The Social Discount Rate

A Baseline Approach

James Broughel

January 2017

MERCATUS WORKING PAPER

MERCATUS CENTER
George Mason University
3434 Washington Blvd., 4th Floor, Arlington, Virginia 22201
www.mercatus.org

James Broughel. "The Social Discount Rate: A Baseline Approach." Mercatus Working Paper, Mercatus Center at George Mason University, Arlington, VA, January 2017.

Abstract

Economists discount future benefit and cost flows for a variety of reasons, including time preference, diminishing marginal utility of consumption, opportunity cost of capital, and risk aversion. Many of these rationales for discounting can be explained using the Ramsey equation found in neoclassical growth theory. This paper argues that Ramsey approaches to discounting are problematic for use in regulatory benefit-cost analysis (BCA) because they are inconsistent with certain foundational principles of BCA. A more useful discounting framework is one that is based on the time value of money, where discounting is used as a way to compare investment projects to a baseline alternative investment. A social discount rate (SDR) used in this manner avoids many ethics controversies that arise in Ramsey discounting approaches with respect to giving preferential treatment to the present generation over future generations, but it still recognizes and accounts for the importance of economic growth. An SDR of about 7 percent appears to be reasonable and is consistent with current guidelines from the Office of Management and Budget.

JEL codes: H43, D6, K2

Keywords: social discount rate, welfare economics, benefit-cost analysis

Author Affiliation and Contact Information

James Broughel
Research Fellow, State and Local Policy Project
Mercatus Center at George Mason University
jbroughel@mercatus.gmu.edu

Acknowledgments

For their helpful feedback on this paper, the author would like to thank Tyler Cowen, Jerry Ellig, Garett Jones, Ted Bolema, attendees at a February 2016 presentation for the Centre for Energy Policy and Economics at ETH Zürich, and two anonymous peer reviewers. Any remaining mistakes are the author's alone.

The Social Discount Rate: A Baseline Approach

James Broughel

1. Introduction

One of the most important decisions in public policy analysis is selection of the social discount rate (SDR). Economists apply an SDR in benefit-cost analysis (BCA) when the benefits and costs of social projects accrue across different time spans. The SDR is the interest rate used to calculate the present value of intertemporally distributed benefit and cost flows, so that these flows can be compared to one another as apples to apples.

A higher SDR will mean a lower present value of future benefit and cost flows. This effect is most profound for benefits and costs that occur in the distant future. Even minor adjustments to the SDR can have huge effects on present-value calculations because of the power of compounding. The more benefits and costs are separated by time, the more sensitive the sign of the net benefits calculation will be to selection of the SDR. Thus, the SDR will matter most in cases where (1) large upfront costs produce flows of benefits in the distant future or (2) large upfront benefits produce flows of costs in the distant future. Some of the most pressing issues of our time, including mitigation of global climate change and the growing national debt, have long-run intergenerational consequences. As a result, the SDR is closely tied to questions of ethics, such as how much society should care—and be willing to spend—for the welfare of future generations.

To understand the practical relevance of the selection of the SDR, table 1 presents estimates of the social cost of carbon (SCC) for the year 2020 using three SDRs. The SCC is an estimate of the monetized damages associated with the emission of an additional ton of carbon dioxide into the atmosphere. The numbers were calculated by averaging the output of three models, known as integrated assessment models, which estimate the effects of climate change. These estimates are inputs into the calculation of benefits of carbon dioxide reductions in government BCAs, where benefits are the avoided damages resulting from carbon dioxide emissions.

Table 1. Estimated Average Social Cost of Carbon Dioxide in 2020 at Various Social Discount Rates

Social discount rate (%)	Cost (2007$ per metric ton of carbon dioxide)
5.0	12
3.0	43
2.5	65

Source: Interagency Working Group on Social Cost of Carbon (2013).

As should be clear from the numbers presented in table 1, selection of the SDR is critical. At a discount rate of 2.5 percent, the SCC is more than five times higher than it is at a 5 percent discount rate. An SCC that is five times higher means it is efficient to spend five times more on climate change mitigation strategies. One reason for this large difference is that the integrated assessment models used to calculate the SCC run simulations out as far as the year 2300, and damages in the year 2300 will be extremely sensitive to selection of the SDR. In fact, the SCC is so sensitive to the SDR that raising the SDR by a mere 0.5 percentage points, from 2.5 to 3.0 percent, causes the SCC estimate to fall by one-third.

There is no consensus among economists as to what the appropriate SDR should be. Economist Martin Weitzman (2001) surveyed more than 2,100 economists and found a wide variation of opinion about the discount rate to be used for projects designed to mitigate global climate change. Responses to his survey ranged from negative discount rates to discount rates over 20 percent, with a mean response of about 4 percent. Even among only the most renowned scholars in the profession, Weitzman found disagreement, although less disagreement than across the larger sample of the profession. More recently, Drupp et al. (2015) surveyed economists about their preferred long-term discount rate and found considerable variation in opinion.

Most economists agree that the SDR should be above zero, although exceptions exist. The surveys of Drupp et al. (2015) and Weitzman (2001) reveal that a minority of economists favors a zero SDR (and sometimes even a negative SDR) for long-term social projects. Cowen (2007), for example, argues for a zero SDR as a means to protect the welfare of future generations. According to Cowen (2004), zero discounting is not necessarily a recipe for government intervention in the economy, but instead it would necessitate policies that foster a rapid rate of economic growth. In earlier work, Cowen and Parfit (1992) also show that discounting is not necessary to address many of the concerns cited by economists who say discounting is required.[1]

More commonly, economists argue for a positive SDR when analyzing potential social projects. Although the connection is not generally made explicit, many rationales for discounting in social project analysis appear to have been carried over from neoclassical models of optimal consumption behavior and economic growth. For example, the rationales for discounting provided by the federal government can be explained with a simple formula known as the Ramsey equation, which states that individual agents that optimize utility discount future consumption flows as a result of impatience and because consumption generates diminishing utility as consumption rises. These rationales form a large part of the theoretical case for discounting in social project analysis.

[1] Viscusi (2007) argues that certain anomalies can arise in economic analyses that do not use discounting. For example, total benefits and costs could approach infinite values as they extend into the infinite future, because costs could recur year after year in perpetuity or because willingness to pay tends to rise with income. However, limiting the time horizon is a much easier and more practical way to solve this problem. In fact, with an infinite time horizon, even using a positive discount rate produces an infinite flow of much smaller costs and benefits in net-present-value terms. Thus, a seeming anomaly in analysis can be resolved without discounting.

However, Ramsey discounting presents a problem for use in BCA for several reasons. As this paper will show, the logic underlying the equation contradicts basic tenets of BCA, such as the assumption that one additional dollar is valued equally by everyone and that a benefit's value should be determined by its recipients. Furthermore, ethics-based value judgments are often necessary to calibrate the parameters in the Ramsey equation, making disagreement among economists virtually inevitable.

In the sections that follow, this paper describes some of the problems associated with Ramsey discounting and why its use in BCA is troublesome. In place of Ramsey discounting, an alternative approach to discounting based on the time value of money is recommended. Here, discounting is used as a way to compare alternative projects to a baseline. This method gives no special treatment to the present generation over future generations, so the approach avoids many of the ethics controversies that can arise with Ramsey discounting. The discounting rule that comes closest to this time-value-of-money approach is known as the *weighted average approach*. This paper recommends a modified version of the weighted average rule with a pure rate of time preference set to zero, and it concludes with suggestions for how policymakers might update government guidelines on regulatory analysis in light of this new information.

2. Ramsey Rationales for Discounting

The current federal guidelines for regulatory analysis, presented in *Circular A-4* (OMB 2003), offer three rationales for discounting: (1) positive time preference, (2) diminishing marginal utility of consumption, and (3) opportunity cost of capital. Perhaps the easiest way to explain these rationales is with the Ramsey formula associated with the Ramsey–Cass–Koopmans economic growth model. The Office of Management and Budget (OMB) does not explicitly mention the formula in its guidance; however, the formula is a convenient vehicle for explaining the agency's reasoning, and many economists believe the Ramsey formula provides a useful framework for discounting (Arrow et al. 2012). The formula is written as

$$r(t) = \rho + \theta \frac{\dot{c}(t)}{c(t)}, \tag{1}$$

where the interest rate r at time t is equal to the representative agent's pure rate of time preference ρ plus the product of the consumption elasticity of marginal utility θ and the instantaneous growth rate of consumption $(\dot{c}(t))/(c(t))$ at time t. This equation is an equilibrium condition, where the marginal cost of using capital—the real interest rate r—equals the marginal physical product of capital, $f'(k)$, that is, the opportunity cost of capital, in equilibrium. The Ramsey rule recommends setting the SDR equal to r in this equation.

Let us consider these parameters one at a time as a basis for discounting. An SDR that is based on a positive pure rate of time preference parameter, ρ, assumes that utility of future citizens is worth less than utility of present citizens because society is impatient. People would rather consume today than in the future, and therefore they lose utility by waiting for future benefits to arrive. Since everyone in society is impatient, any aggregation method will necessarily generate a positive time-preference measure for society as a whole.

This reasoning requires that benefits and costs be valued from the perspective of those alive at the time a policy is implemented. True, individuals in the present are likely to be impatient when it comes to waiting for future benefits to arrive. However, individuals in the future will have no such impatience because the current generation's future will be their present. This problem is relevant across not only generations but also shorter time spans because people are continually being born and dying, and our future selves might be willing to pay different values for benefits than would our present selves. One of the greatest achievements of BCA is its ability to value benefits and costs in terms of what the individuals who receive those benefits and costs would be willing to pay for them (or pay to avoid them). Discounting on the basis of time preference violates this premise because it forces analysts to value benefits and costs in terms of how present members of society value them, rather than how individuals in the future—those who will receive the benefits or bear the costs—value them.

Aside from this problem, serious questions of ethics arise from discounting when it is based on time preference. With the-time preference rationale for discounting, the utility of people in the future is valued less than the utility of people today simply because of the passage of time. The British economist Frank P. Ramsey (1928, 543) himself suggests this is "ethically indefensible." Although individuals are impatient—something not disputed here—society as a whole does not share this impatience. Assuming the *social* discount rate behaves like an *individual's* discount rate is a fallacy of composition.

The parameter θ in the Ramsey equation describes how steeply marginal utility declines as consumption increases and it is also a measure of the representative agent's degree of risk aversion. Thus, it is sometimes referred to as the *coefficient of relative risk aversion*. Because future generations will presumably be richer than people alive today as a result of economic growth, the marginal utility generated by an additional unit of consumption is expected to be lower for people in the future than for those living now. The idea here is that analysts discount dollars rather than utility because dollars generate lower utility for richer citizens. Many environmentalists who have ethics concerns about discounting based on time preference still find this premise for discounting acceptable.[2]

[2] See Gollier (2013) for examples of how risk aversion and consumption smoothing across generations form a basis for discounting in policies that address global climate change.

Even with general agreement that there is diminishing marginal utility *within* individuals (ruling out the possibility of utility monsters with *increasing* marginal utility of consumption), neoclassical economists generally contend that interpersonal comparisons of utility *between* individuals are meaningless because utility is an ordinal concept. Economists can say that the second apple matters less to John than the first, but economists cannot say that the second apple to Susan matters less than the first apple to John. Utilities are rankings, and comparing one person's set of rankings to another person's provides little useful information. This concept has been acknowledged in neoclassical microeconomic theory going back at least to the British economist Lionel Robbins (1938).

Applying any value to the wealth effect term $\theta(\dot{c}(t))/(c(t))$ in the Ramsey equation is akin to applying distributional weights in cross-sectional BCA. Distributional weights are values assigned to benefits and costs accruing to certain subgroups in the population (e.g., the poor) to raise or lower the value of those benefits and costs in an analysis. Distributional weighting is a controversial practice for several reasons. It allows for potential Pareto improvements simply through the redistribution of wealth via pure transfers. This result leads to the conclusion that efficiency improvements are possible simply by equalizing the level of wealth across the population. Distributional weights are also hard to defend because they are inherently arbitrary and require singling out certain groups for special treatment.

In general, neoclassical economists do not weight benefits and costs that occur at the same time according to income. They should not do so across time, either, unless they are willing to make interpersonal comparisons of utility and single out groups for special treatment.[3] If important intertemporal distributional issues are raised in an analysis, a more transparent way to

[3] Note, however, that it is *real* dollars that concerns us here, so making adjustments to dollars based on changes in the price level across time is reasonable.

present this information is to present undiscounted flows of benefits and costs in a separate distributional analysis that highlights impacts on subpopulations of interest, such as the present and future generations. Then decision makers who are more accountable to the public can decide what is a fair and equitable intertemporal distribution of wealth.

Aggregation problems are also associated with discounting dollars based on θ. Recall that in the Ramsey–Cass–Koopmans growth model, the Ramsey equation describes the behavior of a single representative agent. To calculate how the marginal utility of society changes in response to wealth increases, every individual's preferences must be aggregated to form a set of social preferences. This calculation is easy enough in a growth model where everyone has identical preferences. In the real world, however, preferences are far from identical. Nobel laureate Kenneth Arrow (1950) proved that it is impossible to convert ranked individual preferences into a function describing community-wide aggregate social preferences without the possibility of certain paradoxes arising. There is no guarantee that, in the aggregate, people's preferences will be as well behaved as they are at the individual level.

Economist Stephen Marglin (1963, 109) describes the problem in more detail:

> There remains the problem of aggregating the time-preference maps of individuals for collective decisions into a single social time-preference map. This problem is a special case of the general problem of aggregating individual utility functions into a social welfare function. The more general problem has been investigated by Kenneth Arrow and others, and Arrow's negative conclusion that "democratic" aggregation is impossible unless we restrict the allowable class of individual preference functions or abandon one or more intuitively appealing axioms about preferences is too familiar to require elaboration.

What is the time-preference map Marglin speaks of? Elsewhere in the article he states that "the term 'time-preference map' should be carefully distinguished from the loosely used expression 'time preference.' The time-preference map refers to the entire functional relationship between, on the one hand, individual marginal rates of substitution of consumption at one time for consumption at another and, on the other hand, the levels of consumption at all different times" (Marglin 1963, 95–96).

Those familiar with the Ramsey–Cass–Koopmans growth model will know that $1/\theta$ is equal to the elasticity of substitution between consumption at any two points in time. So the time-preference map described by Marglin represents the marginal rates of substitution at all possible levels of consumption across time. By contrast, what Marglin calls the "loosely used expression 'time preference'" refers either to a particular marginal rate of substitution at a specific level of consumption or to the parameter ρ in the Ramsey equation.

The Arrow–Pratt coefficient of relative risk aversion,

$$\theta = \frac{-cu''(c)}{u'(c)}, \tag{2}$$

makes explicit that the coefficient θ is a function of utility (Simon and Blume 1994, 363). Because utility is a component of individual discount rates, the aggregation of individual discount rates to form an SDR can lead to the paradoxes identified by Arrow (1950). This result need not always occur, but it is a problem that cannot be ruled out. Indeed, one of the reasons for BCA's popularity is certainly that it replaced the need for aggregated social welfare functions that were unable to resolve the problems identified by Arrow.

These findings suggest that the appropriate value for θ might be zero in BCA. A value of zero for θ also implies that one additional dollar of consumption generates a constant and equal level of utility for all individuals. This result is intuitively pleasing for BCA analysts because this assumption—that the marginal dollar is equally valuable to everyone—is also a foundational assumption of BCA. Indeed, any value of θ other than zero creates inconsistency in BCA unless economists start applying distributional weights to benefit and cost flows.

Some economists address these aggregation problems by assuming a modified "social" version of the Ramsey rule, such as

$$r = \delta + \eta g, \tag{3}$$

where δ is society's rate of time preference, η is a measure of society's inequality and risk aversion, and g represents the growth rate of the economy. Gollier (2013) provides an excellent survey of these approaches.

Such an approach is sensible if we allow ourselves to think that society has a discount rate function and engages in optimizing risk and consumption-smoothing behavior just as individuals do. But if analysts are willing to go down this road, they might also be forced to consider whether BCA itself should be discarded and replaced with a social welfare function that policymakers seek to maximize. BCA has developed over time in part to avoid the subjective value judgments that analysts necessarily embed in the selection of a social welfare function.

Other economists recommend selecting an SDR on the basis of observable market interest rates.[4] These observable market interest rates could correspond to the left-hand side of the Ramsey equation (the opportunity cost of capital), allowing economists to forgo any consideration of the right-hand-side values. Acknowledging that capital has an opportunity cost means considering that it can have alternative uses, but this is also precisely what analysts seek to determine when they conduct BCA. They consider multiple alternative uses of public resources to identify where resources produce the highest social returns (i.e., the most efficient use of resources). As will be shown in the sections that follow, the cost of capital turns out to be more of a *method of obtaining* an SDR than a *rationale* for discounting. The SDR, when applied in this paper's recommended manner, focuses on only one alternative

---

[4] In the past, this approach was sometimes referred to as the *descriptive* approach to discounting. See Arrow et al. (1996).

use of resources, rather than all alternative uses. This single use may or may not represent the opportunity cost of resources.

3. Gamma Discounting

Before discussing the recommended discounting approach of this paper, it is worth mentioning a slightly different approach called *gamma discounting*. Gamma discounting was developed by Martin Weitzman (1998, 2001), who noticed that the distribution of economists' beliefs about the proper SDR (as measured by surveys) resembled a gamma probability distribution. He points out that the discount factor e^{-rt} can be viewed as a special case in a gamma distribution of the form $g(r) = \frac{\beta^{\alpha}}{\Gamma(\alpha)} r^{\alpha-1} e^{-\beta r}$. The first terms of the expression on the right-hand side of the equation represent probability weights to be applied to a set of uncertain discount factors. Because economists cannot agree on the appropriate discount rate to use in project analysis, each discount factor has a certain probability of being the correct discount factor. Critically, Weitzman (1998, 2001) recommends taking the expected value of the probability-weighted discount *factors* (as opposed to discount *rates*). Doing so generates a certainty-equivalent average discount rate that declines over the term structure toward the lowest possible discount rate. This result flows from Jensen's inequality, which applies to concave net present value functions.

Gollier (2004) responds to Weitzman by pointing out that if one instead takes the expected *future* value e^{rt} of social projects and evaluates payoffs in terms of net future value rather than net present value, one finds that uncertainty causes the discount rate to *rise* to its *highest* possible value. This puzzle, known as the Weitzman–Gollier puzzle, persisted until Gollier and Weitzman (2010) together showed that, after adjustments are made for the risk-aversion and consumption-optimization tendencies of individual agents, the discount rate declines over the time horizon, seeming to solve the contradiction.

However, if gamma discounting is to be applied to BCA, and if it is to avoid the Weitzman–Gollier puzzle, economists must again make certain assumptions about the risk-aversion and consumption-optimization tendencies of society. Whereas it is reasonable to make these assumptions about individual agents, especially in an optimal growth model, it is much less reasonable to assume the same about society as a whole. Again, the line between individual and societal preferences is blurred.

Even if economists are inclined to believe that society indeed has such tendencies as risk aversion, inequality aversion, and the like, they will likely never form a consensus as to how to calibrate these parameters in their models. Such calibrations depend more on ethical questions than on empirical ones, and given what is at stake, such as how much to invest in mitigating global climate change, there will likely never be consensus on the matter. Furthermore, gamma discounting tends to require survey results to calibrate the gamma distribution function, and survey results can be unreliable. First, who should be surveyed? As seen earlier, the distribution of SDRs preferred by a select group of economists in Weitzman's 2001 paper differed significantly from the broader profession's views. Further, given the ethical nature of the question, perhaps a broader range of professions beyond just economists should be included in any survey. There is also the problem of time inconsistency. If this year's survey suggests that 3 percent is the appropriate SDR but next year's respondents decide 7 percent is the appropriate rate, should next year's policymakers terminate all projects that started this year but fail a benefit-cost test next year?

In the following section, an alternative discounting approach that bases the SDR on the time value of money is presented. This approach produces intuitively pleasing results in that it is consistent for use in BCA and does not run into the same thorny aggregation and ethical controversies so common in Ramsey and gamma discounting approaches.

4. An Alternative Approach Based on the Time Value of Money

4.1. The Time Value of Money

The core reason for discounting future cash flows in finance is the *time value of money* (TVM). TVM means that income earned sooner is preferable to income earned later. A leading money and banking textbook puts it this way: “If you are promised $1 of cash flow, for certain, ten years from now, this dollar would not be as valuable to you as $1 is today because, if you had the $1 today, you could invest it and end up with more than $1 in ten years” (Mishkin 2016, 112). In other words, people prefer to receive money earlier rather than later because money can be used to generate even more money over time.

Note that TVM does not say that money today provides more utility than money in the future. Nor does it say that people are impatient, so utility matters less in the future. TVM says only that more money is preferred to less and that getting money sooner rather than later is preferable because it results in more money in total (and presumably, by extension, more utility) at the end of the period.

Analysts discount future cash flows according to TVM because there is an implicit alternative asset or account in which money can be invested. Often this investment is thought to be a risk-free asset,[5] although it could be any alternative investment instrument. The role of the risk-free asset here is critical. First, it is an *implicit alternative investment*. Although analysts generally do not explicitly consider the cash flows from this alternative investment (they only discount the cash flows of the investment under consideration), it is possible to do so. Thus, discounting cash flows is about comparing the cash flows from one investment to the cash flows generated by an implicit alternative investment. Discounting is a rule of thumb that makes comparison easier. If the net present value of an investment is negative, that means the internal rate of return on the implicit alternative asset (e.g., a risk-free asset) exceeds the internal rate of return on the investment being evaluated. A negative net present value only means the return is negative relative to the implicit alternative investment. It does not necessarily mean returns are negative relative to no investment.

Next, the risk-free asset is a *displaced investment* whenever a decision is made to embark on a new investment. Thus, the discount rate under TVM accounts for how resources would have been used if an investment had not been made. Putting money in the implicit alternative investment (e.g., a risk-free asset) is the *baseline* investment scenario. The baseline scenario is not zero investment because it would be foolish to give up free interest at no risk. Thus, discounting is a rule for comparing cash flows from alternative investment opportunities to a *baseline alternative investment scenario*. In this way, the practice of discounting can be thought of as an acknowledgment that the world being evaluated is not static. Failing to discount is to assume no investment in the absence of the investment under consideration.

[5] The rate of return on a risk-free asset is not the only discount rate used in financial analysis, of course. Financial analysts may use hurdle rates, the weighted average cost of capital, or the capital asset pricing model to identify a proper discount rate. Often the discount rate is raised above the rate of return on the risk-free asset to account for the riskiness of an investment (as is done in the capital asset pricing model). This kind of risk adjustment is also possible for an SDR. Alternatively, adjustments for risk could be made in an uncertainty analysis of benefits and costs in BCA.

Contrary to what is sometimes claimed, the opportunity cost of putting money in a financial investment is not necessarily investing in an implicit alternative asset like a risk-free asset. The definition of opportunity cost is the value of the *next-best* alternative forgone when undertaking an activity. The cash flows from a risk-free asset are only one alternative that should be considered when investing, and this alternative may or may not be the next-best alternative to the social project being considered. Thus, discounting under TVM only compares investments to one alternative—the most likely alternative. To determine the opportunity cost of an investment, a wide variety of alternatives must be considered.

4.2. Discounting as a Form of Baseline Analysis

One of the main differences between financial analysis of cash flows and BCA of social projects is that the two techniques compare investments against different baselines. Identifying the implicit alternative investment is relatively easy for financial analysts because market rates of return are usually available for risk-free assets or for other assets with comparable risk to the investment under consideration (i.e., whatever other asset is the *most likely* alternative investment). By contrast, the baseline scenario in BCA is the state of the world as it would have evolved in the absence of a social project. To know what this state of affairs looks like, one needs to know how resources would have been consumed and invested—and what the social returns on those uses would have been—in the absence of a policy. This is much harder to estimate than the baseline in financial analysis.

The discounting rule that perhaps comes closest to identifying the relationship between the baseline scenario and the SDR is known as the *weighted average* approach to discounting. Economists who endorse this approach divide resources displaced by social projects into two categories: resources that would have been consumed and resources that would have been invested. These economists use a consumption rate of discount based on the pure rate of time preference, and they use a higher investment rate of discount based on the cost of capital. The SDR is a weighted average of those two rates, weighted based on the mix of resources (consumption and investment) displaced by social projects. In other words, the SDR is weighted based on the sources of funding for social projects.

Harberger and Jenkins (2015) advocate this kind of approach. They note that when considering the weights to use in an SDR, “The profile of net benefits and costs that we analyze is really the difference between two moving pictures—one showing how the economy would evolve ‘with’ our project or program, and the other tracing a similar evolution ‘without’ it” (Harberger and Jenkins 2015, 8). Note the similarity of language between Harberger and Jenkins’s description of an SDR under the weighted average approach and OMB’s description of the baseline as “the best assessment of the way the world would look absent the proposed action” (OMB 2003, 15). The purpose of the SDR under the weighted average approach is to compare a world without a proposed regulation (i.e., the baseline) with a world in which the regulation is enacted.

Harberger and Jenkins (2015, 7) also note that their approach to discounting requires a “reinterpretation of the concept of opportunity cost.” In their assessment, “rather than thinking of the opportunity cost of public funds as their ‘best’ alternative yield, this reinterpretation looked upon it as their ‘likely’ alternative yield” (Harberger and Jenkins 2015, 7). This is why the rationale for using the opportunity cost of capital as a basis for discounting is misleading. The SDR does not compare an investment to its next-best alternative. Rather, it compares an investment to only one specific alternative, the most likely alternative, which is the baseline alternative.

The weighted average approach uses weights that approximate only how social projects displace consumption and investment flows *in general*. Therefore, an obvious problem with this approach is that resources will be used differently in different contexts. The returns on investments lost as a result of regulation will not be the same for every social project. However, it might be reasonable to think such returns are close to the marginal return on average. For example, if the expected return on an investment is significantly above market rates of return, firms can still borrow from capital markets to finance their investments, even after complying with a regulation.

It is therefore reasonable to assume that investments displaced by social projects are the ones on the margin, meaning they earn rates of return that are approximately equal to the cost of capital. No doubt there will be cases where government actions displace investments with higher or lower returns than the marginal rates—for example, when firms or individuals face credit constraints and cannot borrow capital to put in investments with high expected returns, or when businesses would have invested in projects that fail. Nonetheless, identifying a unique SDR for every social project—at least with present knowledge—will likely prove to be too difficult. The most practical way forward is to calculate an average rate of return on lost investments using market interest rates that reflect the cost of capital to firms.

Under the standard form of the weighted average approach, both consumption and investment are preferred sooner rather than later. On the one hand, producer surplus flows (i.e., profits) can be converted into even greater amounts of both producer and consumer surplus flows in the future, so these flows can have a compounding effect over time. Consumer surplus, on the other hand, is discounted because consumers are impatient, just as with the Ramsey rule.

Failure to discount producer surplus flows is a problem because it does not acknowledge the effect that compounding of lost investments has on economic growth. However, because consumption cannot be reinvested, there is no compounding effect over time. For reasons discussed earlier, it is reasonable to think that forgone consumer surplus flows should receive no special treatment on the basis of when the flows arrive. This would suggest setting the time preference parameter under a weighted average rule equal to zero.

This modified version of the weighted average approach produces some intuitively pleasing results. First, no special treatment is given to consumption on the basis of timing, so analysts are not suggesting anyone should die of cancer today just because Cleopatra enjoyed a second helping of dessert—to use an example from Cowen and Parfit (1992). On the other hand, the importance of economic growth is emphasized with the preference that producer profits arrive sooner rather than later. These intuitive results yield an approach that appears to address many of the concerns raised by Cowen and Parfit (1992), and later by Cowen (2007), related to economic growth and equity across generations. Thus, discounting to net present value, rather than being taken literally (e.g., implying that lives saved in the future are less important than the same benefits today), should be viewed as a rule for comparing returns on social projects to returns under a baseline scenario.

5. Implications for Regulatory Policy

The SDRs currently used by regulatory agencies in the United States are 7 percent and 3 percent. According to the Office of Management and Budget, which sets guidelines for regulatory analysis,

> The 7 percent rate is an estimate of the average before-tax rate of return to private capital in the U.S. economy. It is a broad measure that reflects the returns to real estate and small business capital as well as corporate capital. It approximates the opportunity cost of

> capital, and it is the appropriate discount rate whenever the main effect of a regulation is to displace or alter the use of capital in the private sector. (OMB 2003, 33)

Meanwhile, the lower 3 percent rate exists because

> The effects of regulation do not always fall exclusively or primarily on the allocation of capital. When regulation primarily and directly affects private consumption (e.g., through higher consumer prices for goods and services), a lower discount rate is appropriate. The alternative most often used is sometimes called the "social rate of time preference." (OMB 2003, 33)

What OMB describes in its guidance is something like the weighted average approach to discounting combined with the Ramsey rule (recall that OMB lists diminishing marginal utility of consumption as one of its three reasons for discounting). OMB is concerned with resources that are displaced by regulation, as is the weighted average approach. Displaced consumption is discounted at 3 percent, the pure rate of time preference, whereas displaced investments are discounted at 7 percent, the opportunity cost of capital. Any weighted average of OMB's recommended rates would have to be bounded by 3 percent and 7 percent, so any social project that realistically passes a benefit-cost test at the 7 percent rate should improve economic efficiency. Perhaps OMB thinks something like 4 percent is the rate at which society should discount based on diminishing marginal utility (if following the Ramsey rule).

This paper recommends a slightly modified version of OMB's current approach. First, OMB appears to at least partially embrace the Ramsey rule, given its embrace of diminishing marginal utility and time preference as a rationale for discounting. OMB should reject these rationales for discounting because they are not consistent with standard BCA assumptions. By extension, OMB should reject Ramsey approaches to discounting and be explicit that the weighted average approach is the agency's recommended approach.

OMB may also want to consider whether an update to the 7 percent base-case discount rate is appropriate. Harberger and Jenkins (2015) recommend a discount rate of about 8 percent for advanced countries. They arrive at that number through a weighted average formula with an additional component for the marginal cost of foreign funds. For example,

$$SDR = f_1 r + f_2 \rho + f_3 MCFF, \tag{4}$$

where f_1, f_2, and f_3 are weights to be applied to the opportunity cost of capital r, the pure rate of time preference ρ, and the marginal cost of foreign funds (which presumably differs from the marginal cost of domestic funds). The weights represent the sources of funding to pay for social projects. Using the Harberger and Jenkins estimates of 8.6–10.8 percent for r (which are notably higher than OMB's 7 percent rate) and 6 percent for the $MCFF$, as well as weights of 0.5–0.6 for f_1 and 0.3 for f_3, equation (4) is recalibrated by setting the pure rate of time preference parameter ρ equal to zero.[6] The result is an SDR in the range of 6.1–8.3 percent. Thus, the 7 percent base-case estimate of the SDR recommended by OMB is in line with the approach outlined in this paper. If OMB wants to continue to present a range of discount rates, a reasonable lower bound on the SDR would be zero, reflecting the extreme case where only consumption is displaced by social projects. A reasonable upper bound of 10 or 11 percent would reflect the opposite extreme, where only domestic investment is displaced by social projects. Those who want assurance that interventions are improving social welfare can be fairly confident that a proposed project that realistically passes a benefit-cost test at the higher end of discount rates will improve economic efficiency, although there might be reasons for using still higher SDRs, such as the irreversibility of many social projects.

[6] Harberger and Jenkins (2015) use time preference values in the range of 6.0–8.1 percent for advanced countries.

6. Conclusion

Economists both within and outside the federal government seem uncertain as to why they discount. At the very least, economists discount for a variety of reasons. The scope of issues that the SDR is used to resolve is astounding, considering it is just a single number. The SDR is simultaneously used to account for displaced consumption and investments that pay for social projects, the rate of time preference of society, the proper distribution of wealth across generations, society's degree of risk and inequality aversion, and more. Is it any wonder that economists cannot agree whether the appropriate number to account for all these things is 3 percent, 4 percent, or some other number?

This paper presents a case for abandoning discounting approaches based on the Ramsey rule and gamma discounting for social project analysis. This conclusion is consistent with the results of a recent survey that found that "the prominence of the simple Ramsey Rule needs to be revisited" (Drupp et al. 2015, 4). Still, these approaches are useful in other contexts, such as explaining individual behavior or the behavior of representative agents in optimal growth models. This paper claims only that the Ramsey and gamma discounting approaches pose problems for use in benefit-cost analyses of social projects. By extension, the use of declining discount rates in BCA is similarly problematic.

In place of these approaches, this paper recommends using an SDR that is based on the TVM. With this approach, the SDR serves one function only: to compare returns on social projects to returns under a baseline scenario. Therefore, the no-action alternative need not be explicitly considered in an analysis since the alternative of leaving resources in private hands is already implicitly considered through the practice of discounting.

The rule that comes closest to this approach is the weighted average rule. This paper recommends a modified weighted average rule that sets the pure rate of time preference at zero. Following this approach with recent data yields 7 percent as a sensible estimate of the SDR.

If and when OMB decides to update its guidelines on regulatory analysis, the agency should provide clarity on the reasons it believes discounting is necessary in regulatory analysis. The agency should reject Ramsey approaches and endorse the weighted average approach to discounting. OMB's current recommended base-case estimate of 7 percent is still reasonable, but the agency might want to consider recommending 0 percent as a lower bound on the SDR and 10–11 percent or higher as an upper bound. Such a range would address the inherent uncertainty surrounding the sources of funding for social projects.

References

Arrow, Kenneth J. 1950. "A Difficulty in the Concept of Social Welfare." *Journal of Political Economy* 58 (4): 328–46.

Arrow, Kenneth J., Maureen L. Cropper, Christian Gollier, Ben Groom, Geoffrey M. Heal, Richard G. Newell, William D. Nordhaus, et al. 2012. "How Should Benefits and Costs Be Discounted in an Intergenerational Context? The Views of an Expert Panel." Discussion Paper, Resources for the Future, Washington, DC.

Arrow, Kenneth J., William R. Cline, Karl-Göran Mäler, Mohan Munasinghe, Raymond L. Squitieri, and Joseph E. Stiglitz. 1996. "Intertemporal Equity, Discounting, and Economic Efficiency." In *Climate Change 1995: Economic and Social Dimensions of Climate Change,* edited by James P. Bruce, Hoesung Lee, and Erik F. Haites, 125–44. New York: Cambridge University Press.

Cowen, Tyler. 2004. "Policy Implications of Zero Discounting: An Exploration in Politics and Morality." *Social Philosophy and Policy* 21 (1): 121–40.

———. 2007. "Caring about the Distant Future: Why It Matters and What It Means." *University of Chicago Law Review* 74 (1): 5–40.

Cowen, Tyler, and Derek Parfit. 1992. "Against the Social Discount Rate." In *Justice between Age Groups and Generations*, edited by James S. Fishkin and Peter Laslett, 144–61. New Haven, CT: Yale University Press.

Drupp, Moritz, Mark Freeman, Ben Groom, and Frikk Nesje. 2015. "Discounting Disentangled." Working Paper No. 172, Grantham Research Institute on Climate Change and the Environment, London.

Gollier, Christian. 2004. "Maximizing the Expected Net Future Value as an Alternative Strategy to Gamma Discounting." *Finance Research Letters* 1 (2): 85–89.

———. 2013. *Pricing the Planet's Future: The Economics of Discounting in an Uncertain World*. Princeton, NJ: Princeton University Press.

Gollier, Christian, and Martin Weitzman. 2010. "How Should the Distant Future Be Discounted When Discount Rates Are Uncertain? *Economics Letters* 107 (3): 350–53.

Harberger, Arnold C., and Glenn P. Jenkins. 2015. "Musings on the Social Discount Rate." *Journal of Benefit-Cost Analysis* 6 (1): 6–32.

Interagency Working Group on Social Cost of Carbon. 2013. *Technical Support Document:—Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis—Under Executive Order 12866*.

Marglin, Stephen A. 1963. “The Social Rate of Discount and the Optimal Rate of Investment.” *Quarterly Journal of Economics* 77 (1): 95–111.

Mishkin, Frederic S. 2016. *The Economics of Money, Banking, and Financial Markets*, 11th ed. Essex, UK: Pearson Education.

OMB (Office of Management and Budget). 2003. *Circular A-4: Regulatory Analysis*. September 17.

Ramsey, F. P. 1928. “A Mathematical Theory of Saving.” *Economic Journal* 38 (152): 543–59.

Robbins, Lionel. 1938. “Interpersonal Comparisons of Utility: A Comment.” *Economic Journal* 48 (192): 635–41.

Simon, Carl P., and Lawrence Blume. 1994. *Mathematics for Economists*. New York: Norton.

Viscusi, W. Kip. 2007. “Rational Discounting for Regulatory Analysis.” *University of Chicago Law Review* 74 (1): 209–46.

Weitzman, Martin L. 1998. “Why the Far-Distant Future Should Be Discounted at Its Lowest Possible Rate.” *Journal of Environmental Economics and Management* 36 (3): 201–8.

———. 2001. “Gamma Discounting.” *American Economic Review* 91 (1): 260–71.