

Deborah Lucas, MIT

## 1. Introduction

This article complements the other surveys in this special ARFE issue on the social discount rate by expanding on a few main themes. The first is that the insights of financial economics are generally applicable to investment choices both in the public and private sectors. Among its robust lessons, the idea that certain risks have a real and significant cost to individuals and society is emphasized because most policymakers and many economists fail to take the cost of risk into account in public choice settings.

While the logic behind the need to risk-adjust discount rates to evaluate benefits and costs involving intertemporal tradeoffs is well-established, how to best accomplish it in a government setting continues to be debated. The second theme developed here is that for a variety of conceptual and practical reasons, using market prices or fair value approximations to market prices is likely to be the best alternative for risk-adjusting public sector discount rates for most applications. The other leading approach for risk-adjusting public sector discount rates is a utility-based approach, as advocated by Cherbonnier and Gollier (C&G) (this volume). It is safe to say that the proponents of the different approaches agree that it would be better for governments to adopt any reasonable methodology over the status quo of abstracting from the cost of risk and ignoring the biases and inefficiencies that result. Hence, disagreements about which approach is preferred should not be taken as a justification for policymakers to forego risk-adjustment.

The third theme is that discounting over long horizons involves an unavoidably high degree of uncertainty that makes it advisable to seek alternative measures of value that are less sensitive to assumptions about the distant future. Most government investments (e.g., in hospitals, power plants, ports, private sector debt and bridges) involve time horizons that are similar to those of private sector investments, and hence are amenable to a financial economics approach. For policies whose impacts extend over many decades or even centuries--such as those aimed at reducing climate change, biodiversity loss, or other environmental harms--it is much less apparent that discounting projected cash flows or drawing inferences from historical data can provide answers about capitalized values that are dependable enough to be decision-relevant. Among other challenges, the choice of discount rates is complicated by philosophical considerations that are outside of the realm of financial economics such as the weight to place on the welfare of different individuals at a point in time and across generations. A possible alternative for evaluating the benefits of long-horizon policies is willingness to pay, which governments already sometimes employ. Relatively little appears to have been written about the limits of discounting for long-horizon policy analysis and the possible alternatives, and it is hoped the discussion here will encourage further consideration of this issue.<sup>1</sup>

---

<sup>1</sup> There are a number of useful theoretical results about the effect of consumption uncertainty on discount rates over long horizons (e.g., Gollier, 2010). The conclusion of those analyses has generally been that greater uncertainty implies using lower discount rates. However, the idea that uncertainty is a reason for policymakers to take those estimates less seriously is not addressed by those analyses.

## 2. The imperative for risk adjustment

Financial economics provides fundamental and robust insights into the determinants of economic value, and applying these principles to the choice of discount rates for policy evaluation has the potential to significantly improve social welfare. Many of the lessons in a standard corporate finance textbook—the relation between risk and required returns, the first-order irrelevance of capital structure to value, the term structure of interest rates, and the distinction between real and nominal interest rates—have relevance for the choice of discount rates in public sector applications (see, e.g., Brealey et. al., 2023).

Perhaps the most important lesson of financial economics for policymakers is that certain risks represent a real cost to society. Specifically, risks that cannot be eliminated by diversification—aggregate or market risk, or the additional risks borne by individuals when markets are incomplete—are costly. Risk-adjusting discount rates is often the most natural way to incorporate the cost of aggregate risk into policy analyses.<sup>2</sup> Because the amount of aggregate risk varies significantly across projects and policies, an immediate corollary is that it doesn't make sense to seek a single social discount rate that can be applied universally. Rather, the question is how to identify the appropriate discount rate(s) for the analysis at hand. While this viewpoint is fairly uncontroversial among financial economists (Gollier et. al., 2022), risk-adjustment is routinely neglected in most policy applications. The resulting misallocation of resources can be significant, as emphasized by Cherbonnier and Gollier (C&G) in this volume.

The intellectual foundations for risk-adjustment can be understood through the lens of Arrow-Debreu state prices, and that theory is also the starting point for the analyses in Caplin and Leahy (C&L) and C&G in this volume. The key insight is that the value of a unit of future consumption (relative to a unit of current consumption as numeraire) will vary with time horizon, the future state of the world, and the likelihood of that future state occurring. Importantly, future resources are worth more in states of the world where goods are scarce than in states where they're plentiful. Properly chosen discount rates are a means to incorporate Arrow-Debreu state prices into the valuation process.

The concept of state prices also provides the foundation for the methods that have been developed to value contingent claims such as government guarantees and insurance. Effectively, contingent claims valuation methods (which are also referred to as derivative or option pricing methods) sum together the relevant Arrow Debreu state prices to derive the value of any contingent claim. Employing a contingent claims approach is important for certain applications because it captures the greater value of policies that reduce the likelihood of very bad outcomes, and by the same token, the greater cost to the government when it absorbs the risk of very bad outcomes. Technically, because contingent claims entail leverage that varies over time and across states, no constant discount rate can produce accurate valuations for investments where contingencies are important. However, one can infer an implicit discount rate based on the value of a contingent claim and the distribution of its future payoffs. The implicit discount rate associated of claims that provide insurance against aggregate risk are often lower than the risk-free rate or even negative, a possibility not countenanced by policymakers that rely on a

---

<sup>2</sup> Government policies are often aimed at reducing risks that are costly but in principle diversifiable, such as the health impacts of toxic chemicals. The benefits of such policies (e.g., the value of expected lives saved each year) are most easily incorporated via projected cash flows rather than through a discount rate adjustment.

single rate for discounting. Notably, Merton (1977) was the first to emphasize the importance of taking a contingent claims approach for policy applications such as accurately valuing deposit insurance.

### 3. The case for inferring social discount rates from market prices

The two leading contenders for risk-adjustment are either to use comparable market rates, or to take a utility-based approach. This section reviews the conceptual and practical arguments for using market prices (or fair value estimates of market prices) for risk adjustment in most applications. The exception is when very long time-horizons are important, such as when evaluating the benefits of mitigating climate change. The challenges that arise with any discounting methodology in trying to estimate discounted values over long horizons are taken up in Section 4.

The discussion here is related to the analyses of C&G, C&L and Viscusi in this volume. While there is substantial agreement in many dimensions, the emphasis and some of the conclusions are quite different. C&G make the case for a utility-based approach implemented using a calibrated utility function and estimating aggregate consumption betas. A comparison of the arguments in C&G and in this survey should help readers figure out which line of reasoning they find more persuasive. C&L (this volume) also emphasize the importance of market prices. They demonstrate that using market-based hurdle rates for evaluating projects with safe cash flows often improves social welfare relative to using a below-market rate for discounting in situations where the government can also impose lumpsum transfers and taxes. Many of the conclusions drawn would generalize to valuing risky projects. However, C&L don't explicitly address risk-adjustment, nor do they consider any of the practical issues that are emphasized here pertaining to transparency, accountability, and consistency with the budget process. Viscusi (this volume) discusses a complementary set of practical considerations that involve legal constraints on how discount rates are chosen. Like C&L, Viscusi emphasizes some of the unintended consequences that can arise when the discount rates used for policy analysis are assumed to be lower than market rates.

A thumbnail version of the case for market prices over a utility-based approach is this: A commonly held view among economists is that market prices are the best available aggregators of information about preferences, scarcity, and expectations about the future. In fact, most resource allocation, both in the private and public sectors, is mediated by the market price system. There is no need for analysts to agree about preferences, whether people are rational, or to what extent markets are complete. Estimates are disciplined because market prices often are observable and auditable, and when they're not available, valuation experts can be called upon to verify the quality of analysts' estimates. By contrast, utility function-based inferences require assumptions about preference parameters, the weights to assign to different individuals or groups, the completeness of markets, and so forth. These are objects that are unfamiliar to policymakers and the public, and about which there is little consensus among economists. Model-based discount rates effectively delegate an important part of the decision-making process to analysts that have limited accountability. As one example of the sensitivity of the results to assumptions about uncertain quantities, the implied risk premium in a consumption-based model with a CRRA utility parameter of 5, when calibrated with consumption volatility set to 2.5% to reflect aggregate consumption risk, implies a risk-premium of 0.3%. Calibrating the same model with consumption volatility set to 10% to represent what is typical for individuals implies a risk premium of 5%.

Despite the wide acceptance of market prices for allocating resources in society in both the public and private sectors, policymakers and some economists continue to express reservations about using market-based discount rates for public sector valuations. Some of the most common objections to risk-adjusting government discount rates, and the counterarguments to them, are briefly summarized here. The discussion that follows draws on Lucas and Hong (2023), Lucas (2012 and 2014) and references therein.

### 3.1 Evaluating the government's cost of capital

A common misperception, and one that is widely reflected in public sector discounting practices, is that the government's cost of capital is its own borrowing rate. The basic principles of financial economics tell us that is an error—the cost of capital for an investment depends primarily on its cash flow characteristics, not on the entity that is funding it or by what means. It is physically impossible to fully finance an investment that entails undiversifiable risk with risk-free claims. Nevertheless, the idea that the government's cost of capital is its own borrowing rate seems to be confirmed by casual empirical observation: Governments are able to borrow at relatively low rates even when they make large and very risky investments.

Clearly the reason governments can borrow at low rates is because government debt-holders bear little default risk. Taxpayers and other government stakeholders are in a first-loss position for any risk that is not diversified away, i.e., undiversifiable or market risk. Losses are absorbed either by higher future taxes or reductions in future government services. Investment gains are channeled back to citizens through lower taxes or increased spending. A logical conclusion is that taxpayers and other government stakeholders function as conscripted equity holders in risky investments made by governments.

Government debt will be close to risk-free, at least in nominal terms, as long as the government has the ability to raise taxes and a credible commitment to repay the debt. Because of that safety, the interest rate on government debt isn't informative about the cost of the risk associated with any particular government investment or about the cost of the government's aggregate risk exposure. Put simply, using a discount rate inferred from the market price of a safe asset, (i.e., government debt), to assign value to a risky asset (i.e., a government investment or the benefits it provides) is like valuing apples at the price of oranges.

Discounting at government rates has the effect of treating taxpayer equity as requiring no compensation for undiversifiable risk. The practice systematically understates the full economic cost of risky government investments to society, or the full economic benefit of government actions that lower aggregate risk. The size of the resulting biases will differ with the policy being evaluated, with the largest errors arising for investments that entail the most undiversifiable risk.

### 3.2 All discount rates are inferred from market prices

Despite the above observations, some still question the relevance of market prices or returns for government discounting. However, there is really no alternative to relying on market prices; the choice is only over which market prices or rates will be used as points of reference.

Government interest rates are the most common choice for government discount rates. If market rates cannot be trusted to be reflective of social cost, neither can government interest rates, which are determined by supply and demand for government securities in capital markets. The specialness of

government debt casts doubt on the interpretation that its price reflects a pure rate of social time preference that aligns with policy objectives. Government bonds sell at high prices not only because default is unlikely, but also because of their liquidity and usefulness as collateral, and in instances for their preferential tax treatment. Government bond markets are dominated by large investors whose preferences may not be representative of those of policymakers or of the public at large. In sum, government interest rates reflect the value to investors of the particular characteristics of government debt and the participants in debt markets.

Utility-based models are also calibrated using market prices. The subjective rate of time preference is typically chosen so that the interest rates implied by the model match the average real interest rates on government securities. Other parameters, such as those determining the curvature of the utility function, are based at least in part on whether they are consistent with the observed equity premium or the total return on risky assets. Interestingly, the difficulty of explaining the level and cross-section of risky asset returns with utility-based models is one of the issues that causes some economists to trust market prices over utility-based models, and vice versa for others. Placing more confidence in market prices is consistent with the belief that preferences are reasonably well-reflected in market prices via the forces of supply and demand, but that preferences cannot be accurately characterized with a simple functional representation calibrated with aggregate consumption. Confidence in utility-based valuations may be supported by the observation that those models can fairly accurately reflect macroeconomic aggregates, and the belief that market returns are affected by forces that have limited relevance to social welfare.

### 3.3 Market prices may be inconsistent with social objectives

A philosophical objection to market-based discount rates is that they are likely to be inconsistent with social objectives. This observation in some sense has to be true, but its implications for the choice of discount rates is unclear. Any procedure for choosing social discount rates involves an implicit weighting scheme across individuals and across generations. While it is true that project evaluations based on the implicit weights arising from market prices may lead to conclusions that are inconsistent with a particular social objective, it is implausible that any fixed discounting scheme could consistently do better.

The most obvious hurdle to trying to align discount rates directly to social objectives is the difficulty of identifying and agreeing upon those objectives. Policy goals change over time and with the political landscape. Objectives often appear to be inconsistent across policies even in the same place and at the same time. Although economists have long debated how social welfare should be evaluated, choosing the criteria that policymakers should use to assign weights to individuals within and across generations, and across political jurisdictions, is likely to remain an issue that will be largely decided in the legal and political sphere.

The most common reservation about market prices is that they underrepresent certain groups. Although the choices of current market participants are influenced by altruism to future generations, non-citizens, the young, and the poor, policymakers may want to give greater standing to the preferences of those groups. However, as emphasized by C&L (this volume), there are many situations where redistribution within and across generations can offset unintended distributional consequences more efficiently than choosing an alternative to market rates for discounting. Viscusi (this volume) also points out that there are differing viewpoints among legal scholars about to what extent U.S. officials have the legal authority

or responsibility to consider costs and benefits beyond those affecting its current citizens. It is also notable that while economists seem to favor more equitable weighting schemes, revealed preference (e.g., via deficit spending policies) seems to suggest that policymakers and the electorate are comfortable shifting significant costs to future generations, and that the huge inequalities in wealth and income across and within countries suggests a high tolerance for inequality.

With regard to intergenerational tradeoffs, it is often difficult to determine even directionally whether the social rate of time preference or the associated risk premium would be higher or lower than that implied by markets. On the question of what is owed to future generations, there are at least two countervailing effects that affect the rate of social time preference. On the one hand, future generations are expected to be wealthier on a per capita basis, and hence equating the marginal utility of consumption of individuals across generations would entail transferring resources from the future to the present. That suggests an upward adjustment of social discount rates relative to market rates. On the other hand, it is expected that there will be many more people alive in the future. Under a utilitarian weighting scheme that seeks the greatest happiness for the greatest number of people, it may be optimal to impose large sacrifices on the relatively small current population in order to benefit a larger group in the future. That suggests a downward adjustment of social discount rates relative to market prices.

The many difficulties involved in identifying and incorporating social objectives into discount rates seems to favor the use of market rates. However, for policies involving very long horizons, distributional impacts are of first order and market rates are unavailable; we'll return to these issues in Section 4.

### 3.4 Practical considerations

Most public sector investment decisions that involve discounting are made inside government agencies or are based on estimates produced by them. Putting aside differences between public and private sector objectives, private firms and government agencies operate under very different legal and operational constraints. Those include that government agencies have more limited resources for project evaluation and less latitude in the choice of methodologies. Government salaries are more constrained and financial training is typically limited. Accountability is to elected officials and the public rather than to shareholders. There are different standards for transparency. Estimates of project value are especially important in the budget process, which is the mechanism through which legislators decide how to allocate limited public resources across competing uses. There is very limited ex-post information to assess whether the decisions that were made by policymakers on net added value to society or reduced it.

While these differences may seem orthogonal to questions about value, whether risk-adjustment should be incorporated into policy analyses using market rates or a utility-based approach must take into account these sorts of practical considerations.

#### 3.4.1 Putting cash and accrual estimates on a level playing field: the case for fair value cost estimation

In the U.S., the debate over whether to risk-adjust discount rates frequently arises in the context of certain budgetary costs such as the subsidies associated with government loans, loan guarantees or capital leases. By law, these obligations must be accounted for in the budget on an accrual (i.e., present value) basis using Treasury rates for discounting. Since the budget process is used to make tradeoffs

between different uses of public resources, an important objective is to represent accruals in a way that is most comparable with other government expenditures, which typically are accounted for on a cash basis.<sup>3</sup>

The reasons for risk-adjustment that are outlined above clearly apply to budgetary costs. However, in the context of making budgetary tradeoffs there is another useful concept to highlight, which is known as grant equivalence. Grant equivalence is the idea of assigning a cost to a government-provided service (e.g., guaranteeing a loan) that is equivalent to that of providing the program beneficiaries with an upfront cash grant that they could use to purchase the same service in the private market. It is consistent with the idea of using opportunity cost as the best measure of economic or social value. Grant-equivalent costs also have the interpretation of being subsidy costs because they represent the transfer of value from the government to program beneficiaries.

Grant-equivalence can be achieved by discounting at market or fair value rates because those rates reflect what it would cost the government to transfer the risk to the private sector. Discount rates derived from utility-based models could be given the same interpretation, but the connection to observed cash costs for other expenditures would be harder to explain.

At this point it is useful to clarify the distinction between market value and fair value. Accountants developed the concept of fair value to deal with situations where market prices are the desired metric for value. In well-functioning markets, fair values are calculated with reference to observed market prices or interest rates. However, sometimes no directly comparable market data are available. For example, the government may invest in projects that have no obvious counterpart in the private sector. Less frequently, market prices or rates are judged to be poor indicators of fundamental asset value, such as during periods of financial market distress. In such cases, the concept of fair value allows for the use of interpolation and model-based approximations. Such approximations represent a best estimate of what a price or discount rate would have been in a well-functioning financial market.

Adopting a fair value rather than a market value approach helps to address some of the objections that are made to relying on market prices for selecting government discount rates. It helps to ensure that cost estimates are not inflated by distress premiums, which some would argue are not legitimate government costs. It allows for model estimates to be used when comparable market prices are unavailable, as for the evaluation of very long-term projects. However, because most costs that receive accrual treatment are incurred over horizons that are much shorter than a generation, the fact that long-horizon market rates do not exist is rarely a practical impediment.

The issue of complexity is a further consideration. Superficially, it may seem that discounting at government rates is simplest because it avoids having to identify different discount rates for different programs. However, discounting projected cash flows at government rates creates a different price system for the government and the private sector. That makes it much harder to evaluate the accuracy of cost estimates because market prices for similar investments cannot be used as points of comparison. An advantage of adopting a fair value approach is that a robust private sector accounting and valuation infrastructure has been developed that serves to harmonize practices across reporting entities, disseminate best practices, and provide audit and valuation services. Those services are also available to

---

<sup>3</sup> The issue of valuing credit-related assistance comes up less frequently in other countries, many of which treat credit support as off-balance-sheet and outside of the budget process.

governments. The task of producing cost estimates on a risk-adjusted accrual basis can be made more manageable by standardizing procedures, investing in basic training for staff, and possibly centralizing or outsourcing the estimation function.<sup>4</sup>

A further important and practical reason for governments to use risk-adjusted market rates in budgeting is to avoid illusory arbitrage opportunities. The typical example used to make this point is to consider a situation where a government borrows to invest in stocks, and then books their value at a price based on discounting the expected returns at the government's borrowing rate. The expected equity premium is thereby treated as risk-free profit, and the government has created a money machine, at least on paper.

### 3.4.2 Other objections to market or fair value discount rates

It is sometimes argued that while a market risk premium that compensates for undiversifiable market risk may be a legitimate government cost, some of the other factors that affect observed market rates (e.g., liquidity premiums) are not relevant to the government. That raises the possibility of risk-adjusting discount rates, but not necessarily equating risk adjustment with the use of fair value procedures.

There are conceptual and practical reasons to avoid alternatives that require more subtle adjustments to discount rates. Conceptually, if competitive market prices are accepted as the best available indicators of economic value, and also from an opportunity cost perspective, such adjustments are unnecessary and would result in less accurate valuations. As a practical matter, there are no reliable procedures for isolating a pure market risk premium from observed rates and prices. The academic literature that has attempted to decompose market interest rates into their component parts—pure time value, expected losses, risk premium, taxes, liquidity—reports widely varying weights on those different elements. As emphasized earlier, a fair value approach imposes a well-established set of rules and procedures on valuations that provide discipline that would be lost were these other sorts of adjustments incorporated.

The issue of market incompleteness is also frequently mentioned as casting doubt on the adequacy of a fair value approach for estimating value from a government perspective. As discussed in Section 3.3, when markets are incomplete (as they undoubtedly are), purely private transactions at market prices need not result in an efficient allocation of resources, nor are market prices necessarily an accurate reflection of social value. Nevertheless, governments rely on market prices to measure the cost of most of their activities, and it is not clear how or in which direction discount rates should be adjusted to correct for the effects of incompleteness. Where adjusting for incompleteness becomes more important is in evaluating the benefits of certain government policies (e.g., reducing deaths from toxic substances). Generally, a financial economics approach suggests incorporating such benefits via an adjustment to cash flows rather than to discount rates.

## 4. The fragility of long-horizon discounting and alternatives to avoid it

The sensitivity of present values to small changes in interest rates over long horizons is well-known. For example, the present value of \$1 received in 100 years is 37 cents at a 1% discount rate, and 5 cents at a 3% discount rate. A seven-fold change in present value using two similar interest rates that are both

---

<sup>4</sup> The fair value estimates produced by the financial analysis division of the U.S. Congressional Budget Office provides an example of the feasibility of employing a fair value approach inside government, e.g., CBO (20??)



commonly used for long-horizon policy analysis casts doubt on the ability of discounted cash flow analyses to provide decision-relevant signals about the present value of distant outcomes. Over and above the uncertainty created by discount rates, long-horizon payoffs are also very uncertain.

In the private sector, the fragility of long-horizon discounting exercises is mitigated by the standard capital budgeting practice of cutting off explicit projections of cash flows beyond some relatively short horizon, usually taken to be 7 to 10 years. A terminal value is added to the final cash flow that represents the net value of all cash inflows and outflows beyond that date. Although conceptually that terminal value can be thought of as representing the present value of projected cash flows extending into the indefinite future, in practice that number is usually chosen with reference to a more tangible comparison point, such as the typical resale or liquidation value observed for similar projects.

It might be argued that to evaluate the benefits of long-horizon policies such as those aimed at reducing climate change, there is no alternative to discounting projected benefits over very long horizons. As one example, administrative law requires the use of benefit cost analysis, and an estimate of the social cost of carbon is a necessary benchmark against which to evaluate the net value of potential abatement projects or regulatory mandates. It was presumably with this in mind that Professor Michael Greenstone quipped that the social cost of carbon is the most important figure people have never heard of.

However, the requirement to estimate the value of long-horizon costs and benefits does not imply the use of discounted present values. The benefits of most government expenditures, such as investments in education or national defense or preserving national parks, are rarely explicitly quantified. Rather, benefits are implicitly determined by willingness to pay. That willingness is decided in the political process, often absent any quantitative estimate of the social value of the policy.

In the same way, commitments to holding the rise in global temperatures to set limits are presumably motivated by a sense of the very large damages from failing to do so, rather than by capitalized estimates of the harms. Nevertheless, the EPA relied in part on a discounted present value approach in its recent recommendation to increase the value of the social cost of carbon. The EPA estimate was influenced by the careful and detailed analysis of Rennert et. al. (2021), which incorporates risk adjustment through a utility-function approach calibrated with aggregate consumption. Despite the care taken, the discount rate used in their methodology is highly sensitive to their assumption about the distribution and volatility of future consumption growth and its correlation with climate change. As Pindyck (2022) and others have noted, there is disagreement among economists about the sign of that correlation, which depends whether aggregate growth or insurance value effects dominate. The sensitivity of this type of analysis to small changes in assumptions leaves the EPA open to the criticism that its proposed increase of the social cost of carbon is politically motivated. An alternative approach to setting the price of carbon, such as adjusting it dynamically in response to whether decarbonization is happening more or less quickly than a target pace consistent with willingness to pay, might be less susceptible to such criticisms.

Fortunately, for many policies whose benefits will be realized in the distant future, the related investment decisions often involve much shorter time horizons. For example, government investments to expand solar versus wind capacity, or to install scrubbers on coal plants, are made over horizons similar to those of most private sector investments. The costs are realized over the lifetime of most citizens, and comparable investments are made in the private sector. Hence, choosing between them

only requires assessing the relative costs of the different alternatives, and that can be accomplished using a standard fair value approach.

## References

Brealey, Richard, Stewart Myers, Franklin Allen and Alex Edmans (2023), *Principles of Corporate Finance*, 14<sup>th</sup> edition, McGraw Hill

Caplin, Andrew and John Leahy (this volume) “Missing Participants, Missing Markets, and the Social Discount Rate: Borrowing Constraints, Intergenerational Transfers, Altruism and the Desire for Legacy”

Cherbonnier, Frédéric and Christian Gollier (this volume) “Fixing our Public Discounting Systems”

Gollier, Christian, Frederick van der Ploeg and Jiakun Zheng (2022), “The Discounting Premium Puzzle: Survey evidence from professional economists,” Toulouse School of Economics Working Paper

Lucas, Deborah (2012). “Valuation of Government Policies and Projects,” *Annual Review of Financial Economics*, vol 4.

\_\_\_\_\_ (2014), “Evaluating the Cost of Government Credit Support: The OECD Context,” *Economic Policy*, Vol. 29, Issue 79, pp. 553-597

Lucas, Deborah and Gee Hee Hong (2023), “Evaluating the Costs of Government Credit Support Programs during COVID-19: International Evidence,” IMF Working Paper 23-14.

Merton, Robert (1977), “An analytic derivation of the cost of deposit insurance and loan guarantees: An application of modern option pricing theory,” *Journal of Banking and Finance*, Vol. 1, Issue 1.

Pindyck, Robert (2022), *Climate Future, Averting and Adapting to Climate Change*, Oxford University Press

Rennert, Prest, Pizer, Newell, Anthoff, Kingdon, Rennels, Cooke, Raftery, Ševčíková, and Errickson (2021), *The Social Cost of Carbon: Advances in Long-Term Probabilistic Projections of Population, GDP, Emissions, and Discount Rates*, Brookings Papers on Economic Activity

Viscusi, W. Kip (this volume) “The Social Discount Rate: Legal and Philosophical Underpinnings”