-TERM BUDGET OUTLOOK 32 tbl.3-1 (Mar. 2025), https://www.cbo.gov/file-download/download/private/164631.

[62] *See, e.g.*, *CBO Economic Forecasts Are Similar to the Consensus, Unlike OMB*, COMMITTEE FOR A RESPONSIBLE FEDERAL BUDGET (June 1, 2018), https://www.crfb.org/blogs/cbo-economic-forecasts-are-similar-consensus-unlike-omb.

[63] *See* F.P. Ramsey, *A Mathematical Theory of Saving*, 38 ECON. J. 543 (1928).

in the future will be richer than people today. If it takes 1 unit of consumption to produce 1 util today, then it will take $1 + \gamma g$ units of consumption to produce 1 util when the level of consumption has increased by g. Thus, a project that requires us to sacrifice 1 unit of consumption today will increase total utility across time if and only if it generates more than $(1 + \gamma g)^t$ in benefits at time t. Setting δ equal to γg ensures that the social planner pursues only projects that increase total utility across time.

Ramsey acknowledged the possibility that policymakers might choose to discount future benefits by more than γg. He introduced the idea of a pure time-preference rate, denoted by the Greek letter rho (ρ), which would reflect the discounting of future *utilities* (and not just future consumption).[64] One can conceive of ρ as a type of selfishness (or at least self-interestedness) favoring people who are living now, as opposed to people who will live at some point in the future. The social discount rate then can be expressed as $\delta = \gamma g + \rho$, with the term γg reflecting discounting due to the diminishing marginal utility of consumption and ρ reflecting the additional role of pure time preference.

Ramsey believed that ρ should be set equal to zero, but a second school of economists has advocated positive values of ρ based on the argument that a ρ of zero imposes excessive demands on the present. The Nobel laureate economist Kenneth Arrow articulated one prominent version of the argument for a positive rate of time preference. Arrow's argument is based on a thought experiment initially offered by another Nobel laureate economist, Tjalling Koopmans. Arrow wrote:

> Imagine initially that output consists of a constant stream of completely perishable goods. There can be no investment by definition. Now imagine that an investment opportunity occurs, available only to the first generation. For each unit sacrificed by them, a perpetual stream of α per unit time is generated. If there were no time preference, what would the optimal solution be? Each unit sacrificed would yield a finite utility loss to the first generation, but to compensate there would be a gain, however small, to each of an infinity of generations. Thus any sacrifice by the first generation is good. Strictly speaking, we cannot say that the first generation should sacrifice everything, if marginal utility approaches infinity as consumption approaches zero. But we can say that given any investment, short of the entire income, a still greater investment would be preferred.[65]

Arrow believed that Koopmans's thought experiment established the necessity of setting ρ above zero. He considered it "unacceptable" to say that the present generation should sacrifice all but an infinitesimal amount of consumption in order to facilitate the happiness of future generations.[66] While acknowledging the "very artificial" structure of the economy imagined in the

[64] *See id.* at 554.
[65] See Kenneth J. Arrow, *Inter-Generational Equity and the Rate of Discount in Long-Term Social Investment*, *in* CONTEMPORARY ECONOMIC ISSUES 89, 96 (M.R. Sertel ed., 1999).
[66] *See id*.

thought experiment, Arrow believed that "this is not a valid objection." As he put it: "If a preference ordering is suitable and meaningful, then it should explain behaviour under any physically possible opportunity set, not merely 'realistic' ones."[67] Setting ρ above zero would serve to cap the amount that the present generation would be required to sacrifice in the thought experiment because the stream of gains to future generations would converge to a finite number (α/ρ).[68]

On reflection, it is not so clear why demanding a large sacrifice from the first generation should be "unacceptable," at least where that sacrifice facilitates an infinite stream of consumption utility for future generations. Classical utilitarianism often leads to the conclusion that the interests of a small group of people should yield to the greater good. Indeed, the very premise of cost-benefit analysis for life-saving interventions is that some people's interest in avoiding premature death can be weighed against—and sometimes can be subordinated to—society's interest in higher consumption levels. In Koopmans's thought experiment, the sacrifice demanded of the first generation—penury—seems even less severe than premature death. Thus, if we are comfortable with the notion that, under certain circumstances, a small number of people should be allowed to die so that a larger number can be richer, we should be willing to countenance the possibility that one generation should be allowed to live in poverty so that an infinite number of future generations can be richer.

Another response to Arrow's argument is that Arrow's own approach—setting ρ above zero—also leads to reductio ad absurdum conclusions in artificial thought experiments. Ricky Revesz has offered one counterexample that illustrates this point.[69] Revesz imagines an economy with two generations—one that will live for 50 years, and a second that will live for the next 50 years. A social planner can allocate 100 units of consumption across the two generations. There is no probability of extinction, no possibility of economic growth, and no population growth. Many people have the strong intuition that the social planner should allocate 50 units of consumption to the first generation and 50 units to the second generation. After all, nothing distinguishes the first generation from the second other than chronological order.[70] But setting ρ above zero would suggest that the first generation should receive more than half of the 100 units simply because it comes first. And while Revesz's thought experiment is, like Arrow's, highly artificial, this is—by Arrow's own logic—not a valid objection.

We are thus unpersuaded by Arrow's argument for setting ρ above zero (and thus setting δ above γg). Arrow is correct that $\delta = \gamma g$ may lead to large demands on the present generation under

---

[67] *Id.* at 97.
[68] Even when ρ is zero, extinction risk would cap the amount that the present generation would be required to sacrifice. Extinction risk is distinct from pure time preference because even a Ramsey-style utilitarian would—in calculating the amount of utility that a future generation is likely to experience—account for the possibility that the future generation might never exist.
[69] Richard L. Revesz, *Environmental Regulation, Cost-Benefit Analysis, and the Discounting of Human Lives*, 99 COLUM. L .REV. 941, 998 (1999).
[70] *See id.*

certain circumstances. At the same time, Arrow's alternative (a positive value for ρ) leads to similarly unattractive theoretical results—results that cannot be rationalized by the standard utilitarian appeal to the greatest good for the greatest number.

Rejecting Arrow's a priori argument for a positive ρ does not, however, address an alternative rationale for setting δ apart from γg—the argument from opportunity costs. This argument from opportunity costs, which represents a third school of thought regarding social discounting, advocates setting δ equal to the real rate of return on investment (r) regardless of γg. Dexter Samida and David Weisbach offer one smart version of this argument. They point out that any value of δ other than r is Pareto dominated: there is an alternative that could make at least one affected person—or affected generation—better off without making anyone worse off.[71] The reason is that the government always can invest at a rate of r, and thus it always has the opportunity to transform one dollar today into $(1 + r)^t$ dollars in t periods. Therefore, the government should not pursue a project that would cost one dollar today if the project will yield benefits at time t of less than $(1 + r)^t$ dollars. Rather than pursuing such a project, the government could make a future generation better off without making the present generation worse off if it chose instead to invest the dollar and earn a rate of return r, in which case the future generation would have $(1 + r)^t$ dollars instead of benefits that it values at less than that. Likewise, the government should always pursue a project that costs one dollar today if the project will yield benefits at time t of more than $(1 + r)^t$ dollars. The government could make a future generation better off without making the present generation worse off if it chose to borrow the dollar at an interest rate of r. At time t, the debt will have grown to $(1 + r)^t$. If the benefits from the project will be greater than $(1 + r)^t$, the future generation will be able to pay off the debt and still be better off—all with no harm to the present.[72]

As Samida and Weisbach stress, their argument for setting δ equal to r does not imply that the *utility* of future generations should be discounted. That is, they do not necessarily advocate for a pure time preference (ρ) greater than zero. Their argument simply implies that dollars in the future should be discounted at a rate of r to reflect the possibility of investing those dollars today. Whether this results in a value of δ greater than or less than γg is an empirical question.

One potential objection to this argument is that the decisionmakers in charge of setting δ for cost-benefit analysis purposes do not necessarily have control over large-scale saving and

[71] Dexter Samida & David A. Weisbach, *Paretian Intergenerational Discounting*, 74 U. CHI. L. REV. 145 (2007).
[72] A similar argument would potentially justify setting δ equal to r when $r < \gamma g$. Suppose r is 2% and γg is 3%. Suppose, moreover, that a project costs \$1 and would yield benefits of \$1.025 in one year (i.e., a 2.5% rate of return). The government could make people alive today and people alive in one year both better off by borrowing, say, \$1.003, and spending \$1 on the project. Consumption of people alive today would rise by \$0.003, leaving them better off. People alive in one year would owe a debt of \$1.003 × 1.02 ≈ \$1.023, which is less than their \$1.025 in project benefits, leaving them better off too. To be sure, when $r < \gamma g$, a utilitarian government would want to borrow more for present consumption, but as discussed in text, the Samida-Weisbach argument assumes that the overall amount of intergenerational wealth transfers is the result of a political equilibrium that the regulatory policymaker cannot unilaterally override.

borrowing decisions.[73] To illustrate: Imagine that $r > \gamma g$. In that case, instead of spending 1 today for benefits of $(1 + \gamma g)^t$ in t years, the social planner could make the generation alive in t years better off by investing 1 today and earning a rate of return r, thus yielding $(1 + r)^t$ in t years. However, the EPA Administrator and the Secretary of Transportation do not have the authority to—for example—forgo a costly clean air or vehicle safety standard and use the savings to buy government or corporate bonds. Some of the savings from forgoing the costly regulation will be consumed by the present generation rather than benefitting future generations. And that consumption might be less valuable (in utils) to the current generation than investment would be to the future generation. So while in theory the generation alive at time t could be made better off if the present generation forgoes the regulation and invests the savings, there is no assurance that the generation alive at time t actually *will* be better off as a result. In that case, an equal regard for the utility of future generations requires the government to pursue the project as long as it yields a return greater than γg.

Advocates for setting δ equal to r are not without a rebuttal. They concede that executive agencies (e.g., EPA) choose the discount rate for regulatory analysis and that Congress controls government saving and borrowing. By choosing a lower value of δ, EPA is engaged in *implicit* intergenerational saving via regulation (i.e., it is imposing present costs for future benefits). If Congress is unhappy with that amount of intergenerational saving, it can respond by borrowing more. For example, imagine that EPA adopts a regulation that reduces consumption for the present generation by 1 unit and yields benefits for the next generation of $(1 + \delta)^t$, where δ is whatever discount rate EPA has chosen. If Congress is unhappy with EPA's decision to deprive the present generation of 1 unit of consumption, Congress can borrow 1 unit at an interest rate of r. The net effect will be to maintain the present generation's consumption and to reduce the future generation's consumption at time t by $(1 + r)^t - (1 + \delta)^t$, or in other words, by the amount of government debt that the future generation must repay minus the yield on EPA's regulatory investment. Accounting for Congress's response, EPA will have made the future generation worse off.

Whether Congress will in fact respond to regulatory investments by adjusting the amount of government borrowing is, again, an empirical question.[74] Samida and Weisbach suggest that we can assume the political process produces some "equilibrium amount" of intergenerational transfers, which a single agency cannot unilaterally alter.[75] Alternatively, one might argue—in a more institutionalist vein—that whether or not Congress actually does adjust federal borrowing in response to regulatory investments, agencies still should proceed on the assumption that δ equals

[73] This argument runs parallel to arguments made by Lee Fennell and Richard McAdams, and John Bronsteen, Christopher Buccafusco, and Jonathan Masur, regarding the choice between a project that increases welfare but decreases wealth, on the one hand, or declining the project and redistributing wealth, on the other. *See* Lee Ann Fennell & Richard McAdams, *The Distributive Deficit in Law and Economics*, 100 MINN. L. REV. 1051 (2016); Bronsteen et al., *supra* note 19, at 1664-69.

[74] We know of no study that substantiates the claim that Congress adjusts borrowing in response to regulations with short-term costs and long-term benefits—but no study that disproves the claim either.

[75] *See* Samida & Weisbach, *supra* note 71, at 162.

r. The argument could be framed in terms of institutional competence: for example, officials at EPA are the experts on environmental regulation, and officials at NHTSA are the experts on motor vehicle safety, but neither EPA nor NHTSA has any particular advantage over members of Congress in resolving questions of intergenerational equity.[76] Or the argument could be framed in terms of democratic legitimacy: although a representative democracy may allocate highly technical questions to expert agencies, the larger (and less technical) question of how we ought to value present versus future generations ought to be resolved by elected officials.[77] As is already apparent, the question of whether agencies should set δ equal to γg or equal to r implicates difficult issues of institutional design and political theory. As we shall soon see, though, the answer to the question may not matter hugely if r roughly equals γg.

In any event, advocates for setting δ equal to r still face the separate question of deciding which rate of return to use. A logical answer, at least for purposes of regulatory cost-benefit analysis across U.S. federal agencies, is that the social discount rate should be set equal to the real interest rate on Treasury bonds. First, the interest rate on Treasury bonds is the rate at which the government actually borrows and "saves" (with "saving" taking the form of repaying already-issued debt). Second, the argument for applying a discount rate higher than the Treasury rate applies only when the relevant investment is risky (provided that the risk is not diversifiable). Investors rationally demand higher interest rates for riskier assets because of the diminishing marginal utility of consumption. For example, with diminishing marginal utility of consumption, an investment of $100,000 that pays back $102,000 with 100% certainty is preferable to an investment of $100,000 that will pay back $204,000 with 50% probability and $0 with 50% probability. The reason is, again, that $204,000 of consumption yields less than twice as much utility as $102,000, so expected utility is higher for the investment that yields $102,000 with certainty even though the expected dollar value of both investments is the same. Thus, if the risk-free rate of return is 2%, an investor who experiences diminishing marginal utility of consumption should demand a rate of return north of 2% for the risky bet.

By contrast, there is no obvious reason why society should demand a rate of return higher than the risk-free rate in order to invest in life-saving projects. A project that saves 102,000 lives with 100% certainty is not necessarily preferable to a project that saves 204,000 lives with 50% probability—at least when extinction of the human species is not an issue. The 204,000th life is not necessarily less valuable than the 102,000th. The logic of diminishing marginal utility of consumption, which justifies the use of higher interest rates for risky financial bets, does not translate straightforwardly into an argument for higher interest rates for projects with life-saving benefits.

---

[76] For an analogous argument that executive agencies should defer to Congress' decisions about the distribution of resources across income groups (rather than across generations), see David A. Weisbach, *Distributionally Weighted Cost-Benefit Analysis: Welfare Economics Meets Organizational Design*, 7 J. LEGAL ANALYSIS 151, 176 (2015).

[77] *Cf.* Daniel J. Hemel, *Wealth, Schmealth, Welfare, and Schmelfare*, 59 WAKE FOREST L. REV. 1103, 1138–42 (2024) (interrogating the idea that "distributional decision-making ought to lie with the legislature").

A more subtle version of the argument for choosing a value for *r* different from the risk-free rate relies on the capital asset pricing model (CAPM). The CAPM equation states that the rate of return on any asset should equal the risk-free rate of return plus the product of the asset's "beta" and the market risk premium, where "beta" (β) represents the correlation between the asset's expected return and the expected return of the market.[78] The logic underlying the CAPM equation is that investors can effectively eliminate any asset-specific (or "idiosyncratic") risk by holding a well-diversified portfolio, but diversification does not eliminate risks affecting the market as a whole (e.g., the risk of a recession). Thus, a well-diversified investor should demand a risk premium only for risks that are positively correlated with the market (i.e., for which beta is greater than zero).[79] We can, moreover, think of the government as an exceptionally well diversified investor: through its taxing power, it holds a potential claim on a portion of returns from nearly all assets in the economy. Thus, the logic of CAPM would suggest that the federal government should apply a discount rate higher than the risk-free rate when the return on a project is positively correlated with the broader market.[80]

The correlation between the return on a lifesaving regulation and the market return is likely to vary from regulation to regulation, and often it will be theoretically ambiguous whether a regulation's beta is positive (in which case the discount rate should exceed the risk-free interest rate) or negative (in which case the discount rate should be lower than the risk-free rate). Consider, for example, a regulation that improves road safety. Traffic—and in particular, the movement of large commercial trucks, which pose unique risks to other road users—increases when economic activity is robust.[81] We therefore might think that a regulation that improves road safety will save more lives when economic growth is strong (and fewer lives when economic growth is weak). In that case, beta would be positive, and the discount rate should exceed the risk-free rate. At the same time, economic growth may enable people to purchase newer and safer vehicles.[82] If road safety and vehicle safety are substitutes, then a regulation that improves road safety will save more

---

[78] For example, beta equals 1 for an asset that moves up and down with the market; beta equals 2 for an asset that rises 2% when the market rises 1% and that falls 2% when the market falls 1%; and so on.

[79] For an overview, see generally Eugene F. Fama & Kenneth H. French, *The Capital Asset Pricing Model: Theory and Evidence*, 18 J. ECON. PERSP. 25 (2004).

[80] *See* Mark Harrison, *Valuing the Future: The Social Discount Rate in Cost-Benefit Analysis* 39-48 (Productivity Comm'n of the Austl. Gov't, Visiting Researcher Paper, Apr. 2020). To elaborate on the CAPM logic: Imagine that a regulatory project yields a return in every state of the world equal to the market return minus some small amount (epsilon). If we calculated the return on the regulatory project using the risk-free interest rate as the discount rate, then the project would appear to pass the cost-benefit test (because the expected rate of return on market-correlated assets exceeds the risk-free rate). But clearly, society would be better off if the government just bought another unit of the market-correlated asset rather than investing in a project that always performs epsilon worse than the market-correlated asset. Thus, if we evaluate regulatory projects against the baseline of the government investing the money instead, we should choose a baseline with a similar risk profile as the regulatory project.

[81] *See* Monica M. He, *Driving through the Great Recession: Why Does Motor Vehicle Fatality Decrease When the Economy Slows Down?*, 155 SOC. SCI. & MED. 1 (2016).

[82] *See* David Bishai, Asma Quresh, Prashant James & Abdul Ghaffar, *National Road Casualties and Economic Development*, 15 HEALTH ECON. 65, 80 (2006). Another possibility is that economic growth leads to improvements in post-injury ambulance transport and trauma care. *See id.*

lives when economic growth is weak (and fewer lives when economic growth is strong). In that case, beta would be negative, and the discount rate should be lower than the risk-free rate. The correlation between regulatory returns and the market return is similarly ambiguous in other contexts—for example, regulations that reduce emissions of greenhouse gases.[83] Thus, if agencies set the discount rate (δ) equal to the rate of return (r), we think they should choose a discount rate equal to the real risk-free rate of return absent a reason to think that beta is positive or negative.

As alluded to above, the choice between setting δ equal to γg and setting δ equal to r may not matter much when r is interpreted as the risk-free rate of return on government debt. Recall that estimates for γ from wage elasticity and subjective wellbeing studies fell in the 1 to 2 range, while CBO projects a long-term income growth rate of 1.3%. Using those parameters, if $\delta = \gamma g$, then δ falls somewhere between 1.3% and 2.6%. Meanwhile, as of this writing, the interest rate on inflation-indexed Treasury securities ranged from approximately 1.36% for a five-year term to 2.67% for a thirty-year term.[84] Thus, the range of plausible values for γg and the range of values for r roughly match.

E. The Race Between δ and $\theta_{VSL} \times g$

To sum up so far: We have argued that the diminishing marginal utility of consumption implies that the VSL will rise as consumption levels rise. Theory strongly suggests that the income elasticity of the VSL (θ_{VSL}) should equal or exceed the coefficient of relative risk aversion (γ). Estimates of γ arrive at a wide range of values, but the most reliable studies seem to cluster in the range of 1 to 2.[85] Estimates of θ_{VSL} are similarly scattered, but the most comprehensive longitudinal analysis based on U.S. data arrives at a lower-bound estimate of 1.5 to 1.7.[86] Although there is considerable variance across estimates, the empirical evidence is broadly consistent with the notion that θ_{VSL} equals or exceeds γ.

We also have explained why the social discount rate (δ) should be set equal either to γg (where g is the growth rate) or r (where r represents the real rate of return on government bonds). And we have shown that in practice, the plausible range of values for γg and r are quite close.

What, then, should federal agencies do? Based on economic theory, agencies could justify setting $\theta_{VSL} \times g$ slightly above δ (on the ground that as income rises, consumption utility falls *and* life becomes more pleasant in utility terms). On the basis of empirical studies of γ and the income elasticity of the VSL, agencies could justify setting $\theta_{VSL} \times g$ above, equal to, or below δ, though we are doubtful that the empirical evidence really lends itself to the interpretation that $\theta_{VSL} \times g$ is less than δ. But in any event, a large discrepancy between $\theta_{VSL} \times g$ and δ—or even a small

[83] For a summary of the literature on the "climate beta," see Peter Howard & Jason A. Schwartz, *Valuing the Future: Legal and Economic Considerations for Updating Discount Rates*, 39 YALE J. ON REG. 597, 629-31 (2022).
[84] *See Daily Treasury Rates*, U.S. DEP'T OF THE TREAS., https://home.treasury.gov/resource-center/data-chart-center/interest-rates (last updated Apr. 10 , 2026) (reporting Treasury par real yield curve rates).
[85] *See supra* Section II.A.
[86] *See* Costa & Kahn, *supra* note 45, at 172–73; *supra* notes 45–50 and accompanying text.

discrepancy compounded over a long time period—would be hard to justify on economic or ethical grounds. For example, setting $\theta_{VSL} \times g$ equal to 2% and δ equal to 3% would imply that preventing one death today has the same present value as preventing 2.65 deaths a century from now, preventing 17,260 deaths a millennium from now, or preventing approximately 298 million deaths two millennia from now. The thin empirical evidence suggesting that θ_{VSL} might possibly be less than γ—a suggestion unmoored from economic theory—offers a weak justification for such a weighty discounting of future lives.

We propose, therefore, that agencies adopt a presumption that the VSL growth rate ($\theta_{VSL} \times g$) equals the discount rate (δ). We acknowledge that economic theory could support setting the VSL growth rate slightly *above* the discount rate,[87] in which case future lives would count for *more* in cost-benefit analysis than lives saved today. To justify the current practice of setting the VSL growth rate substantially *below* the discount rate, however, agencies would need to muster theoretical arguments or empirical evidence that are much more robust than what has been offered so far.

F. Discounting From First Principles

Thus far, our analysis has operated primarily within frameworks of utilitarianism and Kaldor-Hicks efficiency. Recall that utilitarians set the discount rate (δ) equal to γg to reflect the diminishing marginal utility of consumption: one unit of consumption in the present period yields the same utility as $1 + \gamma g$ units of consumption in the next period. For utilitarians, then, the VSL growth rate ($\theta_{VSL} \times g$) straightforwardly equals the discount rate when γ equals θ_{VSL}. Whether γ actually equals θ_{VSL} is, for utilitarians, an entirely empirical question.

Kaldor-Hicks efficiency, meanwhile, rests on the premise that the social planner should pursue changes that could be combined with compensating transfers so that at least one person is better off and no one is worse off.[88] Those changes are known as "potential Pareto improvements."[89] Kaldor-Hicks efficiency demands that we set δ equal to r based on the argument we attribute to Dexter Samida and David Weisbach above: Whenever a regulation yields a return less than r, a future generation would be better off (and the present generation would be no worse off) if the government scrapped the regulation and reallocated the regulation's cost toward a capital asset yielding r.[90] Any discrepancy between δ and r suggests that the government is leaving some potential Pareto improvements on the table. The actual values of r, γ, g, and θ_{VSL} depend, again, exclusively on the empirics.

[87] *See* Kaplow, *supra* note 40, at 25; *supra* notes 40–43 & accompanying text.
[88] For technical reasons addressed elsewhere, Kaldor-Hicks efficiency and utilitarianism converge under certain assumptions regarding the optimality or responsiveness of the tax-and-transfer system. *See* Hemel, *supra* note 2, at 674-76.
[89] *See* Russell Hardin, *Magic on the Frontier: The Norm of Efficiency*, 144 U. PA. L. REV. 1987, 1996-97 (1996).
[90] *See* Samida & Weisbach, *supra* note 71; *see also supra* notes 71-72 and accompanying text.

Regulatory cost-benefit analysis in the United States today operates largely under a Kaldor-Hicks efficiency framework,[91] with the significant exception that real-world regulatory cost-benefit analysis—in a deviation from both Kaldor-Hicks and utilitarian principles—assigns the same value of a statistical life to everyone in the United States today regardless of age or income.[92] An alternative approach to cost-benefit analysis is prioritarianism. Whereas utilitarians assign the same value to each individual in the social welfare function and choose policies that maximize the total amount of wellbeing, a prioritarian social welfare function assigns greater weight to the wellbeing of the worse off.[93] An extreme version of prioritarianism—the "difference principle" associated with the philosopher John Rawls[94]—assigns nearly dispositive weight to the wellbeing of the worst off members of society and rejects any policy that reduces the wellbeing of the worst off, even if it leads to large utility gains for everyone else.[95]

A full exploration of prioritarianism and the valuation of present and future lives lies beyond our scope, but for now, it is important to observe that incorporating prioritarianism into the valuation of future lives can yield a broad range of outcomes, depending on how extreme one's prioritarian views are. Prioritarians who follow a less extreme version of the theory than the difference principle could potentially justify assigning an equal value to present and future lives even if γ and θ_{VSL} do not, empirically, turn out to be equal. The argument would go as follows: Economic reasoning leads us to believe that θ_{VSL} exceeds γ (on the theory that richer people derive more happiness—or greater wellbeing—from their lives than poorer people).[96] However, prioritarianism assigns greater weight to the wellbeing of the worse off—and the present generation is worse off than future generations (who presumably will have higher incomes).[97] A prioritarian social welfare function might assign greater weight to the wellbeing of the present generation so as to offset, precisely, the greater utility value of lives saved in the future.[98] This approach would comport with the ethical intuition that society should assign equal value to present and future lives.

---

[91] *See* John D. Graham, *Incorporating Environmental Justice Into Benefit-Cost Analysis of Federal Rulemakings*, 25 RICHMOND PUB. INT. L. REV. 149, 159 (2022); Hemel, *supra* note 2, at 673 n.65.
[92] *See* Hemel, *supra* note 2, at 678.
[93] *See* Matthew D. Adler, Maddalena Ferranna, James K. Hammitt & Nicolas Treich, *Fair Innings? The Utilitarian and Prioritarian Value of Risk Reduction Over a Whole Lifetime*, 75 J. HEALTH ECON. 102412 (2021).
[94] *See* JOHN RAWLS, A THEORY OF JUSTICE (1971).
[95] *See* Matthew D. Adler, *Theory of Prioritarianism*, *in* PRIORITARIANISM IN PRACTICE 37 (Matthew D. Adler & Ole F. Norheim eds., 2022).
[96] *See* Kaplow, *supra* note 40; *supra* notes 40-42 and accompanying text.
[97] Depending on how strongly one decides to prioritize the lives of people who are less well off, a prioritarian approach also might lead one to conclude that future lives should be valued more than present lives—but by a smaller margin than a purely utilitarian calculus. *See* Matthew D. Adler & Nicolas Treich, *Prioritarianism and Climate Change*, 62 ENVTL. & RESOURCE ECON. 279, 296 (2015). We thank Matthew Adler for discussion of this point. For a rich and nuanced prioritarian perspective on intergenerational discounting, see generally Matthew D. Adler, *Future Generations: A Prioritarian View*, 77 GEO. WASH. L. REV. 1478 (2009).
[98] *See* James F. Hammitt & Nicolas Treich, *Prioritarianism and Fatality Risk Regulation*, *in* PRIORITARIANISM IN PRACTICE, *supra* note 95, at 336.

On the other hand, prioritarianism could—in theory—lead one to adopt a social welfare function that more than offsets any empirical discrepancy between θ_{VSL} and γ.[99] At the extreme, one might argue—on Rawlsian grounds—that policymakers should assign nearly dispositive weight to the lives of members of the present generation because the present generation is worse off than all generations to come.[100] Notably, Rawls himself rejected this implication of the difference principle,[101] partly on the ground that members of the present generation care about their progeny and would not want to disregard their descendants' lives.[102]

Or, in between these alternatives, the current regulatory practice of assigning greater weight to lives in the present than to lives in the future could be understood as adding a dash of prioritarianism to what is otherwise a primarily Kaldor-Hicksian cost-benefit analysis framework. To be sure, the dramatic discounting of future lives that is sometimes required by regulatory cost-benefit analysis might be a bridge too far, even for relatively extreme prioritarians. Even a prioritarian like Rawls would likely balk at the EPA's approach, which—in the version with a 7% discount rate—implies that it is better to save one life today than to save forty lives in our grandchildren's generation six decades hence.[103] It is also important to note that the federal government has never claimed to be adopting a prioritarian approach to cost-benefit analysis, and thus it would be tendentious for it to defend its choices on that ground. Nevertheless, in theory prioritarianism could be used to justify something akin to the approach agencies have taken to date, though a relatively extreme version of prioritarianism would be required.

For present purposes, we set aside the question of whether regulatory cost-benefit analysis should be reoriented completely around prioritarian principles. Rather, our goal is to demonstrate that current agency practice is inconsistent with utilitarian and Kaldor-Hicks frameworks, and that if properly applied, those frameworks would lead to substantially different regulatory results. Our project is necessarily incrementalist, but sometimes progress can only be made incrementally or not at all.

G. Contextualizing Critiques of Discounting

Our discussion of the relationship between $\theta_{VSL} \times g$ and δ sheds light on the larger debate over the practice of social discounting—a debate that, despite its inherently technocratic subject matter, has been more heated than the debate over any agency practice in the past several decades (other than perhaps the monetization of life itself). Commentators have focused on the discounting of lives as a thin edge of the wedge in an effort to topple discounting more generally. The approach, deployed on multiple occasions, has been to offer what amounts to a reductio ad absurdum

[99] *See id.*
[100] *See* Adler, *supra* note 97, at 1499.
[101] *See* RAWLS, *supra* note 94, at 291. For a discussion of the bases of Rawls's rejection, see Parath Dasgupta, *Three Conceptions of Intergenerational Justice*, *in* RAMSEY'S LEGACY 149, 157-59 (Hallvard Lillehammer & D.H. Mellor eds., 2005).
[102] *See* RAWLS, *supra* note 94, at 288.
[103] As noted above, EPA's framework implies a VSL growth rate of 0.627%. *See supra* note 7. At a 7% discount rate: 1/((1.00627/1.07)^60) ≈ 40.

regarding discounting lives, and then argue that because the conclusion regarding discounted lives is so preposterous, discounting in all of its forms must be misguided. Our explanation clarifies that a positive social discount rate does *not* imply the devaluation of future lives. It is only when the social discount rate (δ) exceeds $\theta_{\text{VSL}} \times g$ that lives in the future count for less than lives today.

Take, for instance, noted cost-benefit analysis critics Frank Ackerman and Lisa Heinzerling. Ackerman and Heinzerling write that with "a discount rate of five percent, for example, the death of a billion people 500 years from now becomes less serious than the death of one person today."[104] The theory behind this statement is that we should discount the 1 billion lives by 5% for each of the 500 years, yielding a value today of approximately 0.025 lives (roughly 1/40th of a life).[105] Ackerman and Heinzerling find this result untenable and assume that the reader will do the same; they thus take this example to imply that discounting in general is mistaken. Ackerman and Heinzerling point to the "absurdity" of discounting as one ground for indicting the entire cost-benefit analysis enterprise.[106]

Our analysis above offers a response to this type of reductio ad absurdum rhetoric. For all the reasons we described, VSL should increase at roughly the same pace as the discount rate. When δ equals $\theta_{\text{VSL}} \times g$, one billion lives in 500 years will be worth the same as one billion lives today. Ackerman and Heinzerling's critique of discounting applies only to discounting at a higher rate than $\theta_{\text{VSL}} \times g$.

Somewhat more colorful is the argument against discounting offered by economist Tyler Cowen and philosopher Derek Parfit: "Imagine finding out that you, having just reached your twenty-first birthday, must soon die of cancer because one evening Cleopatra wanted an extra helping of dessert."[107] This is meant to strike the reader as similarly absurd, thus demonstrating that it must be error to engage in discounting. The implicit theory behind this statement is that—from the perspective of someone conducting cost-benefit analysis when Cleopatra was Queen of Egypt (51 to 30 BCE)—a life in the 1990s (when Cowen and Parfit wrote) would, due to discounting, have a value less than the value of a slice of cake. For example, if we take $1 million to be the global VSL at the time that Cowen and Parfit were writing, then the present value of $1 million roughly 2020 years prior—applying a 3% discount rate—would be approximately 0.000000000000000001 cents.

But this, too, is based on a fallacy. Let's imagine that the Ptolemaic cost-benefit analyst also was clairvoyant and could foresee the future path of economic growth. According to the economic historian Angus Maddison, per capita global GDP circa 2000 was approximately 13

---

104 Frank Ackerman and Lisa Heinzerling, *Pricing the Priceless: Cost-Benefit Analysis of Environmental Protection*, 150 U. PA. L. REV. 1553, 1571 (2002)

105 $1{,}000{,}000{,}000 / (1.05)^{500} = 0.025$.

106 *See id.* at 1556.

107 Tyler Cowen and Derek Parfit, *Against the Social Discount Rate*, *in* JUSTICE BETWEEN AGE GROUPS AND GENERATIONS 144, 145 (Peter Laslett and James S. Fishkin, eds., 1992).

times per capita global GDP in the year zero.[108] When θ_{VSL} equals 1, VSL grows proportionately with income. So if the global VSL were $1 million in 2000, it would have been—given a θ_{VSL} value of 1—approximately 1/13 of $1 million (roughly $77,000) in Cleopatra's time. If the Ptolemaic cost-benefit analyst adopted our proposal that future lives should be valued the same as present lives, then she would weigh the value of Cleopatra's dessert against the VSL in ancient Egypt ($77,000, in our example). Cowen and Parfit's imagined absurdity—in which the value of an extra helping of dessert for Cleopatra would exceed the value of saving a life in 2000—only arises if cake was extraordinarily expensive in Ptolemaic times.[109]

To be sure, setting the social discount rate (δ) equal to $\theta_{VSL} \times g$ invites a separate critique: that the entire practice of adjusting VSL upwards to reflect income growth and then discounting by the same rate amounts to an elaborate exercise of running in circles. This critique stings if the *only* benefit of the relevant regulation is a reduction in premature deaths. But often that is not the case. For example, vehicle safety regulations generate non-monetary benefits (e.g., reduced costs of car repair and replacement) as well as monetary benefits. The same is true for greenhouse gas regulations. If allowed to continue unchecked, climate change is likely to produce significant *non-life* financial costs: coastal cities could be deluged by rising sea levels; arable farmland could be rendered valueless due to drought; and so forth. Discounting reflects the intuition that reductions in consumption matter more to utility when consumption is lower. And as long as the economy continues to grow—even if at a slowing rate—this intuition implies that reductions in consumption should weigh more heavily in cost-benefit analysis when they occur sooner.

All of which is to say that setting the social discount rate (δ) equal to $\theta_{VSL} \times g$ does not render the debate over discounting irrelevant. But it does make discounting irrelevant to the relative valuation of *lives*. Moreover, it shifts the terrain on which debates over the discount rate occur. When δ equals $\theta_{VSL} \times g$, the devaluation of future lives no longer operates as a strike against social discounting (or against cost-benefit analysis in general). If critics of social discounting seek to discredit the practice nonetheless, they will have to offer different arguments.

III. Agency Practice

As we explained in Part II, there is a strong theoretical case for increasing VSLs over time to keep pace with the social discount rate, such that a life saved in 2046 is worth the same as a life saved in 2026. The empirical evidence is more mixed but not inconsistent, and certainly not decisive enough to counsel deviating from the strong theoretical presumption in favor of setting $\theta_{VSL} \times g$ equal to or greater than δ.

---

[108] *See* Angus Maddison, 1 The World Economy: A Millennial Perspective 30 tbl.1-2 (2006) (estimating that global GDP per capita increased from $444 in the year zero to $5709 in 1998).

[109] Moreover, if the VSL growth rate equals the discount rate, then a slice of cake in 30 BC would bear the same relation to the value of a life in the 1990s as a slice of cake in 30 BCE bears to the value of a life in 30 BCE. The timing of death would be irrelevant to the cost-benefit analysis.

Yet this is not at all what agencies have done. In the sections that follow, we examine four separate agency approaches: those of the Department of Transportation, Environmental Protection Agency, Occupational Safety and Health Administration, and Department of Health and Human Services. In all four cases, we show that the agencies have settled on parameters for θ_{VSL} and g that do not keep pace with the chosen social discount rate and thus yield diminishing values of life over time. We explain that the selection of these values appears to have been driven by considerations that run contrary to the weight of theory and evidence. And we also document the ways in which agency practice appears to have deviated from agencies' own public guidance documents. We also briefly describe a short-lived reform effort by the Biden administration that would have dramatically altered agency practice but was rescinded by President Trump in January 2025.

We then recalculate the cost-benefit analyses of recent regulations promulgated by the Department of Transportation and the EPA. We show that both agencies would have arrived at meaningfully different results had they increased VSL over time to keep pace with the discount rate.

A. Agency Guidance on Increasing VSL

1. The Department of Transportation

Twice within the last fifteen years, in 2013 and 2021, the Department of Transportation has produced guidance documents outlining its approach to calculating increases in VSL over time.[110] Both of these documents involve analytic decisions that have significant impacts on the eventual calculations. Here, we raise doubts about several of those agency choices.

We begin with the DOT's 2013 approach. For its value of g, DOT used estimates of the growth in the Real Median Wage for Workers Covered by Social Security, which is calculated by the Congressional Budget Office.[111] These estimates were unpublished at the time the DOT promulgated its 2013 guidance, but the DOT reported that the Congressional Budget Office estimated real median wage growth of 1.07% over the next 30 years.[112] The choice to use that growth figure—rather than, for example, the growth rate of real GDP per capita—was consequential: at the time, the Congressional Budget Office estimated a significantly higher growth rate for real GDP per capita (approximately 1.6%).[113]

---

[110] DOT 2013 Guidance, *supra* note 2, at 7; U.S. DEP'T OF TRANSP., DEPARTMENTAL GUIDANCE: TREATMENT OF THE VALUE OF PREVENTING FATALITIES AND INJURIES IN PREPARING ECONOMIC ANALYSES 8 (2021), https://www.transportation.gov/sites/dot.gov/files/2021-03/DOT%20VSL%20Guidance%20-%202021%20Update.pdf (hereinafter DOT 2021 Guidance).
[111] DOT 2013 Guidance, *supra* note 2, at 7, at 8-9.
[112] *Id.*
[113] CONGRESSIONAL BUDGET OFFICE, ECONOMIC VARIABLES UNDERLYING THE LONG-TERM BUDGET PROJECTIONS AND PROJECTIONS OF GDP AND POPULATION (June 2012), https://www.cbo.gov/sites/default/files/recurringdata/51119-2012-06-longtermbudgetprojections.xls.

DOT's choice to use the growth rate of real median wages rather than real GDP per capita as its value for *g* is curious for several reasons.[114] First, as DOT foresaw, real wage growth was likely to be slowed by a shift in compensation from wages to health insurance benefits—a shift accelerated by the implementation of the Affordable Care Act starting in 2014.[115] Second, economists in recent years have raised serious concerns about the methodology used by government statistical agencies to impute wage income to self-employed individuals, noting that current methods likely underestimate the growth of labor income.[116] Third, while DOT said it believes that real median wages "are more relevant to estimating the VSL of the average person affected by transportation-related safety risks,"[117] the age group with the highest motor vehicle death rate is people over 75, for whom wages constitute less than half of income.[118] Notably, DOT's decision to use the growth rate of real wages rather than real GDP per capita would, on its own, reduce the present value of a life saved in 30 years by approximately 15% based on the figures reported by the agency.[119]

To calculate θ_{VSL}, DOT canvassed a series of meta-analyses, which found as follows:

> In our [2008] guidance, we cite research by Viscusi and Aldy (2003) that estimated the elasticity of VSL with respect to increases in real income as being between 0.5 and 0.6 . . . . We accordingly have increased VSL by 0.55 percent for every one-percent increase in real income. More recent research by Kniesner, Viscusi, and Ziliak (2010) has derived more refined income elasticity estimates ranging from 2.24 at low incomes to 1.23 at high incomes, with an overall figure of 1.44. An alternative specification yielded an overall elasticity of 1.32. Similarly, Costa and Kahn (2004) estimated the income-elasticity of VSL to be between 1.5 and 1.6. These empirical results are consistent with theoretical arguments suggesting that the income-elasticity of VSL should be greater than 1.0.[120]

If DOT had merely taken the average of these three studies, it would have arrived at a value of θ_{VSL} equal to 1.16.[121] Indeed, there are reasons to support even higher values of θ_{VSL}. The lowest figure (from the 2003 Viscusi and Aldy study) is also based on the oldest study—and one that involved meta-analysis of existing VSL estimates (an approach subject to serious methodological

[114] DOT's choice to use the median rather than the mean also leads to the result that an increase in labor income inequality (manifesting through faster growth of the mean rather than median) corresponds to *less* stringent environmental, health, and safety regulations.
[115] *See id.* at 8.
[116] *See, e.g.*, Michael W.L. Elsby, Bart Hobijn & Ayşegül Şahin, *The Decline of the U.S. Labor Share*, BROOKINGS PAPERS ON ECON. ACTIVITY, Fall 2013, at 15 (suggesting that part of the decline in the labor share of GDP is "spurious").
[117] DOT 2013 Guidance, *supra* note 2, at 7, at 8 n.19.
[118] *See* SOCIAL SECURITY ADMINISTRATION, INCOME OF THE AGED CHARTBOOK, 2014, at 16 (Apr. 2017), https://www.ssa.gov/policy/docs/chartbooks/income_aged/2014/iac14.pdf.
[119] I.e.: $(1.016^{30} - 1.0107^{30})/1.016^{30} \approx 15\%$.
[120] DOT 2013 Guidance, *supra* note 2, at 7.
[121] The average of the Viscusi & Aldy study was 0.55; the average of the Kniesner, Viscusi & Ziliak study is 1.38 (i.e., the average of the overall estimate of 1.44 and the alternative estimate of 1.32), and the average of the Costa & Kahn study is 1.55. The average of these three numbers is 1.16.

doubts discussed above[122]). In addition, as DOT acknowledges, and as we described above in Part II, there are theoretical reasons to believe that the income elasticity of VSL should be above 1. So one might have imagined that DOT would treat $\theta_{\text{VSL}} = 1.16$ as a floor and potentially select a higher value.

Instead, the agency concluded that "[i]n view of the large increase in the income elasticity of VSL that would be suggested by these empirical results, and because the literature seems somewhat unsettled, we will increase our suggested income-elasticity figure only to 1.0."[123] The agency explained:

> While this figure is lower than the elasticity estimates of Kniesner et al. and Costa and Kahn, it is higher than that of Viscusi and Aldy, the basis for our previous guidance. It is difficult to state with confidence whether a cross-sectional income elasticity (such as those estimated in these empirical analyses), representing the difference in sensitivity to fatality risks between low-income and high-income workers in a given population, corresponds to a longitudinal elasticity, representing the way in which VSL is affected by growth in income over time for an overall population. Consequently, we adopt this more moderate figure, pending more comprehensive documentation.[124]

In the passage above, DOT correctly notes that the valuation of future lives depends upon the longitudinal elasticity rather than the cross-sectional elasticity of VSL (i.e., changes in VSL over time as income rises, rather than differences in VSL between the rich and the poor). But DOT incorrectly states that all three of the cited empirical studies are cross-sectional. As explained above, the Costa and Kahn study is longitudinal—precisely the sort of evidence that DOT suggests would be most compelling.[125] And the lower-bound estimate for θ_{VSL} in the Costa and Kahn study was 1.5.[126] The end result of these choices regarding θ_{VSL} and g was that DOT increased VSL by 1.07% per year, well below the lower of the two OMB-recommended discount rates (3%). The effect of this slow growth in VSL, coupled with much higher discount rates, is to dramatically reduce the value of future lives relative to current ones.

DOT's 2021 guidance largely tracks its 2013 approach.[127] DOT again uses the same measure of wage growth for g and states that "we do not take into account changes in non-wage income" when updating VSL, "on the grounds that this non-wage income is not likely to be significant for the average person affected by our rules"[128]—a puzzling statement for an agency

[122] *See supra* Section II.D.
[123] DOT 2013 Guidance, *supra* note 2, at 7.
[124] *Id.*
[125] *See* Costa & Kahn, *supra* note 45.
[126] *Id.* at 173.
[127] DOT's 2021 guidance appears to still be in effect as of 2025. *See* U.S. DEP'T OF TRANSP., *supra* note 2 (noting an "Effective Date: Monday, April 28, 2025").
[128] DOT 2021 Guidance, *supra* note 110, at 7.

whose rules affect the young and the old as well as working-age adults.[129] Again, DOT credits "theoretical arguments suggesting that the income-elasticity of VSL should be greater than 1.0,"[130] and again, DOT cites empirical studies consistent with those theoretical arguments. But once again, DOT opts for the "more moderate" approach of setting θ_{VSL} equal to 1.[131] The agency does not address (or even acknowledge) the fact that this approach—in combination with low growth rates and the 3% and 7% discount rates recommended by OMB—will result in significant devaluation of future lives.

2. *The Environmental Protection Agency*

In contrast to DOT's relatively transparent (though questionable) approach to increasing VSL, EPA's documentation is more opaque. According to its most recent guidance document, EPA assumes a triangular distribution of values of θ_{VSL} with a low end of 0.08, a central value of 0.4, and a high end of 1.0.[132] In one publicly posted white paper, EPA states that "[m]any analyses characterize this range with a triangular distribution with a resulting mean estimate of approximately 0.48."[133] (Curiously, the mean of 0.08, 0.40, and 1.0 actually rounds to 0.49, not 0.48.) Yet in a footnote to a recent proposed rule, EPA stated that it was using an income elasticity value of 0.4, the median rather than the mean of the triangular distribution.[134] We have detailed the reasons why this income elasticity for VSL is improbably low, and we will not rehash them here. Suffice it to say that our doubts regarding DOT's decision to use an income elasticity of 1.0 apply with even greater force to EPA's apparent parameter of 0.4.

For *g*, recent documentation from EPA indicates that the agency uses real GDP per capita rather than real wages.[135] As noted, the use of real GDP per capita is more easily justifiable than DOT's choice to use real wages when regulations affect all members of society.[136] Still, EPA's

---

[129] Even if children are in families that receive only wage income, changes in family size—which would be reflected in real GDP per capita but not real wages—certainly will affect consumption levels.
[130] DOT 2021 Guidance, *supra* note 110, at 7.
[131] *Id.* at 9.
[132] U.S. ENVIRONMENTAL PROTECTION AGENCY, GUIDELINES FOR PREPARING ECONOMIC ANALYSES, at B-4 (Dec. 2010), https://www.epa.gov/sites/default/files/2017-09/documents/ee-0568-22.pdf [hereinafter EPA 2010 GUIDELINES).
[133] U.S. Environmental Protection Agency, *Valuing Mortality Risk Reductions for Environmental Policy: A White Paper* 44 (Scientific Advisory Review Board Draft, Dec. 10, 2010), https://www.epa.gov/sites/default/files/2017-08/documents/ee-0563-1.pdf.
[134] National Primary Drinking Water Regulations for Lead and Copper: Improvements (LCRI), 88 Fed. Reg. 84878, 85002 n.23 (Dec. 6, 2023); *see also* H. Spencer Banzhaf, *The Value of Statistical Life: A Meta-Analysis of Meta-Analyses*, 13 J. BENEFIT-COST ANAL. 182, 183 (2022) ("EPA adjusts for changes in income using an income elasticity of 0.4, which is lower than other agencies . . . .").
[135] U.S. ENVIRONMENTAL PROTECTION AGENCY, BENMAP: ENVIRONMENTAL BENEFITS MAPPING AND ANALYSIS PROGRAM, USER'S MANUAL, VERSION 1.0, at H-31 (July 2025), https://www.epa.gov/system/files/documents/2025-07/benmap_user_manual_v1.0_508.pdf [hereinafter BENMAP USER'S MANUAL]; *see also* Banzhaf, *supra* note 134, at 183 (noting that EPA measures income for VSL purposes "with per-capita GDP, which has grown much faster than the measure of median earnings used by other agencies").
[136] *See infra* notes 114–119 and accompanying text.

choice of a lower value for θ_{VSL} more than offsets its choice of a higher value for g. Illustratively, from 2013 to 2025, EPA's mid-point VSL estimate increased by approximately 6% in real terms,[137] whereas DOT's VSL estimate appears to have increased by approximately 12% in real terms.[138]

The approach we have detailed above applies to EPA's valuation of harms from all pollutants *other than greenhouse gases*. EPA's approach to valuing future harms from climate change exposes the flaws in the agency's calculations even more starkly. In a guidance document related to the social cost of greenhouse gases, the agency reports that its best estimates for γ range from 1.02 to 1.42.[139] This is consistent with EPA's survey of the literature, reported in that same document, which finds values of γ ranging from 1 to 4, with a mean of 1.35.[140] As noted, γ sets a theoretical lower bound on the income elasticity of the VSL.[141] Yet when EPA actually gets around to attaching a value to future U.S. lives, it uses the same low rate as it does for other pollutants, implying that (against all evidence) it has again adopted a θ_{VSL} well below 1.[142]

In addition, because climate change is a global problem, EPA must calculate the value of lives in countries other than the United States. VSL studies do not exist for most countries, however. Accordingly, in order to value lives saved abroad, EPA adjusts international VSLs up or down based on the per capita income in those other countries relative to the United States.[143] In countries with higher per capita income, EPA calculates higher VSLs; in countries with lower per capita income, EPA calculates lower VSLs. This calculation is in some sense identical to the calculation EPA performs when valuing future lives: as people in the future (or in other countries) become wealthier (or poorer), their willingness to pay for life-saving interventions goes up (or down). Here too, then, EPA must choose an income elasticity of VSL to use when translating American VSLs into international VSLs. In this case, however, EPA states explicitly that it will use a VSL of 1.0, rather than the value of 0.4 that it appears to use when valuing future lives within

---

[137] *See* BENMAP USER'S MANUAL, *supra* note 135, at H-33.
[138] With $\theta_{\text{VSL}} = 1$ and g based on the real median wage, the increase in real VSL is simply the increase in the real wage, which was 12% from January 2013 to January 2025. *See Employed Full Time: Median Usual Weekly Real Earnings: Wage and Salary workers: 16 Years and Over*, FEDERAL RESERVE BANK OF ST. LOUIS: FRED ECONOMIC DATA, https://fred.stlouisfed.org/series/LEU0252881600A (last visited Feb. 16, 2026).
[139] U.S. ENVIRONMENTAL PROTECTION AGENCY, REPORT ON THE SOCIAL COST OF GREENHOUSE GASES: ESTIMATES INCORPORATING RECENT SCIENTIFIC ADVANCES 60 tbl.2.4.2 (Sept. 2022), https://www.epa.gov/system/files/documents/2022-11/epa_scghg_report_draft_0.pdf [hereinafter REPORT ON THE SOCIAL COST OF GREENHOUSE GASES].
[140] *Id.* at 61.
[141] *See supra* notes 40–43 and accompanying text.
[142] *See* REPORT ON THE SOCIAL COST OF GREENHOUSE GASES, *supra* note 139, at 129 (noting that "projected changes in premature mortality in the U.S. are monetized using the same value of mortality risk reduction as in the EPA's regulatory analyses").
[143] *Id.* ("Valuation of mortality risk changes outside the U.S. is based on an extrapolation of the EPA value . . . .").

the United States.[144] Apparently the rule at EPA is θ_{VSL} consistent with theory and evidence for thee, but not for me. The result is a devaluation of the lives of future Americans.

Furthermore, the choice to use a higher income elasticity for international comparisons than for future American lives is not only inconsistent, but it also biases estimates of the social cost of greenhouse gases downward. The reason for this is that most deaths from climate change are projected to occur in low and lower-middle-income countries.[145] A higher θ_{VSL} for international comparisons results in a lower VSL for countries that are poorer than the United States and a higher VSL for countries that are richer than the United States. Since the mortality costs of climate change are concentrated in the former, EPA's choice to use a value for θ_{VSL} of 1 rather than 0.4 for international comparisons reduces its mortality cost estimate. Then when projecting domestic deaths (where a higher θ_{VSL} would produce a higher mortality cost estimate), EPA flips to the lower income elasticity value.

At the end of the day, the effect of EPA's approach is to dramatically reduce the value of future lives relative to current ones in EPA cost-benefit analyses. Historically, EPA's approach has produced increases in the VSL of approximately 0.627% per year in real dollars.[146] Increasing the real VSL by 0.627% per year and discounting by 3%, a life saved ten years in the future would be worth 21% less than a life saved today;[147] a life saved thirty years in the future would be worth 50% less than a life saved today;[148] and a life saved fifty years in the future would be worth 69% less than a life saved today.[149] The devaluation of future lives is even more dramatic at a 7% discount rate: 46% over a ten-year timeframe;[150] 84% over thirty years;[151] and 95% over fifty years.[152] The figures for longer timeframes are of particular importance because many of the greatest harms from environmental degradation are anticipated to materialize decades down the road.

3. The Department of Health and Human Services

The Department of Health and Human Services' approach to increasing VSL over time is similar to DOT's current approach. The agency writes in its 2016 guidelines:

> Adjusting for real income growth is a separate step that requires two inputs: an estimate of the change in population-wide real income per person, and an estimate of the extent to which WTP is expected to change in response to the income change. The latter is generally

[144] *See id.* (noting "an assumed income elasticity of 1" for the valuation of foreign lives).
[145] *See* R. Daniel Bressler, *Breaking Down the Social Cost of Carbon* 4–5 (Columbia Univ., Working Paper, Jan. 21, 2025), https://rdanielbressler.com/s/JMP.pdf.
[146] *See supra* note 7.
[147] $1 - (1.00627/1.03)^{10} \approx 0.21$.
[148] $1 - (1.00627/1.03)^{30} \approx 0.50$.
[149] $1 - (1.00627/1.03)^{50} \approx 0.69$.
[150] $1 - (1.00627/1.07)^{10} \approx 0.46$.
[151] $1 - (1.00627/1.07)^{30} \approx 0.84$.
[152] $1 - (1.00627/1.07)^{50} \approx 0.95$.

expressed as the percentage change in the VSL associated with a one percent change in real income; i.e., the income elasticity.[153]

Like DOT, HHS decides to use an income elasticity of 1.0. The agency's explanation is similarly unsatisfying:

> Some research suggests that a one percent change in income leads to less than a one percent change in the VSL (e.g., Viscusi and Aldy 2003), and other research suggests that it leads to more than a one percent change (e.g., Kniesner, Viscusi and Ziliak 2010, Viscusi 2015). Given this uncertainty, HHS analysts should apply an income elasticity of 1.0 in their analyses.[154]

The agency fails to discuss the impact of this choice in light of the OMB-recommended discount rates of 3% and 7%.[155]

For its estimates of future wage growth (*g*), HHS uses the Congressional Budget Office's projections for real earnings per worker, which—as of the latest HHS guidance in March 2026—CBO expected would grow at a rate of 1.117%.[156] This choice, which mirrors the approach taken by DOT, is itself curious and significant for the reasons we described above.[157]

In any event, coupled with a 1.0 income elasticity of VSL, this means that HHS will increase VSL by 1.117% yearly.[158] Against a 3% discount rate, this in turn means that VSL will steadily erode over time. HHS's central VSL estimate for 2026 was $14.1 million (in 2025 dollars).[159] From a 2026 vantagepoint, given HHS's parameter choices, the present value of a life saved 10 years in the future would be only $11.8 million;[160] thirty years in the future (2056) it would be $8.2 million;[161] and 50 years into the future (2076) it would be only $5.8 million[162] (with all figures in 2025 dollars).

4. *The Occupational Safety and Health Administration*

Unlike the three prior agencies, the Occupational Safety and Health Administration (OSHA) does not publish official guidance documents in which it lays out its approach to valuing future lives. Rather, the agency's methodology must be extrapolated from its regulations, many of

---

[153] OFFICE OF THE ASSISTANT SECRETARY FOR PLANNING AND EVALUATION, U.S. DEPARTMENT OF HEALTH AND HUMAN SERVICES, GUIDELINES FOR REGULATORY IMPACT ANALYSIS 16 (2016), https://aspe.hhs.gov/sites/default/files/migrated_legacy_files//171981/HHS_RIAGuidance.pdf.
[154] *Id.*
[155] In its 2025 guidance, HHS adheres to an income elasticity of 1.0 without further discussion. *See* Kearsley, *supra* note 2, at 2.
[156] *Id.*
[157] *See supra* notes 114–119 and accompanying text.
[158] *See* Kearsley, *supra* note 2, at 2.
[159] *Id.* at 2 tbl.1.
[160] $14,100,000 x (1.0117/1.03)^10 = $11,785,943.79.
[161] $14,100,000 x (1.0117/1.03)^30 = $8,234,834.42
[162] $14,100,000 x (1.0117/1.03)^50 = $5,753,675.67

which involve short discussions of its approach. Notably, however, the agency's approach appears to have shifted over time in ways that make it less reliable and often harder to parse.

As of 2016, OSHA employed a type of analysis that was better aligned with economic theory than the other agencies we have surveyed. In the economic analysis accompanying a 2016 regulation of workplace exposure to silica compounds (which are known carcinogens),[163] the agency began by surveying the relevant literature on g and selected 1.4% as its best estimate of the growth rate of real income per capita.[164] OSHA then reviewed the same empirical literature on θ_{VSL} canvassed by the other agencies and concluded that the Costa and Kahn study provided "[p]robably the most direct evidence" of $\theta_{\mathrm{VSL}} \times g$.[165] OSHA announced that it will use a value for $\theta_{\mathrm{VSL}} \times g$ of 2%, implying a θ_{VSL} of slightly more than 1.4[166]—less than Costa and Kahn's estimate of 1.5 to 1.7, but higher than any other agency.

By 2024, however, the agency's approach appeared to have shifted. In December 2024, the agency issued a final rule revising its personal protective equipment standard for construction to explicitly state that protective gear must fit properly.[167] So far as we were able to determine, this rule contains the agency's most recent statement of its methodology for calculating the value of future lives. In its analysis of the rule's benefits, OSHA stated that it was unable to quantify the number of deaths averted by the rule "due to data limitations."[168] The agency then added: "If OSHA were to monetize fatalities potentially avoided by this final rule, the analysis would use the Department of Transportation (DOT) 2023 value-of-a-statistical-life (VSL) estimate of $13.2 million per avoided fatality."[169] From that starting point, OSHA explained that it would use an income elasticity of the VSL of 1 and a growth rate based on the growth of real gross domestic product over a ten-year timeframe (1.7% per year according to then-current CBO projections).[170] OSHA explained that its resulting estimate for the VSL growth rate would be 1.7%,[171] still higher

---

[163] Occupational Safety & Health Admin., Final Economic Analysis and Final Regulatory Flexibility Analysis: Supporting Document for the Notice of Proposed Rulemaking for Occupational Exposure to Crystalline Silica (2016), https://www.regulations.gov/document/OSHA-2010-0034-4247 [hereinafter OSHA Final Economic Analysis).
[164] *Id.* at VII-44.
[165] *Id.*
[166] This figure comes from dividing OSHA's 2% estimate for $\theta_{\mathrm{VSL}} \times g$ by its 1.4% estimate for per-capita income growth (2%/1.4% ≈ 1.43). That figure roughly matches the estimate for θ_{VSL} in the quantile regression analysis by Kniesner et al., *supra* note 52 (1.44). *See* OSHA Final Economic Analysis, *supra* note 163, at VII-45 (noting that Kniesner et al.'s estimate, "[a]pplied to a long-term increase in per capita income of about 1.4 percent a year," would suggest a VSL growth rate of "about 2 percent a year").
[167] U.S. Department of Labor, Occupational Safety and Health Administration, Personal Protective Equipment in Construction, 89 Fed. Reg. 100,321 (Dec. 12, 2024).
[168] *Id.* at 100,343.
[169] *Id.*
[170] *Id.* at 100,343 n.30 (citing CONG. BUDGET OFF., THE BUDGET AND ECONOMIC OUTLOOK: 2022 TO 2032, at 23 tbl.2–1 (May 2022)).
[171] *See id.* at 100,343 n.31).

than DOT (1.07%),[172] EPA (0.627%),[173] and HHS (1.117%),[174] but below the 2% rate that the agency had chosen in 2016 and below the discount rate that OSHA uses.[175]

5. *OMB's Short-Lived Attempt at Uniformity*

For a brief period of time, it appeared that the federal government might have harmonized these discordant approaches to valuing future lives. In November 2023, the Office of Management and Budget published a significant set of revisions to Circular A-4, the document that instructs federal agencies on how to conduct cost-benefit analysis.[176] The aspect of the revision that garnered the most attention was OMB's introduction of distributional weighting into cost-benefit analysis.[177] But in order to implement distributional weighting, OMB needed to choose a value for γ.[178] The value that it chose was 1.4.[179] The revised Circular A-4 also established a new default discount rate (δ) of 2%.[180] OMB did not tell agencies what value to use for the income elasticity of the VSL, but if agencies had followed the strong theoretical arguments for setting θ_{VSL} at or above γ, they would have arrived at a VSL growth rate that was close to the discount rate. For example, with a value for θ_{VSL} of 1.4 and an income growth rate equal to CBO's then-current estimate of the long-term trend for real GDP per capita (1.3% per year[181]), the VSL growth rate would have been 1.82% (i.e., 1.4 x 1.3% = 1.82%). The end result would have been to value future lives almost equivalently to present lives.

Yet that is not what occurred. In the few months following the release of the new Circular A-4, two agencies published updated indications of their approach to valuing future lives—the EPA in the Regulatory Impact Analysis accompanying a December 2023 rule and the Department of Health and Human Services in a January 2024 guidance document. In both cases, the agencies settled on values that diverged from what OMB has required. And in both cases, the agencies selected values that made future lives significantly less valuable than present ones. In the course of a December 2023 Clean Water Act rule, EPA quietly mentioned in a footnote that it was

---

172 *See supra* Section III.A.1.
173 *See supra* Section III.A.2.
174 *See supra* Section III.A.3.
175 The personal protective equipment regulation used a 2% discount rate, in line with Biden administration guidance. 89 Fed. Reg. at 100,339–41. Presumably the agency will revert to the use of 3% and 7% discount rates going forward. *See infra* note 184 & accompanying text.
176 Off. of Mgmt. & Budget, Circular No. A-4, https://bidenwhitehouse.archives.gov/wp-content/uploads/2023/11/CircularA-4.pdf (November 9, 2023) [hereinafter Revised Circular A-4].
177 *See, e.g.*, Arthur G. Fraas, John D. Graham, Kerry Krutilla, Randall Lutter, Jason F. Shogren, and W. Kip Viscusi, *Improving Distributional Analysis in Regulatory Evaluation: An Assessment of the 2023 Circular A-4*, 19 REV. ENVT'L ECON. & POL'Y 2010 (2025).
178 For an explanation of the relationship between the absolute value of the income elasticity of marginal utility and distributional weighting, see Daniel J. Hemel, *Wealth, Schmealth, Welfare, and Schmelfare*, 59 WAKE FOREST L. REV. 1103, 1122 (2024).
179 *See* Revised Circular A-4, *supra* note 176, at 67.
180 *Id.* at 76-77.
181 *See* CONGRESSIONAL BUDGET OFFICE, THE 2023 LONG-TERM BUDGET OUTLOOK 28 tbl.3-1 (June 2023), https://www.cbo.gov/system/files/2023-06/59014-LTBO.pdf.

continuing to use a value of θ_{VSL} equal to 0.4, not 1.4.[182] And in a guidance document published in January 2024, HHS doubled down on its prior approach as well. Without even referencing the new OMB guidance, it announced that it would maintain its "assumption" that θ_{VSL} is equal to 1.0.[183]

A year later, this issue became moot: Less than two weeks after taking office for a second time, President Trump directed OMB to repeal the revised version of Circular A-4 and revert to its prior instantiation, which did not set any value for γ (and which required discount rates of 3% and 7%).[184] The valuation of future lives is thus more uncertain than ever: it is impossible to know what framework future administrations will adopt, much less which particular values individual agencies will employ.

B. Regulations, Recalculated

The problems we have described with how agencies value future lives are not merely of theoretical interest. Instead, as we have emphasized, they have had (and continue to have) a meaningful impact on agencies' regulatory choices, and thus on individuals' lives. This section provides two illustrations of that point. Below, we recalculate the cost-benefit analyses for two relatively recent regulations, using parameters for valuing future lives that we believe are better supported by theory and evidence. As the following discussion will illustrate, the debate surrounding these regulations could have been very different—and the regulatory outcomes might have been different as well—had agencies more accurately valued future lives.

1. DOT's Rearview Camera Regulation Rethought

For the most part, agencies publish cost-benefit analyses only for regulations they adopt, and they adopt only regulations that pass a cost-benefit test. A notable exception is DOT's 2014 rear visibility rule, which the agency adopted pursuant to a congressional mandate notwithstanding the fact that the regulation flunked DOT's own cost-benefit analysis. To show how choices regarding θ_{VSL}, g, and δ can have significant effects on whether a regulation passes a cost-benefit test, in this section we recalculate DOT's analysis of its 2014 rear visibility rule using different parameters for θ_{VSL}, g, and δ.

DOT's 2014 rule requires new motor vehicles to have rearview cameras starting with the model year 2018 fleet.[185] These cameras—familiar to most readers who have driven or ridden in

---

[182] *See* National Primary Drinking Water Regulations for Lead and Copper: Improvements (LCRI), 88 Fed. Reg. 84878, 85002 n.23 (Dec. 6, 2023).

[183] *See* Aaron Kearsley, *HHS Standard Values for Regulatory Analysis, 2024*, OFF. OF THE ASS'T SEC'Y FOR PLANNING & EVALUATION, U.S. DEP'T OF HEALTH & HUMAN SERVS. (Jan. 25, 2024), https://aspe.hhs.gov/sites/default/files/documents/cd2a1348ea0777b1aa918089e4965b8c/standard-ria-values.pdf.

[184] Exec. Order No. 14,192, 90 Fed. Reg. 9065, 9067 (Jan. 31, 2025) ("The Director shall revoke OMB Circular No. A-4 of November 9, 2023 (Regulatory Analysis), and all accompanying appendices, guidelines, and documents, and shall reinstate the prior version of Circular A-4, issued on September 17, 2003").

[185] Technically, the rule does not specify that car manufacturers must install cameras. But as one report explains:

cars manufactured in the last several years—are designed to reduce death and injury from accidents that occur when cars are backing up.[186] The regulation was controversial at the time it was issued because it was one of the few Obama-era regulations for which the projected costs exceeded the projected benefits. (That is, it did not pass a cost-benefit test.) This fact reportedly caused Obama's Office of Information and Regulatory Affairs (OIRA) to hold up the rule for several years while demanding that the agency investigate cheaper alternatives.[187] Cass Sunstein, who was OIRA administrator for some of the relevant period, later wrote a thoughtful article defending the regulation—even though quantified costs exceeded quantified benefits—on the ground that the rule also generated significant additional benefits that are "difficult to monetize."[188] However, as our analysis will make clear, if projected future lives saved are valued the same as present lives (i.e., if δ is set equal to $\theta_{\mathrm{VSL}} \times g$), the conclusion that this regulation fails a cost-benefit test is cast into significant doubt.

As with many regulations, the costs of DOT's rule were predominantly economic expenses that would be borne in the near term: here, the cost of installing backup cameras in new vehicles starting in 2018. And as with many regulations, the benefits took the form of lives saved and would be realized at some point in the future. The issue of the proper discount rate thus arises because installing a backup camera in a 2018 automobile will create installation costs in 2018, but the benefits may not arise until years later. Someone could be driving that same car in 2028 or even 2038 and use the backup camera to avoid an accident.

To calculate the expected benefits from the regulation, DOT assumed that the number of backup accidents avoided would be proportional to the number of miles a car was driven in a given year. Accordingly, the agency first estimated the number of miles that a given 2018 car would be driven in each year of its projected lifespan. Not surprisingly, an automobile's usage is weighted toward its first years in existence. DOT estimated that roughly 67% of the total miles ever driven by a 2018 passenger car would be driven in the first ten years that car was on the road, 30% would be driven in the next decade, and only 3% of the total miles would be driven in years 21 through 30.[189] This means that this DOT regulation was comparatively weighted toward the present day,

---

> Under the new rule, all vehicles will have to come equipped with the ability for the driver to see a 10-foot by 20-foot zone directly behind the vehicle. There are also requirements involving image size and other factors that pretty much ensure that rear-view cameras are the only solution that will work.

Chris Woodyard, *NHTSA To Require Backup Cameras on All Vehicles*, USA TODAY (Mar. 31, 2014), https://www.usatoday.com/story/money/cars/2014/03/31/nhtsa-rear-view-cameras/7114531.

186 U.S. Department of Transportation, National Highway Transportation Safety Administration, Final Regulatory Impact Analysis—Backover Crash Avoidance Technologies, FMVSS No. 111, at i (2010), https://www.regulations.gov/document/NHTSA-2010-0162-0255 [hereinafter Final Regulatory Impact Analysis].

187 *See* Scot J. Paltrow, *How a Small U.S. Agency Stalls Life-Saving Regulations*, REUTERS (Oct. 29, 2015), https://www.reuters.com/investigates/special-report/usa-regulations-oira.

188 *See* Cass R. Sunstein, *Rear Visibility and Some Unresolved Problems for Economic Analysis (With Notes on Experience Goods)*, 10 J. BENEFIT-COST ANAL. 317, 347 (2019).

189 Final Regulatory Impact Analysis, *supra* note 186, at 88 tbl.VII-1a.

to a far greater degree than regulations meant to curb climate change or even regulations meant to address more conventional air pollutants, all of which will likely have their greatest effects decades in the future.

In accordance with its 2013 guidance document, which was applicable to this 2014 regulation, DOT employed a VSL of $8.86 million in 2010 dollars, which translates to a VSL of $9.65 million in 2018 using DOT's 1.07% rate of VSL growth.[190] For every life saved in the future, DOT increased the value of that life by 1.07% per year and discounted the future benefit by, alternatively, 3% and 7%. (We focus here on the 3% discount rate calculation for ease of explication.)[191] DOT's preferred system involved a 180-degree camera with a dashboard-mounted display screen. It estimated that the system would prevent approximately 35.5 deaths over the lifetime of the model year 2018 cars in which it was installed.[192] DOT thus estimated that, for model year 2018, the regulation would generate $316.6 million in benefits due to lives saved.[193] The agency also estimated an additional $79.4 million in benefits from other sources, for a total of $396 million in benefits from installing cameras in 2018 cars.[194] It weighed this against an estimated $556.7 million in installation costs for 2018 cars, thus arriving at expected net benefits of *negative* $160.7 million for model year 2018 alone.[195] This negative cost-benefit result is what spurred the controversy surrounding this regulation.

---

[190] DOT expressed all costs in 2010 dollars (i.e., not adjusting for post-2010 inflation). Thus, the $9.65 million figure reflects a VSL in 2018 of $9.65 million *in 2010 dollars*, or $11.04 million in 2018 dollars using the Consumer Price Index for All Urban Consumers (CPI-U) to adjust nominal prices. U.S. BUREAU OF LABOR STATISTICS, CPI INFLATION CALCULATOR, https://www.bls.gov/data/inflation_calculator.htm (last visited Sept. 6, 2022) (reporting that $9.65 million January 2010 had the same buying power as $11.04 million in January 2018).

[191] In the course of this calculation, DOT made a further error. DOT should have used a single VSL, namely $9.65 million in 2018 dollars at DOT's 1.07% rate of increase. Instead, however, it calculated a new VSL for every year that a 2018 automobile was expected to be operating, out to 2047, and then took the weighted average of those VSLs. It then applied the 1.07% inflation rate and 3% discount rate. The effect was that DOT increased the VSL by 1.07% twice, rather than once. *See* Final Regulatory Impact Analysis, *supra* note 186, at 92. This meant that DOT's reported benefits number was incorrect given its stated methodology. However, the effect of this error was relatively small—much smaller than the effect of eliminating the discounting of future lives.

[192] *Id.* at 94 tbl.VII-3. The actual DOT calculation reported in this table is different (and peculiar). DOT applied its discounted VSL to the population of lives that might be saved by the regulation before calculating how many lives would *actually* be saved by the regulation. In order to determine the correct figure (which is the 35.5 lives reported in the text), it is necessary perform the same calculation without the discounting. Accordingly, in Table VII-3, the column for the 180 degree camera at a 3% discount rate reports a target population of 170 people killed by passenger cars, with 40% of those cars lacking cameras and a final system effectiveness of 33%. 170 people × 40% × 33% = 22.44 lives saved. Performing the same calculation for light trucks yields 264 people × 15% × 33% = 13.068 lives saved. 22.44 + 13.068 = 35.508 lives saved in total.

Note that we refer in text to "lives," but the actual unit of analysis for DOT is "equivalent lives." DOT translates non-fatal injuries into "equivalent lives" using the Maximum Abbreviated Injury Scale, which rates nonfatal injuries from MAIS 1 (minor) to MAIS 5 (critical). For example, avoiding an MAIS 5 injury (e.g., a ruptured liver with tissue loss) translates to 0.593 equivalent lives. *See* Hemel, *supra* note 1, at 687.

[193] Final Regulatory Impact Analysis, *supra* note 186, at tbl.VII-7.

[194] *Id.*

[195] *Id.* $396 million - $556.7 million = - $160.7 million.

If DOT had instead set $\theta_{\mathrm{VSL}} \times g$ equal to δ, the results of its analysis would have changed markedly. We can recalculate DOT's cost-benefit analysis in two steps for a value of $\theta_{\mathrm{VSL}} \times g$ equal to δ, using the 3% rate for δ that DOT applied.[196] First, in order to translate DOT's $8.86 million VSL for 2010 into a VSL for 2018 (to match the year that the costs are accrued), the VSL would be increased yearly by 3%, rather, than 1.07%. This yields a VSL in 2018 of $8.86 million $\times (1.03)^8 =$ $11.2 million. That VSL is then applied to the 35.5 lives that the regulation is expected to save for each model year, without any further discounting. This yields expected benefits of $398.4 million. That is, even with respect to this relatively present-oriented auto safety regulation, increasing VSL over time to keep up with the discount rate would increase the expected benefits by $81.8 million, or approximately 25%. Adding this recalculated figure of $398.4 million for lives saved to the $79.4 million in non-life benefits estimated by DOT yields overall regulatory benefits of $477.8 million. That is substantially closer to DOT's projection of $556.7 million in costs.

As it turns out, DOT's $556.7 million cost figure may well have been wide of the mark too. The agency estimated that the cost per vehicle of installing rearview camera systems would be approximately $43 to $45 in cars that already have display screens and $132 to $142 in cars without such screens.[197] The agency estimated that only 5% to 7% of model year 2018 vehicles subject to the regulation would have display screens. In fact, the figure ended up being close to 99%.[198] To be sure, the promulgation of the rear visibility rule likely accelerated the adoption of display screens. But if DOT had estimated that even half of all model year 2018 vehicles would have had display screens regardless of the rule, the cost would have fallen to around $380 million.[199]

All of this is to say: If DOT had set the VSL growth rate equal to the discount rate, the rearview camera regulation would have looked much better in cost-benefit terms. The controversy that surrounded the regulation's significant negative net benefits might have been substantially more muted. While DOT's devaluation of future lives did not—on its own—tip the cost-benefit outcome, it came close. And if DOT had set the VSL growth rate equal to the discount rate and made different defensible assumptions about display screen adoption, the rear visibility rule would have passed a cost-benefit test by a comfortable margin.

2. *Particulate Matter Standards Revised*

---

[196] Alternatively, we could have reduced δ to be in line with a better estimate of the VSL growth rate ($\theta_{\mathrm{VSL}} \times g$). Whether we raise the VSL growth rate to match the discount rate or lower the discount rate to match the VSL growth rate will not affect the net present value of future lives, though it will have implications for future non-mortality benefits and costs.
[197] Final Regulatory Impact Analysis, *supra* note 186, at 83 tbl.VI-10.
[198] *See* Hemel, *supra* note 1, at 689.
[199] *See id.*

In January 2023, the EPA announced that it would be updating the National Ambient Air Quality Standards related to fine particulate matter, sometimes known as soot.[200] Particulate matter can be an especially deadly air pollutant: when breathed in, it can cause heart disease and cancer, leading to substantial numbers of premature deaths.[201] EPA estimated that tightening the standards for particulate matter emissions would save hundreds or even thousands of lives.[202] Indeed, avoided premature deaths accounted for over 98% of the monetized benefits of the regulation.[203]

Importantly, many of the premature deaths avoided by more stringent regulation will occur years into the future. Assuming the regulation remains in effect, it will reduce particulate matter emissions for years to come. And pollution avoided in a given year will mean fewer deaths ten or even twenty years later due to the time lag between when someone breathes in particulate matter and when it begins to affect their health.[204] Accordingly, EPA calculated expected net benefits and costs for the regulation for the 20-year period from 2032 to 2051.[205] In addition, the agency estimated the net benefits and costs for four regulatory options of varying stringency.[206] Particulate matter standards involve both an annual emissions standard—the maximum amount of particulate matter permitted in the air, averaged over a year—and a 24-hour standard—the maximum amount of particulate matter permitted in the air during any given 24-hour period.[207] The four options EPA considered differed in various ways along both of these two dimensions (the annual limitation and the 24-hour limitation).[208] Nonetheless, in the interest of clarity we will collapse the two dimensions and simply refer to the four options as Options 1,2, 3, and 4, in order of increasing stringency.[209]

---

[200] U.S. Environmental Protection Agency, Notice of Proposed Rulemaking, Reconsideration of the National Ambient Air Quality Standards for Particulate Matter, https://www.epa.gov/system/files/documents/2023-01/PM%20NAAQS%20NPRM%20-%20prepublication%20version%20for%20web.pdf.
[201] U.S. Environmental Protection Agency, EPA Proposes to Strengthen Air Quality Standards to Protect the Public from Harmful Effects of Soot, https://www.epa.gov/newsreleases/epa-proposes-strengthen-air-quality-standards-protect-public-harmful-effects-soot.
[202] U.S. Environmental Protection Agency, Regulatory Impact Analysis for the Proposed Reconsideration of the National Ambient Air Quality Standards for Particulate Matter, December 2022, https://www.epa.gov/system/files/documents/2023-01/naaqs-pm_ria_proposed_2022-12.pdf, at ES-17 tbl.ES-6.
[203] *Id.* at ES-18 tbl.ES-7.
[204] *E.g.,* U.S. Environmental Protection Agency Science Advisory Board, *Advisory Council on Clean Air Compliance Analysis Response to Agency Request on Cessation Lag*, December, 2004 (explaining official EPA policy that deaths caused by particulate matter will occur up to 20 years after the particulate matter is released into the atmosphere).
[205] *See* Regulatory Impact Analysis, *supra* note 202, at 8-7 tbl.8-4.
[206] *Id.*
[207] *Id.* at ES-1.
[208] *Id.* at ES-17 tbl.ES-6.
[209] For the interested reader, Option 1 corresponds to an annual standard of 10 micro-grams per cubic meter and a 24-hour standard of 35 micro-grams per cubic meter, Option 2 corresponds to an annual standard of 10 micro-grams per cubic meter and a 24-hour standard of 30 micro-grams per cubic meter, Option 3 corresponds to an annual standard of 9 micro-grams per cubic meter and a 24-hour standard of 35 micro-grams per cubic meter, and Option 4 corresponds to an annual standard of 8 micro-grams per cubic meter and a 24-hour standard of 35 micro-grams per cubic meter. *See id.* at 8-7 tbl.8-4 (listing the benefits of the various regulatory options).

In its cost-benefit analysis, EPA calculated the following costs and benefits for the four regulatory options, in 2017 dollars, discounted to the year 2022 based on a 3% discount rate:

Table 2: EPA-Calculated Costs and Benefits of Particulate Matter Regulation (billions 2017$)

	Option 1	Option 2	Option 3	Option 4
Benefits	$200	$220	$490	$1,100
Costs	$1.1	$2.9	$4.5	$21
Net Benefits	$198.9[210]	$217.1	$485.5	$1,079

As Table 2 indicates, EPA's own calculations predicted that Option 4 would generate the greatest net benefits, $593.5 billion more than Option 3. In addition, the majority view of the agency's Clean Air Scientific Advisory Committee was that the agency should adopt Option 4 or perhaps even a more stringent rule.[211] Notwithstanding those facts, the agency announced that it was proposing to adopt either Option 1 or Option 3.[212] The EPA did not explain why it favored the weaker standard, citing only "uncertainties" in the available evidence[213]—uncertainties that will be present any time an agency regulates.[214]

These EPA calculations are based on EPA's practice of increasing VSL at less than 1% per year, rather than setting the yearly increase in VSL equal to the discount rate, as we have argued is better supported by theory and evidence.[215] Accordingly, we recalculate the expected

---

[210] Our calculated net benefits do not precisely match EPA's figures due to rounding. *See id.* at 8-10, Table 8-7 (offering net benefit figures for two regulatory options). However, EPA does not report net benefits for two of the four regulatory options, and the net benefit figures it does offer are peculiar. For instance, EPA's reported figures indicate that $490 billion minus $4.5 billion equals $490 billion, which seems odd even as a matter of rounding error. *Id.* Accordingly, we will use the figures we have calculated by hand.

[211] Environmental Protection Agency, Reconsideration of the National Ambient Air Quality Standards for Particulate Matter, Jan. 5, 2023, at 14.

[212] *Id.* at 9.

[213] *Id.* at 14-17.

[214] *See* Jonathan S. Masur & Eric A. Posner, *Unquantified Benefits and the Problem of Regulation Under Uncertainty*, 102 CORNELL L. REV. 87 (2016) (discussing how agencies should handle the uncertainties that attend essentially all types of regulation).

[215] *See supra* Section III.A.2. Above, we reported that a 2024 EPA regulation implied annual inflation of the value of life of 0.627%. *See supra* note 7. However, the regulation we examine here predates that 2024 regulation, and that regulation's statement about VSL appears to be predominantly forward-looking. In order to estimate EPA's yearly increases to the VSL in 2022, we look to past EPA practice. EPA valued life at $4.8 million in 1990 (1990 dollars) and $7.4 million in 2006 (2006 dollars). *See* U.S. ENVIRONMENTAL PROTECTION AGENCY, GUIDELINES FOR PREPARING ECONOMIC ANALYSES 90 (Sept. 2000), https://www.epa.gov/sites/default/files/2017-09/documents/ee-0228c-07.pdf; EPA 2010 GUIDELINES, *supra* note 132, at B-1. Once inflation is taken into account using EPA's preferred implicit price deflator method, the VSL for 1990 is approximately $6.776 million (2006 dollars). *See Gross Domestic Product: Implicit Price Deflator*, FEDERAL RESERVE BANK OF ST. LOUIS: FRED ECONOMIC DATA, https://fred.stlouisfed.org/series/GDPDEF (last visited Oct. 3, 2022). Real VSL growth from $6.776 million in 1990 to $7.4 million in 2006 implies annual growth of approximately 0.55%.

costs and benefits of the regulation using 3% annual increases in VSL and a 3% discount rate.[216] That recalculation yields the following results:

Table 3: Recalculated Costs and Benefits of Particulate Matter Regulation (billions 2017$)[217]

	Option 1	Option 2	Option 3	Option 4
Benefits	$313	$353	$770	$1,702
Costs	$1.1	$2.9	$4.5	$21
Net Benefits	$311.9	$350.1	$765.5	$1,681
Percentage increase in net benefits compared with EPA calculation	56.8%	61.3%	57.7%	55.8%

This re-analysis does not alter the ordinal ranking of the four options. Option 4, the most stringent option, generated the greatest net benefits in EPA's calculations, and it generates the greatest net benefits in our recalculation as well. But our recalculation reveals that the net benefits from the most stringent regulatory option would have dwarfed the net benefits from the other options EPA was considering, by a margin of $915.5 billion, compared with the $593.5 billion that EPA calculated.

Ultimately, EPA adopted Option 3, promulgating a final rule in February 2024 that took effect two months later.[218] In the interim, twenty-four states and several industry groups sued EPA in the D.C. Circuit, challenging the rule's scientific basis and economic impact.[219] And after President Trump took office, EPA—in an unusual about-face—asked the D.C. Circuit to vacate the rule,[220] a request that the court has yet to grant. The challengers' argument—and the Trump EPA's—is that the rule is too stringent.[221] But proper valuation of future lives would suggest the

[216] EPA has helpfully provided year-by-year estimates of benefits, which enable us to recalculate them based upon the proper yearly increases to VSL. *See* Regulatory Impact Analysis, *supra* note 202, at 8-7 tbl.8-4. To recalculate the regulation, we inflate EPA's VSL an additional 2.4% each year. $1.0055 \times 1.024 \approx 1.03$, which is the discount rate used in these calculations.

[217] This calculation assumes that all of the benefits are from avoidance of premature deaths. That appears to be a safe assumption. *See* Joseph E. Aldy & Kip Viscusi, *Valuing Mortality Risk Reduction Benefits of Federal Regulations: Comments on Circular A-4*, 19 REV. ENVT'L ECON. & POL'Y 200, 200 (2025) (noting that "reduced mortality represented 99 percent . . . of annual benefits expected under full implementation of the [PM2.5 NAAQS] rule").

[218] Reconsideration of the National Ambient Air Quality Standards for Particulate Matter, 89 Fed. Reg. 16202 (Mar. 6, 2024).

[219] *See* Kentucky v. EPA, No. 24-1050 (D.C. Cir. filed Mar. 6, 2024); Lisa Friedman, *Nearly Half the States Sue E.P.A. Over New Limits on Deadly Pollution*, N.Y. TIMES (Mar. 6, 2024), https://www.nytimes.com/2024/03/06/climate/epa-soot-lawsuit.html.

[220] Respondent's Motion to Vacate and Remand, *Kentucky v. EPA*, No. 24-1050 (D.C. Cir. Nov. 24, 2025).

[221] The Trump EPA's decision not to defend this rule also dovetails with its choice to cease monetizing the mortality benefits of reducing particulate matter emissions. *See* U.S. Envtl. Prot. Agency, *supra* note 13, at 34.

opposite. By choosing Option 3 rather than Option 4, the agency appears to have left nearly $1 trillion in net social benefits on the table.

3. The Consequences of Undervaluing Future Lives

As the foregoing examples make clear, agencies' decisions regarding how to value future lives are not of merely academic interest. Rather, they likely exert significant impact on how agencies elect to regulate threats to health and safety. The question of how to value future lives is of particular importance because it is transsubstantive. A substantial proportion of agency regulations, across a wide variety of agencies, exist to save lives. How those lives are valued thus has ramifications for the entire regulatory apparatus of the administrative state. If agencies valued future lives more accurately, we would likely have safer cars and cleaner air, as well as cleaner water, safer workplaces, and greater protection from many other everyday hazards. By undervaluing them, agencies are leaving on the table life-saving regulations that could benefit future generations at an acceptable cost to present generations.

The result is a set of normative regulatory choices that many people would likely find undesirable. Agencies' decision to undervalue of future lives represents a preference for present-day Americans over future Americans. It thus represents a preference for the old over the young, and for the people who happen to live in the United States today over the people who will happen to live here in the future. These normative choices are not the product of democratic debate and decision-making. Instead, they are embedded in the nearly inscrutable technocratic choices of regulatory agencies. That is not how such weighty normative decisions should be undertaken.

C. The Political Economy of Future Lives

Once one understands the magnitude and normative effect of undervaluing future lives, the next natural question is *why* agencies have adopted such an approach. From the perspective of modern theories of administration, agencies' decisions to undervalue future lives present a puzzle. Canonical theories of administrative behavior predict that agencies will attempt to increase their importance and authority (and thus their budgets).[222] But the agency approaches we have canvassed do precisely the opposite. By devaluing future lives, they make regulations appear less justified and thereby reduce the number and scope of the regulations they can promulgate. The result is a diminished agency with less freedom of action (and potentially a smaller budget) than would obtain if the agency did not discount future lives so dramatically.

Why might agencies have adopted this approach? We cannot know for sure, but we can speculate. Based on the substance of agency decisions as well as the manner in which they describe and justify their choices, the agencies appear to be restraining themselves (or being restrained by

[222] *See* WILLIAM A. NISKANEN, JR., BUREAUCRACY AND REPRESENTATIVE Government 36-42 (1971); Ernest A. Young, *Executive Preemption*, 102 NW. U.L. REV. 869, 878 (2008) ("Federal agencies, after all, have no mandate to represent state interests and possess strong countervailing incentives to maximize their own power and jurisdiction."). *But see* Daryl J. Levinson, *Empire-Building Government in Constitutional Law*, 118 HARV. L. REV. 915, 932 (2005).

outside actors) precisely in order to appear more amenable to regulated parties. To begin with, it is notable that every agency's approach results in future lives being valued less than theory and evidence suggests they ought to be. If agencies were simply making random mistakes, we would expect some to overshoot and some to undershoot. But here the mistakes appear systematic, not random. In the same vein, the Department of Transportation refers to its own decision to set θ_{VSL} equal to 1 (in the face of theory and evidence to the contrary) as the "more moderate" choice—a distinctly political way of describing it.[223]

Why might agencies have trimmed their own sails? Agency approaches to valuing future lives have all the hallmarks of a bureaucracy staying its hand in order to avoid upsetting regulated entities. If administrative agencies were concerned that excessive regulation would turn powerful interests against them, and as a result they wanted to limit their own reach, they might well decide to take steps like these.[224]

If an agency's goal is to limit its own regulatory reach, why target $\theta_{\text{VSL}} \times g$? It may be that, in the context of cost-benefit analysis, those were the most readily available levers available for the agency to pull. When calculating the future benefits of a regulation, there are really only four parameters at play: VSL, the discount rate (δ), θ_{VSL}, and the growth rate (g). The discount rate is set by OMB, is consistent across agencies, and has remained constant (at 3% and 7%) for decades, except for a fleetingly short period in which the Biden administration sought to lower the rate to 2%. VSL is at least moderately politically salient: the choice of VSL sends signals to the public, and the news media writes articles for the general public about the VSL.[225] If an agency curbed its power by using a lower VSL, it might trigger unwanted negative attention. Adjusting θ_{VSL} and g, by contrast, allows agencies to limit the projected benefits of regulations while flying well under the radar.

This hypothesis is consistent with the fact that DOT and OSHA have increased their reported VSLs substantially in recent years. Those VSL figures are reported publicly—on a website accessible to the public (DOT) and in the proposed rulemaking, printed in the Federal Register (OSHA). If any number will be salient to the public, it is those VSL numbers. But as we

---

[223] DOT 2021 Guidance, *supra* note 110, at 9.
[224] Some agency decisions could be explained by the practical politics of the moment, though we tend to doubt these ad hoc explanations. For instance, perhaps DOT's decision to set θ_{VSL} equal to 1 in its 2013 guidance was undertaken in order to shield a recovering auto industry from the more stringent safety and fuel economy standards. But by 2021, the auto industry crisis was well in the past, and yet the agency did not adjust its approach, so it is hard to credit this decision as a move that was sensitive to the economic and political conditions following the Great Recession.
[225] Hemel, *supra* note 1, at 652-55; Dylan Matthews, *The Tricky Business of Putting a Dollar Value on a Human Life,* VOX (Dec. 22, 2022), https://www.vox.com/future-perfect/23449849/social-cost-carbon-value-statistical-life-epa. One counterexample is the Consumer Product Safety Commission's recent decision to reduce the VSL used for children. Previously, the Commission set the VSL for children under the age of 18 at twice the VSL for adults. In February 2026, the Commission announced that—going forward—it would use the same VSL for children and adults. *See* Consumer Product Safety Commission, Withdrawal of Final Guidance for Estimating Value per Statistical Life, 91 Fed. Reg. 8,845 (Feb. 24, 2026). The change appears to have elicited little public reaction.

have demonstrated, the outcome of a cost-benefit analysis depends significantly on VSLs calculated for years far into the future, and thus on δ, θ_{VSL}, and g as well as the VSL. It might be that some agencies have decided to publicly report higher present VSLs and then, behind the scenes, use artificially low values of θ_{VSL} and g to suppress the calculated benefits of regulations.

Accordingly, the obscurity of θ_{VSL} and g may be critical to the internal logic of agencies' decisions. The choices of θ_{VSL} and g would seem to be the apotheosis of technocratic decision-making. These decisions are described in esoteric guidance documents that are not even published in the Federal Register. Certainly, if one wanted to push policy in an anti-regulatory direction without eliciting backlash from the broader public, θ_{VSL} and g—not VSL—would be logical levers to pull. For their part, regulated parties are much more likely to notice the consequences of choices of θ_{VSL} or g—more or less stringent regulation—rather than the details of those choices. Thus, agencies may have seen manipulation of θ_{VSL} and g as the most effective way to placate regulated parties without raising alarms within the broader public. Better to tamp down the value of future lives (θ_{VSL} and g) than present ones (VSL).

Of course, as we have said, this is all speculation. We have no direct evidence as to why agencies have selected such low values for θ_{VSL} and g. Perhaps there is a more banal explanation—maybe it is simply the fact that agencies have used lower values for a long time and don't want to admit error by adjusting them, or maybe it is that agencies have reviewed the same evidence that we have and arrived at a very different conclusion. But when multiple agencies adopt unjustifiably low values for important regulatory parameters in a manner that looks like some type of political compromise, and then they describe that decision in political terms, it probably makes sense to listen to what they are saying.

CONCLUSION

The question of how to value future lives might seem esoteric and relatively unimportant, and it is—unless you care about climate change, or toxic air pollutants, or vehicle safety, or any of a wide variety of other regulatory issues. Whether agencies will promulgate regulations with large life-saving benefits far in the future depends substantially on the values of δ, θ_{VSL}, and g. As we have argued, there is a strong theoretical case for valuing future lives at least as much as present lives, such that 1000 lives saved twenty years from now should be treated as no less of a benefit to society than 1000 lives saved tomorrow. The empirical data that reads on this question is more mixed, but it still points generally in the same direction. And existing empirical studies, limited as they are, do not provide a compelling reason to deviate from strong theoretical and normative priors. Yet agency after agency across the administrative state has adopted parameters for the income elasticity of the VSL and the growth rate that generate very gradual increases in VSL, at a rate that lags the social discount rate by a substantial margin. And since this is not a new phenomenon, it affects not only regulatory decisions today with life-or-death consequences in the future, but also regulatory decisions in the past with life-or-death consequences today. It is hard to estimate the number of people who might still be alive had agencies placed appropriate weight on their existences and acted to protect them years in the past. But what is clear is that it is now well past time for agencies to assign the same value to future lives that we would have wanted agencies years ago to accord to our own.