Rehabilitating the Opportunity Cost of Capital in Cost–Benefit Analysis

James Broughel

MERCATUS WORKING PAPER

All studies in the Mercatus Working Paper series have followed a rigorous process of academic evaluation, including (except where otherwise noted) at least one double-blind peer review. Working Papers present an author's provisional findings, which, upon further consideration and revision, are likely to be republished in an academic journal. The opinions expressed in Mercatus Working Papers are the authors' and do not represent official positions of the Mercatus Center or George Mason University.

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

James Broughel, "Rehabilitating the Opportunity Cost of Capital in Cost–Benefit Analysis," Mercatus Working Paper, Mercatus Center at George Mason University, Arlington, VA, September 2021.

Abstract
The opportunity cost of capital describes how invested resources would appreciate in value over time. Yet, currently, no consensus exists among economists as to how to account for this issue in a cost–benefit analysis, because disagreements are wrapped up in a broader, seemingly intractable, controversy surrounding the social discount rate. The social discount rate debate, in turn, forms part of a larger debate about what the normative aims of public policy should be. This paper provides background on these debates, explains how adherents of the two mainstream social discounting approaches each deal with the opportunity cost of capital, and describes limitations of both approaches, which are substantial. The paper argues that absent special circumstances, the conceptually sound approach to accounting for the opportunity cost of capital is using a shadow price, not a social discount rate. The paper reviews various methods of calculating this shadow price and discusses the policy implications of changing various parameters in the shadow price equation. The policy implications of cost–benefit analysis change dramatically depending on the social rate of time preference parameter selected in the shadow price equation, because at low rates, the convergence condition underlying cost–benefit analysis is violated. When this occurs, attention in analysis turns away from nonpecuniary benefits and costs and toward those benefits and costs exchanged in markets.

JEL codes*:* D04; D61; H43

Keywords: opportunity cost of capital; cost–benefit analysis; social discount rate; shadow price of capital; transversality condition

Author Affiliation and Contact Information
James Broughel
Senior Research Fellow
Mercatus Center at George Mason University
jbroughel@mercatus.gmu.edu

© 2021 by James Broughel and the Mercatus Center at George Mason University

Rehabilitating the Opportunity Cost of Capital in Cost–Benefit Analysis

James Broughel

1. Introduction

In a cost–benefit analysis (CBA), the *opportunity cost of capital* refers to how invested resources would appreciate in value over time. Yet economists have not reached a consensus on how to account for this issue in CBA. A long-standing debate in economics is about the appropriate method for selecting the social discount rate, and one of the primary ways in which the competing camps in the social discount rate controversy disagree is in their treatment of the opportunity cost of capital.

Views on social discounting are divided between two mainstream approaches, known as the *social opportunity cost* (SOC) method and the *social time preference* (STP) method (Spackman 2004). The SOC approach uses a social discount rate to account for the opportunity cost of capital. The STP approach uses a shadow price for this purpose, which is a conversion factor by which the market value of a capital asset is multiplied to convert the value of the capital asset into equivalent units of consumption. Under the STP approach, the social discount rate represents society's rate of time preference, or the rate at which society would trade present for future consumption. This time preference social discount rate goes into the calculation of the shadow price of capital assets, but it does not account for the opportunity cost of capital directly like the social discount rate does following the SOC approach.

Both approaches have their limitations, and as this paper will argue, both fail to completely address the issue of capital's opportunity cost in a comprehensive and

satisfactory manner. The SOC approach depends on extremely rigid conditions holding for the method to be appropriate—conditions unlikely to be met in most circumstances. In contrast, the STP approach in theory accounts for the opportunity cost of capital well, yet in practice its adherents tend to ignore this opportunity cost for reasons that are not compelling. The practical result of this state of affairs is that CBA does an inadequate job of accounting for the opportunity cost of capital regardless of which method is followed.

The disagreements between the SOC and STP proponents actually go beyond discounting, because the two methods have competing conceptions of what the goal of public policy should be. SOC adherents assert that they are measuring allocative efficiency, whereas STP proponents define the objective of policy as maximizing a particular social welfare function. Thus, disagreements about social discounting and opportunity cost are wrapped up in an even larger disagreement about what CBA should measure and what the normative aim of public policy should be.

This paper seeks to clarify the issues surrounding these disputes and is organized as follows. Section 2 provides a background on the social discount rate controversy in CBA, and argues that the appropriate way to account for the opportunity cost of capital in most situations is using a shadow price, as is the case with the STP approach. Section 3 presents three methods of estimating the shadow price of capital, drawing from the literature on social discounting, and describes how the different methods can explain the evolution of investments that take place in market versus nonmarket settings. Section 4 explains the policy implications of selecting a low social rate of time preference, which can result in violations of the convergence condition in the shadow price of capital equation. When this happens, the policy implications of CBA change dramatically, but the task of producing a

CBA becomes administratively simpler, because the focus of CBA turns to those benefits and costs traded in markets. Section 5 concludes with recommendations for the Biden administration as it looks to update existing guidelines on regulatory economic analysis.

2. The Social Discount Rate Controversy

In a CBA, a dollar's worth of investment should generally receive more weight in the analysis than a dollar's worth of consumption, for the simple reason that investment has a higher opportunity cost than an equivalent dollar value of consumption. This higher opportunity cost is usually accounted for in one of two primary ways, both of which relate either directly or indirectly to the social discount rate: the SOC approach and the STP approach.

Under the SOC approach, the social discount rate accounts for the opportunity cost of capital directly. Following this method, the discount rate in CBA is a weighted average of the rate of return on a unit of capital invested in the economy and the rate at which society is willing to trade current for future consumption (known as the *social rate of time preference*). This weighted average interest rate can be thought of as describing a counterfactual rate of return on an investment that is forgone when a project is undertaken. In the words of Arnold Harberger and Glenn Jenkins, "the profile of net benefits and costs [analyzed under the SOC approach] 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 & Jenkins, 2015, p. 8).

The SOC discount rate plays a role similar to the role a financial discount rate plays when discounting cash flows in that it is a device used to compare investments under

consideration with a counterfactual investment scenario that would occur in the absence of any portfolio change. In cash flow analysis, this forgone rate of return—the opportunity cost of capital—is accounted for with a discount rate.

Adherents of the SOC approach also claim to be measuring allocative efficiency. Thus, an important role of the social discount rate in their framework is to identify whether a particular project can return individual agents in the economy to their pre-project level of lifetime utility, thereby satisfying the Kaldor–Hicks potential compensation test. For the SOC discount rate to serve this role, however, a number of fairly strict assumptions must be met (Sjaastad & Wisecarver, 1977).

The first assumption is that benefits are "just like cash" (Burgess, 2018, p. 11). That is, the SOC approach assumes that benefits will accumulate in value at the same rate of return as the rate of return on the counterfactual investment that is displaced by the project. This assumption is reasonable in financial analysis, where all benefits and costs come in the form of money. It might also be reasonable for some public projects, such as building tollbooths, which cost the government money to set up and also generate streams of monetary income in the future. The "just like cash" requirement is clearly not met when benefits and costs come in nonpecuniary form, as is the case with most health or environmental regulations. In such instances, benefits will almost never compound in value at a rate comparable to the rate of return on displaced capital investments.

A second key assumption underlying the SOC method is that no savings externalities cause agents' private rates of time preference and the social rate of time preference to diverge. The SOC approach assumes a single market interest rate can represent the private rates of time preference of all agents in the economy. This assumption is unrealistic for a

few reasons. First, future generations are unable to participate in present markets, so their preferences carry little or no weight in the decision the current generation makes about how much to save.[1] Second, heterogeneity of consumers (including in their respective financial constraints), externalities in financial markets, and disequilibrium ensure that even within the present generation, individuals' marginal rates of substitution for present versus future consumption are often going to diverge from any single market interest rate used in the SOC framework.

The STP social discounting method, by contrast, is a more generally applicable approach because it can be applied in any CBA, not just when projects have benefits that are like cash and when private and social rates of time preference are equivalent. The STP method sees the selection of the social rate of time preference as a normative input in analysis. In other words, the decision of how much weight to place on the future is a value judgment. Accordingly, what CBA measures is normative, which explains STP advocates' willingness to abandon allocative efficiency and the Kaldor–Hicks potential compensation test as the basis for CBA and to endorse a social welfare function approach instead.

The STP approach recognizes that the issue of the opportunity cost of capital and the issue of social time preference must usually be kept separate and dealt with independently. The SOC method conflates these two issues and attempts to deal with both using a single social discount rate, which in most cases is inappropriate (Feldstein, 1972). The SOC

[1] One way to view the SOC method is that it grants standing in analysis only to the current generation.

approach *can* be appropriate if its stringent assumptions are met, but usually some assumptions will not be met.

In fact, even advocates of the SOC approach will sometimes acknowledge that using a shadow price to account for the opportunity cost of capital is the more generally applicable approach compared with using a social discount rate. Peter Abelson, an advocate of the SOC approach, writes that he "agrees … on the principle of the [shadow price of capital] approach, but also that this approach is not practical for general use" (Abelson, 2020, p. 15). During a panel discussion at the 2019 annual meeting of the Society for Benefit–Cost Analysis, three panelists—one of whom was SOC-advocate Arnold Harberger—reached a consensus that using a shadow price is more generally correct (Broughel, 2020). David Burgess, yet another advocate of the SOC approach, acknowledges that an analyst should multiply costs by a "marginal cost of funds" factor "when benefits deviate from the 'just like cash' benchmark that is implicit in the SOC approach" (Burgess, 2018, p. 14). The marginal cost of funds factor is a shadow price factor applied to capital investments to convert them to their consumption equivalent. However, it is one that is applied in the first period of an analysis, rather than in the period in which a benefit or cost is delivered, on the basis of an assumption that agents in the economy will anticipate the arrival of the capital good and adjust behavior beforehand (Liu, 2003).

The STP approach is sometimes viewed as being too difficult to conduct in practice, and the SOC method has the advantage of being administratively simpler. For this reason, the SOC method might be best viewed as a rule of thumb. For example, the U.S. Office of Management and Budget's (OMB) "analytically preferred" approach is to convert benefits and costs to consumption equivalents using the shadow price of capital. The office

nonetheless cautions federal regulatory agencies that "shadow prices are not well established for the United States. Furthermore, the distribution of impacts from regulations on capital and consumption is not always well known. Consequently, any agency that wishes to tackle this challenging analytical task should check with OMB before proceeding" (OMB, 2003, p. 33).

The simplicity of the SOC approach likely explains why some governments have chosen to adopt it despite its producing theoretically unsound recommendations in many, if not most, instances. Parts of the Australian government use the SOC approach (Abelson, 2020), and the U.S. federal government recommends that federal agencies use both a 3 percent and a 7 percent social discount rate in their CBAs (OMB, 2003). The 7 percent rate can be thought of as accommodating the SOC approach.

SOC adherents' resistance to adopting the STP method probably stems from two factors. First, SOC proponents want to measure allocative efficiency, not the well-being of the benevolent social planner whose welfare function is the basis of the STP approach. Second, in practice many STP advocates argue that applying a shadow price to capital investment is unnecessary. As the preceding OMB statement makes clear, little consensus exists on the degree to which capital investment is displaced by government policies. The extent will depend on a variety of factors, including how much a policy raises interest rates, the elasticity of the supply of saving to changes in interest rates, and the degree to which resources marshaled by government projects previously lay idle, for example, because of unemployment. STP supporters often argue that so little investment is displaced by

government projects as to make the use of the shadow price of capital factor unnecessary.[2] Or they argue that government spending and regulations indirectly induce investments that offset returns from displaced investments,[3] which would also make use of the shadow price of capital unnecessary.

Harberger and Jenkins refer to those who would use a low social discount rate, reflecting society's time preference, without applying a shadow price factor to forgone investments as "enemies of sound economics" (Harberger & Jenkins, 2015, p. 9). However, for reasons outlined earlier, the SOC approach is far from sound itself.

OMB's recommended rates of 3 percent and 7 percent can be viewed as a compromise between the SOC and STP approaches. According to OMB, the 3 percent rate is meant to represent the social rate of time preference. One might reasonably conclude then that the 3 percent rate is used to capture the extreme situation whereby no investment is displaced or created by a government policy. If all benefits and costs come in the form of consumption, then one could discount them all at society's rate of time preference without converting any investments into their consumption equivalent. Meanwhile, the 7 percent rate—which OMB claims represents the opportunity cost of capital—might represent the

---

[2] For example, Moore, Boardman, and Vining (2013) argue that government projects are ultimately funded by taxes. Because consumption forms a larger share of the economy than investment, the authors assert that taxes primarily reduce consumption rather than investment. However, the argument overlooks that the marginal dollar spent likely displaces more investment than the average dollar. The argument may also implicitly assume something like Ricardian equivalence, whereby taxpayers increase current saving in anticipation of future higher taxes. Interestingly, Ricardian equivalence is inconsistent with some other aspects of the STP approach, such as the assumption that capital investments are not anticipated.

[3] Arrow (1966) makes this argument. However, Feldstein (1970) shows that the conditions under which induced investments exactly offset displaced investments hold only under special circumstances.

opposite extreme, a situation whereby the costs of a project displace investment dollar for dollar.

This view—that the 3 percent and the 7 percent rates represent opposite extremes—accords with recent draft guidelines for regulatory analysis from the U.S. Environmental Protection Agency (EPA), in which the agency states, "The use of the social opportunity cost of capital as the social discount rate requires a situation where investment is crowded out dollar-for-dollar by the investment costs of environmental policies. This is an unlikely outcome, but it can be useful for sensitivity analysis and special cases" (EPA, 2020, pp. 6–16). Similarly, an earlier version of the EPA guidelines from 2010 states, "In most cases the results of applying the more detailed 'shadow price of capital' approach will lie somewhere between the [net present value] estimates ignoring the opportunity costs of capital displacements and discounting all costs and benefits using these two alternative discount rates" (EPA, 2010, pp. 6–19).

This logic—that the true social discount rate might lie somewhere between OMB's recommended rates (or even below those rates) or, alternatively, that a shadow price approach would produce a net present value between the net present values calculated using these two social discount rates—is incorrect. In fact, analyses using the 7 percent social discount rate will generally underestimate the amount of investment returns displaced by government programs.

Consider, for example, a regulation that costs $1 million and provides a consumption benefit next year. The monetary value of the benefit could be discounted at the rate of return to capital, which is mathematically the same as evaluating projects in future value terms by assuming the $1 million in costs displaces investment dollar for

dollar and compounds in value for one full year. The problem with this approach lies in what happens beyond the one-year time horizon. Implicitly, an analysis conducted in this way assumes that the capital investment is cashed in and the proceeds are consumed next year, when the consumption benefit is delivered. This notion is equivalent to assuming that the shadow price factor applied to the capital asset's terminal value is 1. (Note that a shadow price factor of 1 is also the implicit factor STP advocates use when they conclude that the opportunity cost of capital can be safely ignored.) In reality, the displaced investment is likely to continue growing in value even beyond the date on which the benefit is delivered.

The preceding example illustrates why the opportunity of capital cannot be accounted for using a social discount rate in most cases. Whenever benefits and costs are heterogeneous—with some coming in the form of consumption and some in the form of investment—then these benefits and costs will evolve differently through time, and therefore the opportunity cost of each needs to be accounted for separately. A social discount rate cannot account for these differing opportunity costs because it is bluntly applied to all benefits and costs in an analysis, irrespective of whether they are evolving like consumption or investment.

Despite having a stronger conceptual approach, STP advocates generally ignore the opportunity cost of capital in practice by failing to convert investments into their consumption equivalents. Meanwhile, the SOC method is administratively simpler, but it depends on assumptions so strict that they are unlikely to hold in most real-world policy settings. Regardless of the approach chosen, therefore, the opportunity cost of capital is not addressed satisfactorily in a CBA.

3. Calculating the Shadow Price of Capital

The shadow price of capital method was developed by economists Stephen Marglin, David Bradford, and Robert Lind (Marglin, 1963; Bradford, 1975; Lind, 1982), and it works very similarly in nature to pricing a stock or a bond in that a capital asset can be priced by valuing the stream of future income that it is expected to produce. Rather than a dividend or coupon payment stream, however, here the relevant income stream is the stream of consumption that the capital asset generates.

The shadow price of capital approach conforms with various theories of capital that state that capital's value should lie not in the effort exerted to produce it in the past, but rather in what it will do in the future in terms of being able to satisfy consumer demands (Hayek, 1931; Kirzner, 1976). A capital asset's value is ultimately derived from the subjective value consumers place on the end products that the capital eventually produces, not from anything inherent in the capital itself.

One can begin this exercise of valuing a capital asset with the simple case where all of the returns to capital are consumed each period. In this scenario, the principal value of the capital wealth base never grows and so the capital asset simply produces an infinite stream of equal consumption payments, which can be valued according to the formula

$$\text{SPC} = \frac{\text{ROI}}{\text{SRTP}}, \tag{1}$$

which says that the shadow price of capital (SPC) is equal to the rate of return on investment (ROI), which is the marginal social rate of return to capital net of depreciation, divided by the social rate of time preference (SRTP). In this case, capital is valued like a perpetuity in financial analysis, and often this simple case is the primary way

the shadow price of capital method is explained (e.g., Burgess & Zerbe, 2013; EPA, 2020, pp. 6–14; Li and Pizer, 2021).

One likely reason for this simple model's popularity is that in the Ramsey growth model with a capital income tax, a rational consumer who owns this capital asset will want to consume all of the return each period, consistent with equation (1). This equilibrium assumes that the consumer is optimizing utility and the consumer's marginal rate of time preference is equal to the market interest rate. If these conditions do not hold, however—because of externalities in financial markets, because of disequilibrium, or because the agent is not perfectly rational—then some nonzero fraction of capital's return will likely be reinvested each period. Some reinvestment seems likely in a second-best world where market inefficiencies are common. Indeed, studies that examine the marginal propensity to consume out of income or wealth often find that most of the marginal dollar earned is invested, not consumed (Carroll et al., 2017).

Equation (1) can be thought of as appropriate for valuing returns from natural capital, however. Natural capital's returns are often ongoing, but they cannot be reinvested because they are nonpecuniary (i.e., they do not come in the form of money). For physical capital, or human or natural capital that generates pecuniary benefits because the returns are traded in markets, the more general case is going to be that some of the return to capital is reinvested, whereas the remaining portion is consumed. In that case, the shadow price of capital can be estimated according to the equation

$$\mathrm{SPC} = \sum_{t=0}^{\infty}(1-f)\mathrm{ROI}\,\frac{(1+f\ \mathrm{ROI})^t}{(1+\mathrm{SRTP})^t}. \qquad (2)$$

Here, f is the fraction of the return that is invested each period, so $(1 - f) \times \text{ROI}$ is what is consumed in the initial period out of the first period ROI.[4] The consumption stream that capital generates grows at a rate of $f \times \text{ROI}$ each period, which for simplicity we will denote as g.[5] The consumption stream is then discounted each period at the social rate of time preference, which for shorthand we will denote as r^*. Finally, the stream of consumption is aggregated across time periods (t), and the resulting shadow price of capital conversion factor can then be multiplied by investment benefits and costs in a CBA.

Figure 1 illustrates the time stream of consumption forgone when a dollar of capital investment is displaced by public policy. The displaced investment can be thought of as an asset bearing characteristics like the Crusonia plant concept in the writings of Frank Knight (1944). When returns to capital can be reinvested, as is the case in equation (2), the predominant feature of the consumption stream is its compounding nature over time. Returns grow as some of the return is reinvested back in the capital fund each period, a distinctive characteristic of assets with financial returns. Reinvestment grows the base of wealth, which becomes the source of a greater return to capital the next period, and so on.

---

[4] A simplifying assumption here is that the same fraction of the return is reinvested each period. This process might be thought of as realistic at the national level, where roughly the same proportion of gross domestic product is invested each year, or it might be realistic of the marginal dollar earned or lost for the foreseeable future.

[5] Here, g is also assumed to be a constant. However, g could change over time, depending on whether there are diminishing or increasing returns to capital or depending on whether f is changing.

Figure 1. The Time Stream of Consumption Forgone from a Dollar of Displaced Investment

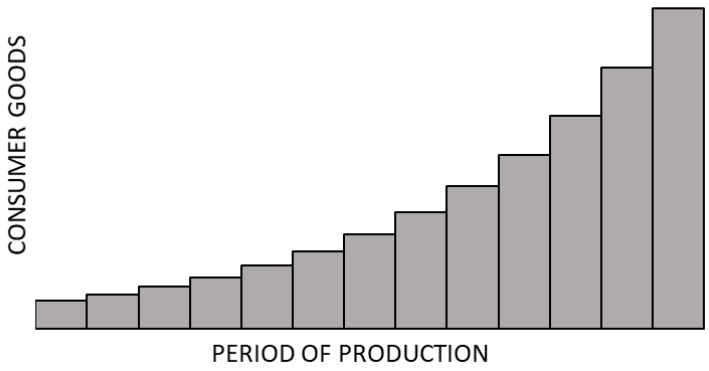

The variable r^* from equation (2) is a special case of a more general parameter, r, the market interest rate. The market interest rate should equal the social rate of time preference, or r^*, in the special case where the economy is operating along an optimal growth path that maximizes social welfare across time. The SOC method takes an implicitly intragenerational perspective since it accepts whatever the current market interest rate is as the social rate of time preference. Because future generations' preferences are neglected in current markets, however, STP advocates correctly note that, from an intergenerational perspective, r^* depends on one's normative values about how much to discount the future. For the sake of mathematical convenience, one could assume $r^* > g$, which is a convergence condition akin to the transversality condition from economic growth theory (Barro & Sala-i-Martin, 2004). If the convergence condition holds and one takes the limit of equation (2) as the time horizon extends to infinity, then equation (2) simplifies to

$$\text{SPC} = \frac{(1-f)\text{ROI}}{\text{SRTP}-f\ \text{ROI}} = \frac{(1-f)\text{ROI}}{r^*-g}. \tag{3}$$

Variants of this equation can be found in the writings of supporters of the shadow price of capital method, including Lind (1982); Moore, Boardman, and Vining (2013); and Boardman et al. (2018).[6] The formula is similar in nature to how a stock with a growing dividend would be priced. Equation (3) is a shorthand version of equation (2) that can be used when the shadow price of capital converges to a finite number. In cases where the convergence condition does not hold because $r^* \leq g$, equation (2) is the more general form that intuitively shows why the shadow price of capital is unbounded when r^* is less than or equal to g.

4. The Policy Implications of the Convergence Condition

If the analyst is unwilling to neglect the interests of future generations, then using market interest rates as the basis for r^* is unsatisfactory. However, it is not sufficient to simply assume that the convergence condition underlying CBA holds for the sake of mathematical convenience. For one thing, a very different set of policy prescriptions will be in order depending on whether the condition holds. When $r^* \leq g$, investing funds is always preferred to current consumption, and so society should seek to spur more investment at least until g falls below r^*. By contrast, when $r^* > g$, this preference will not always be the case. Then, present consumption is sometimes more desirable than investing to consume in the future. In

[6] SOC advocates sometimes criticize versions of equation (3) on the basis that it implicitly assumes that agents in the economy are myopic. The shadow price is based on a simple reinvestment rule that agents follow and is applied in the period in which the capital good is delivered. If Ricardian equivalence holds, consumers will fully anticipate changes in wealth in the future and will adjust behavior in the present. To address this issue when benefits are not like cash equivalents, the SOC approach uses a marginal cost of funds factor to convert capital goods to consumption equivalents in the first period of the analysis (Liu, 2003). Following this approach addresses the issue of benefits often being nonpecuniary, but it does not avoid the problems of the social rate of time preference diverging from private rates of time preference. Furthermore, it does not avoid the questions of convergence that are a focus of this section.

that scenario, the government can proceed with evaluating projects in a manner consistent with the traditional STP method, meaning that it should apply a shadow price factor to investments on the basis of equation (3) or something similar.[7]

When the convergence condition is violated, a problem that arises is that investment benefits and costs produce infinite consumption streams, and so the analyst has to figure out how to compare them. For example, preserving a wildlife species or a scenic view for all posterity could provide utility to members of society for the indefinite future, but paying for these benefits by displacing capital investment leads to an ongoing loss of consumption as well. As the previous section illustrated, however, one way to compare infinite consumption streams is to distinguish between benefits and costs that trade in markets from those that occur outside of market settings. Recall that policies earning nonmarket returns (i.e., returns like those in equation [1] that cannot be reinvested in financial markets) can produce consumption streams that are ongoing, but one would also expect that the value of those returns will eventually be overtaken by the consumption streams from assets that earn pecuniary returns. This result follows from the fact that financial returns can be reinvested, thereby producing the kind of compounding growth seen in Figure 1, whereas utility cannot be reinvested.

---

[7] Note that it is also possible when $r^*>g$ for there to be *dynamic inefficiency*, whereby society has excess capital and would want to reduce investment, since doing so could increase consumption in all periods. This situation occurs when ROI is less than the steady-state rate of growth of output, g. In equation (3), it occurs when the social rate of return to capital net of depreciation, ROI, is negative. The SPC in this case will be negative as well, implying excess capital. Note that this situation is different from the convergence condition being violated, which also produces a negative SPC in equation (3) but has opposite policy implications (i.e., society would want to increase investment rather than reduce it).

This discussion raises questions about when nonmarket returns can compound in value in a manner similar to financial assets. One example would likely be situations involving catastrophic risk. Catastrophic outcomes create significant challenges for CBA (Weitzman, 2011). This paper will not attempt to resolve these issues, except to note that they deserve careful consideration. For ordinary, run-of-the-mill policies, however, it seems reasonable to assume benefits and costs producing nonmarket returns will, in general, have a lower opportunity cost than benefits and costs whose returns are traded in markets, because the former lack the advantages of exchanges involving money. A policy implication, therefore, when the convergence condition is violated, might be that the government should turn its focus toward earning commercial rates of return. One way this might be accomplished is by setting up a sovereign wealth fund (Broughel, 2021).

The critical piece of information needed to answer the question of whether convergence takes place—and by extension to resolve the puzzle of what the corresponding policy implications are—is to know what r^* is—that is, what the social rate of time preference is. However, as already noted, selection of r^* generally depends on one's values (Feldstein, 1964, 1972). Therefore, which of the radically different policy conclusions are reached ultimately depends on value judgments.

Despite the ethical nature of the problem, several factors can still guide the selection of r^*. For example, Kaldor–Hicks efficiency is often said to form the normative foundation for CBA (Graham, 2008), and a trademark characteristic of Kaldor–Hicks efficiency is that it is insensitive to equity and distributional concerns (Adler, 2019). The SOC method takes an *intra*generational approach to Kaldor–Hicks efficiency, when it assumes that the market interest rate, reflecting the current generation's rate of time preference, is an acceptable r^*

value. However, one could also evaluate Kaldor–Hicks efficiency from an *inter*generational perspective. In other words, one need not accept the social welfare function that the STP proponents adopt, which is the welfare function of a benevolent social planner. One could instead accept a welfare function based on the economic principle of efficiency.

The benefit of this approach is that, as in the intragenerational case, it avoids certain value judgments about the desirability of how wealth is distributed. The focus is on whether wealth in the aggregate is increased, irrespective of its distribution. The social rate of time preference can be thought of as an intergenerational weight similar to a distributional weight that is applied within a time period in CBA. Unequal distributional weights are inconsistent with Kaldor–Hicks efficiency because they weight consumption differently depending on who receives it (Harberger, 1978).[8] Thus, it is conceivable that a social welfare function that corresponds with Kaldor–Hicks efficiency in an intergenerational context might have a social rate of time preference of zero.

The condition that $r = r^* = \text{ROI} = 0$ has an intuitive meaning in that it corresponds with an optimum whereby markets are complete and no more socially profitable investment opportunities are left. An additional dollar of investment yields a social rate of return of exactly zero, and therefore society is indifferent between consuming and investing the marginal dollar. At this point, the dollar value of social resources has been maximized across time, a modified golden rule rate of growth akin to the traditional golden rule in

[8] Sometimes, advocates of basing CBA on Kaldor–Hicks efficiency argue that discounting is appropriate because future generations will be richer than our own. However, these same advocates usually reject arguments for distributional weighting within a time period on the basis of income, so it is strange to accept such weights in an intergenerational context. Note that sound reasons may exist for using distributional weights at times. The point here is that such unequal weighting is inconsistent with Kaldor–Hicks efficiency.

economic growth theory.[9] Once this wealth-maximizing rate of growth has been achieved, then the task of policymakers can turn to whether consumer goods are allocated efficiently within the current time period, such that those who value the consumer goods most are also those who possess them. When the dynamic optimum across time and the static optimum within a time period are both reached, the economy can be said to be in a state of general equilibrium, whereby there are no further incentives to trade and the dollar value of all social wealth has been maximized.

Economist Tyler Cowen has argued for a zero social discount rate (Cowen, 2007, 2018),[10] and he has stated that a policy implication that follows from a zero social discount rate is to prioritize economic growth (Cowen, 2004, 2007). Cowen's position could likely be described as falling within the STP camp. What makes him unusual in the broader STP community is that he allows the convergence condition underlying CBA to be violated. Although many adherents of the STP approach are willing to dispense with shadow pricing investments altogether, essentially ignoring the effects of policy on capital investment, Cowen's position is something like the following: capital investment is so important it should be prioritized over virtually everything else policy affects.

---

[9] This paper will not address the myriad, challenging questions that arise when considering optimal population questions. Rather, it takes for granted that the relevant golden rule for policy purposes relates to maximizing consumption in the aggregate, not maximizing per capita consumption.

[10] Aside from Cowen, zero (or near-zero) discount rates are also consistent with the views of some other prominent economists. Martin Weitzman (2001), for example, surveys prominent economists, some of whom favor a zero social discount rate. See also Nicholas Stern (2006), who favors a zero discount rate on utility. Cowen (2007) allows for a small amount of discounting on the basis of the future being richer than the present. Yew-Kwang Ng (2003) argues for applying a small discount factor on the basis of a continuing possibility of the extinction of the human race. Note that little practical difference exists between a zero r^* and a near-zero r^* since both are likely to violate the convergence condition that underpins cost–benefit analysis.

At first glance, having the shadow price of capital be unbounded seems impractical. Administratively, however, it may be easier to evaluate projects under these conditions because in most cases all that needs to be accounted for are the benefits and costs that are traded in markets. These are the benefits and costs that already have prices associated with them, thereby avoiding one of the more challenging aspects of CBA, which is assigning dollar values to nonmarket goods.

5. Conclusion

This paper has explored the continuing challenges associated with accounting for the opportunity cost of capital in cost–benefit analysis. Challenges stem primarily from two factors: First, controversies surrounding the opportunity cost of capital are wrapped up in broader controversies about the social discount rate and the normative aims of public policy more generally. Second, a better accounting of opportunity cost may change the policy implications of CBA dramatically, and those implications may be hard for policymakers, as well as many economists, to accept.

The administrative simplicity of the SOC approach is its primary advantage; however, the stringent assumptions required to make its recommendations technically correct will almost never be met. Thus, net present value calculations using the SOC discount rate will often be misleading, or worse, meaningless.

This fact explains why there appears to be a modest consensus that using a shadow price is technically the correct way to account for capital's opportunity cost, even if some economists find this method to be impractical administratively. Even if one accepts the STP method, however, a further challenge is deciding what r^* should be. SOC adherents would

have us believe that markets are reasonably efficient, such that market interest rates are a good guide to determining the social rate of time preference. STP advocates typically do not find this logic compelling. They believe the selection of this rate ultimately comes down to one's values. Therefore, the choice of social welfare function becomes of critical importance.

If the analyst selects a relatively low social rate of time preference, then attention in CBA turns toward benefits and costs that are traded in markets: those with pecuniary, financial returns. Financial returns can be reinvested and generate compounding growth. Eventually, this growth will tend to overtake the benefits and costs from nonmarket investments, at least for ordinary projects that do not address catastrophic risks or other large-scale market failures. That is not to say that health and environmental benefits are not important. But with a low social rate of time preference, these benefits should be valued according to the returns they produce in markets. Indeed, some recent CBAs have even been structured in just this way, giving emphasis to pecuniary impacts (Broughel & Kotrous, 2021). A benefit of this approach is its relative simplicity.

Alternatively, one could assume a relatively high social rate of time preference, in which case the policy implications of CBA would not change that much. But this assumption raises ethical questions about whether it is appropriate to discount the future at a high rate. Moreover, a number of practical challenges arise, including identifying how much investment is displaced by particular projects, which is an area with very little consensus.

The Biden administration is now updating federal regulatory analysis guidelines and has made distributional concerns and intergenerational equity a centerpiece of its regulatory reform agenda (Biden, 2021). The Biden administration could take several

productive steps to better account for the opportunity cost of capital in CBA. First, it should make an effort to clarify the distinction between the *social rate of time preference* and the *opportunity cost of capital*, two concepts that are routinely conflated. In fact, it might be helpful to dispense altogether with the term *social discount rate*, which is commonly used to describe both of the previous two concepts. Second, analysts in the government should start accounting for the opportunity cost of capital using a shadow price. Doing so is especially critical given recent moves by the Biden administration toward using lower or declining social discount rates (U.S. Interagency Working Group, 2021). If these methods are adopted, analysts might end up ignoring the opportunity cost of capital altogether because these lower rates reflect social time preference rationales for discounting and not the opportunity cost of capital. This inattention to opportunity cost could have severe implications, including for the wellbeing of future generations.

Finally, the policy implications of a full accounting of the opportunity cost of capital should be explained. Unlike some other assumptions in economic analysis that are highly consequential in the context of particular policies, the issues surrounding the opportunity cost of capital are so important that they have ramifications that could affect government's role in the economy. Two policy implications, for example, of a low social rate of time preference are that society might want to extend the domain of market activity across as many areas of life as possible and that it might want to set up intergenerational savings vehicles, such as sovereign wealth funds, to spur growth and accrue wealth for future generations.

These somewhat unusual policy implications may explain why the opportunity cost of capital has never been dealt with adequately in CBA. Government activities often involve

curtailing market activity and capital investment to promote current consumption and activity in the nonmarket sector. CBA routinely says such projects pass muster. In this sense, CBA, as currently practiced, has essentially been constructed in a way to justify what the government is already doing. But this job is not for economists. Rather, their job is to identify trade-offs and to explain what is forgone in the process of choosing among competing ends for scarce resources. Whenever trade-offs arise, in other words, a good economist will always acknowledge the cost.

References

Abelson, P. (2020). A partial review of seven official guidelines for cost-benefit analysis. Journal of Benefit-Cost Analysis, 11, 272–293.

Adler, M. D. (2019). Measuring social welfare: An introduction. New York: Oxford University Press.

Arrow, K. J. (1966). Discounting and public investment criteria. In A. V. Kneese & S. C. Smith (Eds.), Water research (pp. 13–32). Baltimore: Johns Hopkins University Press.

Barro, R. J., & Sala-i-Martin, X. I. (2004). Economic growth (2nd ed.). Cambridge, MA: MIT Press.

Biden, J. R. Jr. (2021). Memorandum on modernizing regulatory review. White House, Washington, DC.

Boardman, A. E., Greenberg, D. H, Vining, A. R., & Weimer, D. L. (2018). Cost–benefit analysis: Concepts and practice (5th ed.). Cambridge, UK: Cambridge University Press.

Bradford, D. F. (1975). Constraints on government investment opportunities and the choice of discount rate. American Economic Review, 65, 887–899.

Broughel, J. (2020). Cost-benefit analysis as a failure to learn from the past. Journal of Private Enterprise, 35, 105–113.

———. (2021). Sovereign wealth funds: A potential solution to market failure and government failure. Economic Affairs, 41, 241–251.

Broughel, J., & Kotrous, M. (2021). The benefits of coronavirus suppression: A cost-benefit analysis of the response to the first wave of COVID-19 in the United States. PLoS ONE, 16, e0252729.

Burgess, D. F. (2018). The appropriate measure of the social discount rate and its role in the analysis of policies with long-run consequences [Paper presentation]. Mercatus Symposium on the Social Discount Rate, Mercatus Center at George Mason University, Arlington, VA.

Burgess, D. F., & Zerbe, R. O. (2013). The most appropriate discount rate. Journal of Benefit–Cost Analysis, 4, 391–400.

Carroll, C., Slacalek, J., Tokuoka, K., & White, M. N. (2017). The distribution of wealth and the marginal propensity to consume. Quantitative Economics, 8, 977–1020.

Cowen, T. (2004). Policy implications of zero discounting: An exploration in politics and morality. Social Philosophy and Policy, 21, 121–140.

———. (2007). Caring about the distant future: Why it matters and what it means. University of Chicago Law Review, 74, 5–40.

———. (2018). Stubborn attachments: A vision for a society of free, prosperous, and responsible individuals. San Francisco: Stripe Press.

Feldstein, M. S. (1964). The social time preference discount rate in cost-benefit analysis. Economic Journal, 74, 360–379.

———. (1970). Choice of technique in the public sector: A simplification. *Economic Journal,* 80, 985–90.

———. (1972). The inadequacy of weighted discount rates. In Layard R. E. (Ed.), Cost–benefit analysis: Selected reading (pp. 311–332). Baltimore: Penguin.

Graham, J. D. (2008). Saving lives through administrative law and economics. University of Pennsylvania Law Review, 157, 395.

Harberger, A. C. (1978). On the use of distributional weights in social cost-benefit analysis. Journal of Political Economy 86, S87–S120.

Harberger, A. C., & Jenkins, G. P. (2015). Musings on the social discount rate. Journal of Benefit–Cost Analysis, 6, 6–32.

Hayek, F. A. (1931). Prices and production. New York: Augustus M. Kelly.

Kirzner, I. (1976). Ludwig von Mises and the theory of capital and interest. In The foundations of modern Austrian economics (pp. 51–65). Mission, KS: Sheed & Ward.

Knight, F. (1944). Diminishing returns from investment. Journal of Political Economy, 52, 26–47.

Li, Q., & Pizer, W. (2021). Use of the consumption discount rate for public policy over the distant future. Journal of Environmental Economics and Management, 107, 102428.

Lind, R. C. (1982). A primer on the major issues relating to the discount rate for evaluating national energy options. In R. C. Lind, K. J. Arrow, G. R. Corey, P. Dasgupta, A. K. Sen, T. Stauffer, … & R. Wilson, R. (Eds.), Discounting for time and risk in energy policy (Vol. 3, pp. 21–94). Washington, DC: Resources for the Future.

Liu, L. (2003). A marginal cost of funds approach to multi-period project evaluation. Journal of Public Economics, 87, 1707–1718.

Marglin, S. A. (1963). The opportunity costs of public investment. Quarterly Journal of Economics, 77, 274–289.

Moore, M. A., Boardman, A. E., & Vining, A. R. (2013). More appropriate discounting: The rate of social time preference and the value of the social discount rate. Journal of Benefit–Cost Analysis, 4, 1–16.

Ng, Y-K. (2003). Appropriate discounting of future utilities need not be the dictatorship of the present: A note on Chichilnisky. Social Choice and Welfare, 21, 113–116.

Sjaastad, L. A., & Wisecarver, D. L. (1977). The social cost of public finance. Journal of Political Economy 85, 513–548.

Spackman, M. (2004). Time discounting and of the cost of capital in government. Fiscal Studies, 25, 467–518.

Stern, N. (2006). The economics of climate change: The Stern review. London: HM Treasury. http://www.lse.ac.uk/GranthamInstitute/publication/the-economics-of-climate-change-the-stern-review/.

U.S. Environmental Protection Agency. (2010). Guidelines for preparing economic analyses. Washington, DC.

———. (2020). Draft guidelines for preparing economic analyses. Washington, DC.

U.S. Interagency Working Group on Social Cost of Greenhouse Gases. (2021). Technical support document: Social cost of carbon, methane, and nitrous oxide interim estimates under Executive Order 13990. Washington, DC.

U.S. Office of Management and Budget. (2003). Circular A-4: Regulatory analysis. Washington, DC.

Weitzman, M. (2001). Gamma discounting. American Economic Review, 91, 260–271.

———. (2011). Fat-tailed uncertainty in the economics of catastrophic climate change. Review of Environmental Economics and Policy, 5, 275–292.