A Primer on Discounting, Climate Change, and Intergenerational Equity

Jon Bakija
Williams College
First draft: October, 2015; This draft: November 2018

Discounting enables us to compare benefits and costs that occur at different points in time. There are two different kinds of "discount rates," and each is useful for answering different questions. Whenever benefits and costs are separated by time, a *market discount rate* is what we need to answer the question "what is the economically efficient thing to do?" A *social discount rate,* on the other hand, is relevant to questions about equity, and is specifically used to answer the question: "what is the *ethical* value of a dollar of benefits or costs in the future, relative to a dollar of benefits or costs today?"

To illustrate the concept of discounting and why it is useful, we'll start with some examples. In the discussion below, assume that all dollar-valued benefits and costs have been adjusted for inflation so that we are measuring everything in real terms (e.g., constant year 2016 dollars), so that when we talk about "dollar-valued" benefits and costs, we are measuring them in such a way that the dollars have already been adjusted so as to have the same purchasing power now and in the future. Given that, the appropriate discount rates should also be measured in real terms, subtracting off the portion that is compensation for inflation.

Discounting and the Market Discount Rate

In order to determine whether a particular choice (among alternative methods of investing for the future) is economically efficient or not, it is necessary to discount future benefits and costs using a *market* discount rate, which is based on the rate of return (e.g. an interest rate) that could be earned on the best available alternative investment that has the same risk characteristics.

Example: discounting at the market interest rate when benefits and costs are separated by one year

Suppose we are making a choice about whether to do something now that would produce a benefit of \$110, measured in today's dollars, one year from now, and that this \$110 benefit is risk free (that is, there is a 100% chance that the benefit will be exactly \$110). Further suppose that the best available alternative risk-free investment is to save our money in a bank account that pays a 10 percent real risk-free interest rate (that is, the interest rate after subtracting off the portion that is compensation for inflation is 10 percent).

The intuition behind discounting at the market interest rate is that a real benefit of \$110 one year from now is worth less than a \$110 benefit today, because if we had the \$110 today we could save it and earn interest for a year, producing a benefit that is larger than \$110 next year. In order to adjust for this, we can "discount" the future benefit of \$110 to compute its "present value." The "present value" of a future amount of money is the amount we would have to put in the bank today at a given market discount rate (interest rate) in order to get that amount in the future. In this example, if we put \$100 in the bank today at a 10% interest rate, it would yield \$100×1.1 = \$110 one year from now. So the present value of \$110 one year from now at a discount rate of 10% is \$100. Mathematically, we can solve the problem this way:

(Present value)×(1+r) = (Future value)

In that equation, r is the market discount rate, which in this example is 0.1 (because 10% = 0.1). The information in the question told us the future value and the market discount rate, so we needed to solve for the present value. To do that, we need to divide both sides of the equation above by (1+r), which yields:

(Present value) = (Future value)/(1+r)

So, in this example, the present value of \$110 to be received one year from now is \$110/1.1 = \$100.

How would we use this information to decide whether a choice was economically efficient or not? If the choice in this example produces a benefit that has a present value of \$100, then it would be economically efficient to make this choice instead of choosing some alternative investment long as the present value of the cost is less than \$100, and it would be economically *in*efficient to accept this choice if the present value of the cost is greater than \$100. If the cost is exactly \$100, then there is no net gain or loss in economic efficiency, and if our goal is to do the efficient thing, we would be indifferent between making this choice or not. If, for example, the cost of the choice today is \$105, it would be economically inefficient to choose to do it, as the present value of the dollar-valued cost is greater than the present value of the dollar-valued benefit.

Another way of seeing why this makes sense is to compare future values. If the present value of the cost is higher than the present value of the benefit, then the future value of the cost will also be higher than the future value of the benefit, because to convert present values to future values, we multiply both the present value of the cost and the present value of the benefit by the same number. If we had saved the \$105 in a bank account at a 10% real interest rate, we would have had \$105×1.1 = \$115.5 one year from now. That is better than \$110 one year from now – we get an extra \$5.50 by saving the \$105 in the bank account at a 10% interest rate instead of making the choice that only produces \$110 in benefit one year from now. So we would be better off rejecting the choice that costs \$105 today and pays a benefit of \$110 one year from now.

Yet another way to think about this which always leads to the same result would be to compare the rates of return of the alternative choices. If incurring a cost of \$105 today only produces a benefit of \$110 next year, then the rate of return is:

$$rate\ of\ return = \frac{\$110 - \$105}{\$105} \times 100\% = 4.55\%$$

In the example in question, your best alternative with similar risk characteristics pays a rate of return of 10%, so you would be better off choosing the alternative with the 10% rate of return than making the choice that only pays a 4.55% rate of return.

An important point to recognize is that present value is fundamentally about opportunity cost and choosing the best among an array of choices for how one might invest one's saving, but it does not by itself tell us how much we should save as opposed to consuming today. For example, a decision to save a little more today might produce future benefits that exceed the present costs when discounted at the market discount rate. But in that case, one might still rationally decide that it is better not to do the saving and to consume today instead, because the utility gained from a little extra consumption today might be higher than the utility gained from a somewhat larger present value dollar gain in consumption in the future. If you're already saving a lot, so that consumption is expected to be higher in the future than it is today, and if there is diminishing marginal utility of consumption, then the fact that the present value of dollar benefits in the future exceed the present value of dollar benefits today must be weighed against the fact that the marginal utility of an additional dollar is lower in the future than it is today, in the scenario just described. If you are impatient, so that you value a util today more than a util in the future, then you would take that into account in your choice as well.

For example, you might voluntarily decide to reject an opportunity to save another $100 at a 10 percent rate of return, and instead consume that $100 today, and in that case it would be economically efficient for you to choose to do so. Implicitly, that would mean your willingness to pay for another dollar of consumption today in terms of consumption sacrificed next year is more than $1.10. The logic of economic efficiency *does* take into account differences in the marginal utility of a dollar across time and across states of the world (as in the case of insurance) when we are talking about voluntary choices made by an individual who bears the costs of and gets the benefits of the decision in question. But regardless of that, should you decide that you want to save, you'd definitely rather do it in an instrument that offers a 10 percent rate of return than in one that offers only a 4.55 percent rate of return, if they both have the same risk characteristics. The concept of present value discounted at the appropriate market interest rate, by itself, helps you decide which of those alternative investments you should choose.

Discounting by the market interest rate when benefits and costs are separated by more than one year

Now change the example so that everything is the same, except that we want to know the present value of $121 received two years from now. To compute the present value here, we need to take account of the fact that interest compounds. In other words, if we put $100 in the bank today at a 10% interest rate, we have $110 in the bank account one year from now, and then in the second year the 10% interest rate will be applied to both the initial $100 and the $10 of interest that has accumulated in the account. The amount in the bank account after two years will thus be $\$100\times1.1\times1.1 = \$121$.

More generally, here's how to compute the present value of a future amount, taking compound interest into account. If r is the annual discount rate expressed in decimal terms, and t is the number of years between now and when the future benefit and cost occurs, then:

$$(\text{Present value})\times(1+r)^t = (\text{Future value})$$

Or, if we divide both sides of the equation above by $(1+r)^t$:

$$(\text{Present value}) = (\text{Future value})/(1+r)^t$$

So in our example, the present value of \$121 received two years from now, evaluated at a 10% annual market discount rate, would be $\$121/(1.1)^2 = \$100$. That's because that's the amount we'd have to put in the bank today at a 10% interest rate to get \$121 in benefits two years from now. In this example, choosing to do something that produces a risk-free benefit of \$121 two years from now can only be economically efficient if the cost today is less than \$100 – otherwise, you'd always be better-off choosing the alternative investment that pays the 10% rate of return.

Risk and Discounting

So far, we've talked about "the" market discount rate, but in fact there are many different market discount rates, each corresponding to different risk characteristics of the choice being made. If people are risk averse, then the appropriate discount rate for determining the economic efficiency of a choice that exposes one to risk is *higher* than the risk-free interest rate – it must include a risk premium to compensate for the fact that risk-averse people dislike risk. Conversely, the appropriate discount rate for determining the economic efficiency of a choice that *reduces* risk (i.e., a choice that provides us with insurance) is *lower* than the risk-free interest rate – a risk-averse person should be willing to accept a rate of return lower than the risk-free rate in exchange for the benefits of being insured against other risks.

In fact, one can think of the decision to purchase insurance as a decision to invest in something that has an expected negative rate of return. The expected value of insurance benefit payouts is generally less than the price paid for insurance, due to the need for the insurance company to cover its administrative costs (including earning a normal economic profit). Yet people voluntarily buy insurance where the price paid exceeds the expected benefit payout, because it protects them from risk. This implies that the discount rate the person is implying to this decision is negative.

To demonstrate the point precisely, suppose that over the next year, your consumption will be \$90,000 with a probability of 90 percent, but there is a 10% chance of a catastrophe that would reduce your consumption to \$10,000 in the absence of insurance. Assume, for purposes of the example, that utility is equal to the square root of consumption, which implies diminishing marginal utility of consumption, and risk aversion, which are both the same thing. In this scenario, what is the maximum amount you would be willing to pay for one year's worth of full insurance, which offers to pay an $80,000 benefit in the event of catastrophe? The expected value in dollars of the insurance benefit in this example is $8,000, which is the 10% probability of the adverse event times the benefit payout of $80,000. Yet, because the insurance protects you from risk, you should be willing to pay more than $8,000 for this insurance. Finally, to keep things simple, assume the interest rate is zero.[1] If you do the math, you'd find the maximum amount you should be would be willing to pay for this insurance in this scenario is $11,600.[2]

So in this example, someone is willing to pay up to $11,600 for insurance that has an expected payout of only $8,000. If you did that, your rate of return on the investment in insurance would be:

$$\frac{\$8{,}000 - \$11{,}600}{\$11{,}600} \times 100\% = -31\%$$

That -31% is the appropriate risk-adjusted discount rate to apply to the cost-benefit analysis of whether this decision to purchase insurance is economically efficient or not in this example. As long as the rate of return on this investment in insurance is greater than -31%, then you make yourself better off by buying it. In a competitive insurance market, the rate of return on that investment is likely to be a lot better than -31%, but still negative. For example, if the administrative cost of providing the insurance is $1,000 per customer, and the insurance market is perfectly competitive, then the price of the insurance will equal the expected marginal cost,

---

[1] If the interest rate were positive, you'd want to take that into account in your decisions too, but in a competitive market the opportunity cost of lost interest that comes from investing your money in insurance instead of in an interest-bearing asset would be offset by the fact that the insurance company can earn interest in the meantime, which lowers its marginal cost. In competitive market that would reduce the price of the insurance.

[2] To see where the $11,600 comes from, first, calculate the person's expected utility without insurance:
$E(U)_{no\ insurance} = 0.1 \times \sqrt{\$10{,}000} + 0.9 \times \sqrt{\$90{,}000} = 280$
Next, convert that expected utility into "certainty equivalent consumption," or *C.E.C.*, which is the amount of consumption that, if you had it with certainty, would give you the same utility as the expected utility from the risky situation described above.
$E(U)_{no\ insurance} = \sqrt{C.E.C.}$
$280 = \sqrt{C.E.C.}$
$280^2 = C.E.C.$
$C.E.C. = \$78{,}400$
The insurance guarantees consumption of: $90,000 – (price of insurance). So the maximum price you should be willing to pay for the insurance (MWTP) would be the price that would reduce your guaranteed consumption with the insurance down to your certainty equivalent consumption of $78,400. So $90,000 – MWTP = $78,400. Solving for MWTP yields MWTP = $11,600.

$\$1,000 + 10\%\times\$8,000 = \$9,000$. In that case, the actual rate of return on the investment in insurance will be:

$$\frac{\$8,000 - \$9,000}{\$9,000} \times 100\% = -11.11\%$$

The actual rate of return of -11.11% offered to you by the competitive market here is higher than your discount rate of -31%, so you should be willing to buy the insurance.

Conversely, if you are presented with an investment opportunity that exposes you to greater risk than you would face otherwise, and if you are risk-averse, well-informed, and rational, then you should only be willing to accept that risk if you are sufficiently compensated for it through a higher expected rate of return. For example, consider corporate stocks, which are shares of ownership of a corporation that entitle the owner to a share of the company's profits. Corporate stocks are riskier than many other types of assets, and the riskiness causes the demand curve for stocks to shift left, pushing the price paid for a given stream of expected future profits down, and pushing the expected rate of return up. In equilibrium, the demand for stocks shifts just enough so that the expected rate of return is pushed up to the marginal investor's discount rate for stocks, where that discount rate is the sum of the risk-free rate of return and the risk premium the marginal investor requires to be willing to hold stocks.

The difference in the average long-run historical real rates of return on assets with different risk characteristics seems to suggest that people are very risk-averse. For example, between 1871 and 2012, the average annual real interest rate on long-term U.S. government bonds, a relatively low-risk asset, was 2.5 percent per year. Over the same period, the average annual real rate of return on the S&P 500 stock market index, including returns in the form of both dividends and capital gains, was 6.5 percent per year.[3]

Those different rates of return make a big difference over long periods of time. For example, suppose you are at the beginning of your career and you are starting to save for retirement 40 years from now. If you save $10,000 at a 2.5 percent rate of return, in 40 years that $10,000 will turn into $\$10,000\times(1.025)^{40} = \$26,851$. If you instead save the $10,000 at a 6.5 percent rate of return, the $10,000 will turn into $\$10,000\times(1.065)^{40} = \$124,161$ in 40 years. But that is just the *expected* or *on average* outcome. If people are rational and well-informed, then the whole reason the reward to saving in the form of stocks is so much higher than the reward for saving in the form of U.S. government bonds is that the stocks are riskier – if you save in stocks, the stock market could crash and you could end up with your $10,000 turning into something considerably *less* than $26,851 in 40 years. Given the historical experience of the stock market in the U.S., the probability of that kind of outcome seems very low, so the risk premium on stocks seems puzzlingly high. If taken at face value, it suggests that people must be *extremely* risk-

[3] Calculated by author using data from Robert Shiller's web site, at <http://www.econ.yale.edu/~shiller/data/ie_data.xls>.

averse. However, it might alternatively be explained by the fact that people think the historical riskiness of the stock market in the U.S. understates its true riskiness, for example because if you look at the global history of stock markets, we can find lots of examples of occasional, rare catastrophes that essentially wiped out almost the entire value of the stock market (e.g., the German stock market in World War II). Or it might be explained by people not fully understanding the stock market.[4]

In any event, the key point for our purposes is that the appropriate discount rate for an investment that increases your exposure to risk ought to be higher than the risk-free interest rate. To evaluate whether an investment that is as risky as the S&P 500 is economically efficient, the historical evidence above suggests that a real market discount rate of around 6.5% might be appropriate.

Application to Determining the Economically Efficient Pigouvian Tax

Burning fossil fuels releases carbon dioxide into the atmosphere. This in turn contributes to a greenhouse effect, resulting in an increase in the average global temperature that lasts a very long time. That is a clear example of a negative externality problem. Economist Arthur Pigou worked out the economically efficient solution to the problem in 1920, and you learned about it in introductory economics. To review, consider the diagram of the market for fossil fuels (such as oil natural gas, coal, etc.) shown below in Figure 1. Suppose we've normalized one unit of "fossil fuel" to be the quantity of fossil fuel that emits one ton of carbon dioxide into the atmosphere.

In figure 1, each unit of fossil fuel that is consumed produces a marginal external damage, which is equal to the vertical distance between the supply curve and the marginal social cost curve. The free market would lead to production of all units for which consumers' marginal private benefit (that is, the height of the demand curve, or marginal willingness to pay) exceeds the producers' marginal private cost (that is, the height of the supply curve), leading to production and consumption of Q_{market}, at a price of P_{market}. The economically efficient quantity is $Q_{efficient}$, where the demand curve intersects the marginal social cost curve.

At the free market outcome Q_{market}, the deadweight loss is the gray shaded triangle. That deadweight loss represents the amount by which the marginal social cost exceeds the marginal private benefit, summed up over all the units between $Q_{efficient}$ and Q_{market}.

The economically efficient solution could be achieved if we imposed a Pigouvian tax that equaled the marginal external damage at the efficient quantity. For example, if the marginal external damage at $Q_{efficient}$ were $40, and one unit of fossil fuels emits one ton of carbon dioxide

[4] See DeLong and Magin (2009) for an informative discussion of these issues.

Figure 1 – The Market for Fossil Fuels

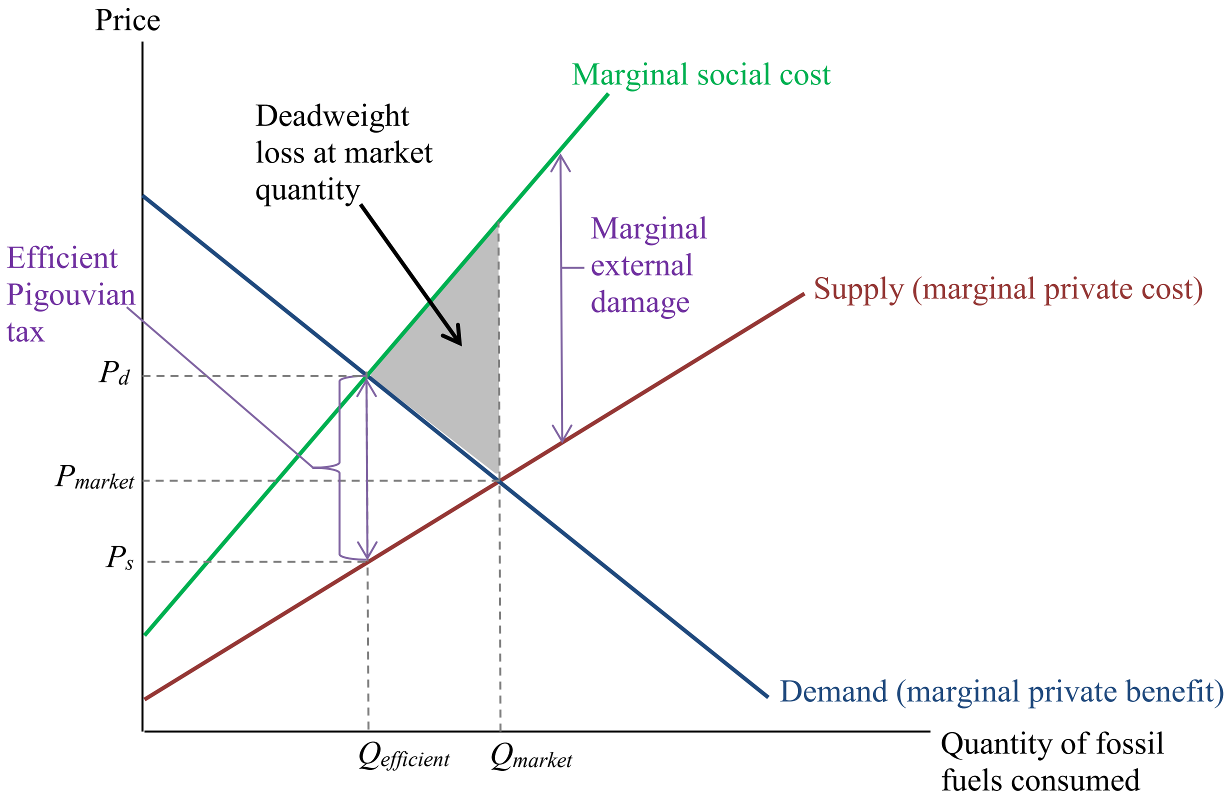

into the atmosphere, then the economically efficient tax would be $40 per unit. That marginal external damage at the efficient quantity is known as the *social cost of carbon*.[5] The tax inserts a "wedge" in between the supply of demand curves with a height equal to the amount of the tax per unit (in the example, $40), pushing the price paid by demanders up to P_d, and pushing down the price received by suppliers after the tax to P_s. The higher price paid by consumers gives them an incentive to cut back on consumption, and the lower price received by suppliers gives them an incentive to cut back on production, until we reach a new equilibrium of $Q_{efficient.}$ Doing that eliminates the deadweight loss that otherwise would have occurred if we had produced and consumed at Q_{market}, which in turn increases economic surplus.

What role does discounting play in this analysis? Conceptually, marginal external damage caused by consuming one unit of fossil fuel in figure 1 is the discounted present value of all future dollar-valued damages of the carbon dioxide emitted when we burn that one unit of

[5] The term "social cost of carbon" is sometimes used to refer to the *average* external damage from emitting a ton of carbon dioxide. However, if our goal is to determine the economically efficient Pigouvian tax, what we really need to know is the *marginal* external damage at the efficient quantity. The average and marginal external damage would only be equal if the marginal external damage happened to be constant, i.e., if the marginal social cost and marginal private cost curves in figure 1 were parallel. There's no obvious reason to expect this to be true.

fossil fuel. If our goal is to compute the *economically efficient* Pigouvian tax, we need to discount those future damages using the market discount rate. The dollar value of the damages in each future year should be the marginal willingness to pay of people in the future to avoid the damage – that is, what is the discounted sum of the maximum amount that people around the globe, over all future generations, would be willing to pay to avoid those damages, if they could avoid them by paying?

To get the efficiency analysis right, when estimating those future marginal external damages, we'd need to account for the likelihood that future generations would tend to have a higher willingness- to-pay to avoid a given amount of damage if economic growth causes them to be richer than we are today. Remember that willingness-to-pay is the joint product of preferences and ability-to-pay (wealth). Moreover, in the future, the combination of economic growth and climate change would cause environmental amenities such as a cool climate to become relatively scarce compared to man-made goods, which would be relatively abundant due to economic growth. In general, when one thing becomes scarcer and another thing becomes more abundant, that causes the marginal willingness to pay for the relatively scarcer thing to rise (it's a movement up and to the left along the demand curve for that scarce thing). That would also raise future generations' willingness-to-pay to avoid the damage from carbon emissions, and other things equal that should increase our estimate of the marginal external damage. Not all efforts to model the economic impacts of climate change take such considerations into account, because it is difficult to do so. But economists agree that in principle such considerations *should* be account for. Regardless, to convert those future damages to present values, we would need to discount at the market discount rate, otherwise we would not be answering the question of what is economically efficient.

In their book *Climate Shock*, Gernot Wagner and Martin Weitzman accept that this is the right way to think about the question, but they argue that the appropriate market discount rate to use in order to compute the marginal external damage (that is, the social cost of carbon), should be very low or perhaps even negative, because fighting climate change is an investment that provides us with insurance against potential catastrophe.

The Social Discount Rate and Intergenerational Equity

The *social discount rate* addresses a completely different question than the market discount rate does. The social discount rate is relevant to the question of what the *ethical* value of a dollar is at different points in time.

To illustrate the idea, imagine that our goal is to maximize utilitarian social welfare, summed up across all people in all generations now and in the future, counting each person's utility equally. One thing we would want to account for in an analysis of what policies would maximize utilitarian social welfare would be the fact that economic growth causes peoples' dollar-valued well-being to increase over time, on average, and that when dollar-valued well-being increases, the marginal utility of an additional dollar declines.

Between 1870 and 2010, average real GDP per person for the world as a whole increased at an average annual rate of 1.6 percent per year.[6] Suppose we assume, perhaps conservatively, that in the future, average real GDP per person will increase by 1 percent per year on average. Further suppose, perhaps conservatively, that peoples' utility functions are such that each 1 percent increase in income is associated with a 1 percent decline in marginal utility, so that a doubling of income would cut marginal utility in half. This is a slightly smaller degree of diminishing marginal utility of income that is suggested by happiness surveys analyzed by Layard, Mayraz, and Nickell (2008). It is also *much* less than the degree of diminishing marginal utility implied by the risk premium in corporate stocks.[7] If we believe that future generations, on average, will experience real income growth of one percent per year on average, and that each one percent increase in income is associated with a one percent decline in the marginal utility of income, it means that on average marginal utility is declining by one percent every year. That would imply that an additional real dollar one hundred years from now is only worth $1/(1.01)^{100} = 0.37$ times as much in terms of utilitarian social welfare as an additional dollar is worth today. The reason is that over 100 years, the average person is predicted have $1.01^{100} = 2.7$ times as much real income as the average person today, and given the assumptions we made above, that person's marginal utility would only be 0.37 times as high as the marginal utility of the average person today.

In the example described above, the 1 percent rate at which marginal utility declines every year, on average, is the *social discount rate*. It represents the rate at which the ethical value of a dollar declines over time, in this case because people are growing richer over time and their marginal utilities are falling as a result, and because we've assumed a utilitarian social welfare function.

So, in this hypothetical scenario, we assume that people will be much richer in the future because of continued economic growth, and therefore future generations will have much lower marginal utilities for an additional dollar, then why would a utilitarian ever favor doing anything that involves making a sacrifice today in order benefit people in the future? One answer is that transferring resources from people today to people in the future involves the *opposite* of Okun's "leaky bucket" problem.

---

[6] Author's calculation based on data from the Angus Maddison Project, available at <http://www.ggdc.net/maddison/maddison-project/home.htm>.

[7] Estimates of the degree of risk aversion necessary to explain the historical risk premium in rates of return on the U.S. stock market suggest that each 1 percent increase in income would be associated with a 30 to 50 percent decline in the marginal utility of income. However, other evidence on demand for insurance and how people respond to other risky situations in their lives seems to suggest something on the order of a 1 percent increase in income being associated with a 1 to 3 percent decline in marginal utility. There are a variety of theories that might reconcile these conflicting pieces of evidence that lean towards the latter conclusion. See Siegel and Thaler (1997) and DeLong and Magin (2009) for accessible discussion of these issues.

Okun compared redistribution from rich to poor through taxation and transfers to carrying water to a thirsty person in a "leaky bucket" – the rich are made worse off by more than the poor are made better off when measured in dollars, because of the administrative costs of operating the tax-and-transfer system, and the deadweight loss that comes from peoples' behavioral responses (e.g., working less) in response to the distortion to incentives caused by taxes and transfers. Those are the "leaks in the bucket.' In that case, the gain in social welfare from additional bits of redistribution from rich to poor is increasingly limited and eventually eliminated by the leak in the bucket. So for example, if each $1 collected in taxes from the rich and transferred to the poor causes $2 in harm once deadweight loss and administrative costs are taken into account, then half the bucket leaks out, and the transfer only raises utilitarian social welfare if an additional dollar is worth at least twice as much utility to the poor as to the rich.

In the case of redistributing resources from the present to the future, we have the opposite of a "leaky bucket" problem, because anything we save and invest for the benefit of people in the future can earn interest at a market rate of return. So for example, abstracting from risk for the moment, imagine that it were possible to save and invest for the benefit of people in the future in a way that offers real rate of return of 2 percent per year. If the social discount rate (the rate at which marginal utility is declining) were 1 percent per year, as in our hypothetical scenario, then a utilitarian who thought this through would argue that we are *ethically obligated* to save more for the benefit of people in the future. That's because the value of our investment is growing over time at a value of 2 percent per year, while the value of a dollar in terms of utilitarian social welfare is only declining at 1 percent per year.

In this scenario, if we make a sacrifice today of $1 billion in order to save and invest for the benefit of people 100 years from now at a 2 percent rate of return, that $1 billion will make those people 100 years from now better off by ($1 billion)×$(1.02)^{100}$ = $7.24 billion. Or to make an analogy to Okun's metaphor, in this example, when we try to transfer resources to people 100 years in the future, instead of some of the initial bucket leaking out, the initial bucket actually multiplies into 7.24 buckets! If we normalize marginal utility today to be equal to one, and the social discount rate is 1 percent, it implies the marginal utility of people 100 years from now is on average $1/(1.01)^{100}$ =0.37. So we are making people 100 years from now better off by (7.42 billion)×0.37 = 2.75 billion utils, and we are only sacrificing 1 billion utils today in order to make that happen. So a utilitarian would argue we should do that.

More generally, a utilitarian would argue that we should save and invest more for people in the future whenever the market rate of return exceeds the social discount rate. Some people view this as an odd implication of utilitarianism – it suggests that if market rates of return on investment are higher than the rate at which marginal utilities are declining over time as people get richer due to economic growth, we are actually ethically obligated to redistribute to people in the future who will be richer than us!

Other things equal, the higher the market rate of return that is available, and the lower the social discount rate, the more we would be ethically obligated to do for the future overall, in the utilitarian framework. Factors that would lead a utilitarian to conclude that we should do *more* saving and investment for the benefit of future generations, holding other things constant, would include:

- A lower predicted rate of economic growth (which means that marginal utilities don't decline so much over time).

- A higher market rate of return on investment, because it would mean that a dollar of consumption sacrificed today would produce larger dollar-valued gains for people in the future.

- A *less* sharply curved utility function (if we predict economic growth in the future will be positive). In that scenario, other things equal, economic growth causes marginal utility to decline at a slower rate in the future. So for example, if utility as a function of income or consumption were an upward sloping straight line, the constant slope implies constant marginal utilities, so that as economic growth makes future generations richer, the gain in utility that they get from an additional dollar does not diminish. That is the most extreme version of what we mean by "less sharply curved," and in that case, a utilitarian would put greater ethical weight on dollar-valued gains to richer future generations than otherwise. Alternatively, if the utility function starts out steep and then gets very flat as someone gets richer, we would call that a very "sharply curved" utility function, and in that case dollar-valued-gains to future generations would get less ethical weight in the utilitarian framework.

- Exposure of future generations to greater risk. This is because if there is diminishing marginal utility, the upside of the risk will raise their utilities by less than the downside of the risk will lower them. This particular factor matters more if utility functions are *more* sharply curved.

There's a limit to how much we'd be ethically obligated to do for the future, however. As we save and invest more and more for the benefit of future generations, that will push market rate of return down (due to diminishing marginal returns to investment), and it will also make people in the future even richer, which lowers their marginal utilities and raises the social discount rate. Eventually, the market rate of the return and the social discount rate would be equalized, and that point the utilitarian would say we'd done enough for the future.

This may all seem impossibly altruistic, and maybe it is. With that said, individuals who value the well-being of their own descendants might behave in a way that bears some resemblance to this. The fact that on average, the rate of growth of real GDP per person has been 1.6 percent per year between 1870 and 2010 means that the average income of people on earth in 2010 was about nine times higher in 2010 as it was in 1870, after adjusting for inflation. We are, on average, *much* richer than our ancestors, largely because past generations left us a very valuable capital stock, excellent technology, and a stock of knowledge that makes us, on average, much more productive than they were.

Given how much richer we are than past generations, it would be hard to argue that past generations were not generous enough to us in ethical terms, on average. They were much, much poorer than us on average, yet made sacrifices to leave us considerably better off than they were. They could have been completely selfish, consumed everything, and left us with nothing. Of course, the bounty of the generosity of past generations is extremely unevenly distributed across people around the globe today, so these are just statements about what happened *on average*. Moreover, as they say about the stock market, past performance is no guarantee of future results.

Is the Social Discount Rate Relevant to Deciding What to Do About Climate Change?

Some economists and legal scholars, such as David Weisbach, a law professor at the University of Chicago, argue that we should discount future external harms from climate change at the market discount rate, and not the social discount rate, in order to determine what Pigouvian tax we should put on carbon. [8] Their claim is that even if we subscribe to a utilitarian ethic, or to any other ethic that posits ethical obligations to future generations (or to the poor, or to anyone else that might be affected by our decisions), those ethical considerations should guide *how much* we do overall for the people who we are morally obligated to help or compensate, but should not determine *which* projects we do for them. The decision about how much to invest in climate change mitigation efforts as opposed to other possible investments should be governed by considerations of economic efficiency, political effectiveness, and so forth, so as to achieve whatever our ethical obligation is at the lowest possible cost (or conversely, to provide as much help as possible at any given cost).

The point that Weisbach and other like-minded scholars are making is *not* that efficiency is the only thing that matters, but rather that we ought to use our overall level of saving and investment to deal with social welfare maximization across generations, and should choose which particular projects to invest in, including efforts to fight climate change, based on which ones have the highest rates of return, after adjusting for risk. Because it corrects an externality, the Pigouvian tax on carbon will be the highest rate-of-return investment up to a point, and that point will be where the tax is just equal to the marginal external damage discounted at the market discount rate.

By contrast, setting the tax on carbon by discounting damages at the *social* discount rate suffers from what Weisbach calls "climate change blinders" problem. A carbon tax equal to marginal external damage discounted at the social discount rate only appears to us to be social-welfare-maximizing because we are blinding ourselves to other alternative means of helping the people

[8] See Posner and Weisbach (2010) and Gardiner and Weisbach (2016) for an articulation of this viewpoint and a debate over whether or this approach is ethically sound. A prominent report on climate change that did its analysis of policy using social discount rates is Stern (2007).

that we are morally obligated to help. This approach would only be social welfare maximizing if we assume that that actions to fight climate change are the *only* way to help our intended beneficiaries. If we make this assumption, then we are failing to consider alternatives that would provide the same amount of help at a lower cost, or that would provide more help at a given cost. Failing to implement those alternatives, if they exist, unnecessarily wastes resources, and the waste of resources does nothing to further our ethical goals.

To someone whose goal is to maximize utilitarian social welfare, for example, the first-best solution would be to set the Pigouvian tax at the economically efficient level, which is the marginal external damage discounted at the market discount rate that reflects the risk characteristics of the investment in fighting climate change. If we think that does not do enough to discharge our ethical obligations to future generations, then we should consider helping those future generations through some alternative method that has a higher rate of return, such as increasing our overall level of saving for the future, and investing that saving in the best available investments with the highest risk-adjusted rates of return. If we've already set the Pigouvian tax at its efficient level, then by definition, a further increase in the Pigouvian tax would not be the investment with the highest available risk-adjusted rate of return. If, by contrast, we were to set the Pigouvian tax at a rate higher than the efficient level, that would be like making an investment for the future that pays less than the market discount rate. Compared to setting the tax at the efficient level, that would either mean providing a given amount of benefit to the future at a higher cost to ourselves today than necessary, or it would mean providing less benefit to the future at the same cost to ourselves today.

Think of the cost of an investment to fight climate change as the consumer surplus and producer surplus in figure 1 that are sacrificed when we reduce consumption of fossil fuels. For all reductions in fossil fuel consumption between Q_{market} and $Q_{efficient}$ in figure 1, the cost in terms of lost consumer and producer surplus is less than the gain in terms of reduced damage to victims of the externality in the future. But if we push fossil fuel consumption below $Q_{efficient}$, then those additional reductions in fossil fuel consumption cost more in terms of consumer and producer surplus lost than they gain in terms of damage avoided.

To see why you'd want to use the market discount rate to set the carbon tax, imagine there are just two points in time, today, and 100 years from now. Suppose we've already set the carbon tax at the economically efficient level, and we are considering whether we should provide additional help to the future through further increases in the carbon tax, or if we should provide additional help to the future through an alternative investment that pays the market discount rate. Further suppose the market discount rate is 2 percent, and the social discount rate is 1 percent. Finally, imagine that in this scenario, by investing an extra $1 billion in fighting change today, we could avoid an extra $5 billion of damage from climate change 100 years from now.

Discounted at the market discount rate of 2 percent, the present value of $5 billion 100 years from now is ($5 billion)/($1.02^{100}$) = $0.69 billion. At the social discount rate of 1 percent, the discounted present value of $5 billion 100 years from now would be $1.85 billion. Since the present value of the cost of the additional investment in fighting climate change is $1 billion, if we were to make our decision based on the *market* discount rate we would *not* do it, but if we were to make our decision based on the social discount rate, we *would* do it.

The logic for why we should make the decision based on the *market* discount rate is an opportunity cost argument. In the example, we have a choice between making a further investment in fighting climate change that provides $5 billion of benefit to people 100 years from now and costs $1 billion today, or making an alternative investment to benefit the future that pays a 2 percent rate of return (that's what a 2 percent market discount rate means). If we put the $1 billion today in the alternative investment, we could make people 100 years from now better off by:

$(\$1 \text{ billion})\times(1.02)^{100} = \$7.24 \text{ billion}$

That makes people in the future better off by more than the $5 billion gain they'd get if we invested more in fighting climate change instead.

Alternatively, the calculations we already did suggest that if we made the alternative investment instead of the $1 billion additional investment in fighting climate change, we could have invested just $0.69 billion today to provide a $5 billion benefit 100 years from now, and we could have pocketed the remaining $0.31 billion ourselves, leaving the future no worse off and making ourselves better off. If we make our decision about which investment to choose based on the market discount rate instead of the social discount rate, it's a potential Pareto improvement. In this scenario, deciding climate change policy based on the social discount rate rather than the market discount rate would only make sense if climate change policy were the *only* reliable way to help future generations.

This argument faces all the same potential objections as any other potential Pareto improvement. Even if we set the Pigouvian tax at the economically efficient level, we are still imposing some costs on people in the future through exacerbating climate change. The fact that the tax is at the efficient level just means the dollar-valued harms to people in the future from the climate change that still happens are worth less than the dollar-valued benefits from burning fossil fuels today, when compared according to the standards of economic efficiency (which includes discounting at the market rate). It could be that even when we impose the efficient Pigouvian tax, the remaining harms to people in the future from climate change are still so serious as to ethically justify additional efforts, beyond the efficient Pigouvian tax, to compensate people in the future. The argument above says that in that scenario, the first-best response would be to provide that compensation in the form of some alternative investment that offers a higher risk-adjusted return than further efforts to fight climate change. But will that compensation happen? That's a good question. This comes down to a question of politics, and the answer is not obvious.

One might object that the relevant choice is not between: (1) a higher-than-efficient Pigouvian tax; and (2) an efficient Pigouvian tax on carbon plus the best available additional investments, the proceeds of which are set aside to compensate future generations for the harm from climate change. Rather, one might argue that choice is between a higher-than-efficient carbon tax and nothing at all, because there is no feasible way to ensure that the alternative investment that is intended to compensate future generations will not be squandered by intervening generations. But the same problem applies to efforts to fight climate change – if we set the carbon tax at higher than the efficient level in an effort to provide ethical compensation to future generations, there's no guarantee that won't be undone by intervening generations either. If one approach to benefitting future generations is significantly more likely to work as a political strategy than the other, then Weisbach would agree that is certainly a relevant consideration that ought to affect which approach is chosen. The point is just that it is not obvious which approach actually is more likely to work politically – and it is a question of the relative effectiveness of different political strategies rather than a question of ethics.

Similar issues arise regarding the distributional consequences of climate change in different parts of the globe. Climate change is likely to have its most severe negative impacts in the countries that already have hot climates, which tend to be disproportionately poor.[9] A utilitarian would argue that the first-best policy would be to help people in those countries not by setting taxes on carbon higher than the economically efficient levels, but rather to set the tax at the efficient level across the globe, which would be the same tax in each and every country – this is because climate change is a global problem, so the marginal external damage of an additional ton of carbon emissions is identical regardless of where in the globe it happens. Then we would address the distributional concerns by transferring resources in some other manner, such as foreign aid cash transfers. Is that likely to happen? That's also a good question. Another good question is whether it would be any more likely to happen if we tried to do it through a higher-than-efficient carbon tax.[10]

Despite the complications, many economists, including Wagner and Weitzman, are sufficiently persuaded by the argument for discounting at the market discount rate that they just take it as given. But Wagner and Weitzman argue that if you take the insurance-like aspect of fighting climate change seriously, even the cold hard calculus of economic efficiency seems to support a very high Pigouvian tax on carbon emissions.

[9] Burke, Hsiang, and Miguel (2015) provide estimates of the economic impacts of climate change in different parts of the globe.

[10] Posner and Weisbach (2010) and Gardiner and Weisbach (2016) provide further discussion of these issues.

References

Burke, Marshall, Solomon M. Hsiang, and Edward Miguel. 2015. "Global Non-Linear Effect of Temperature on Economic Production." *Nature*. Vol. 527 (November), pp. 235-239.

DeLong, J. Bradford and Konstantin Magin. 2009. "The U.S. Equity Premium: Past, Present, and Future." *Journal of Economic Perspectives*. Vol. 23, No. 1 (March), pp. 193-208.

Gardiner, Stephen M. and David A. Weisbach. 2016. *Debating Climate Ethics*. Oxford University Press.

Layard, Richard, S. Nickell, and G. Mayraz. 2008. "The Marginal Utility of Income." *Journal of Public Economics*. Vol. 92, no. 8-9 (August): 1846-1857.

Okun, Arthur. 1975. *Equality and Efficiency: The Big Tradeoff*. Washington, DC: Brookings Institution Press.

Posner, Eric, and David Weisbach. 2010. *Climate Change Justice*. Princeton University Press.

Siegel, Jeremy J. and Richard H. Thaler. 1997. "Anomalies: The Equity Premium Puzzle." *Journal of Economic Perspectives*. Vol. 11, No. 1 (March), pp. 191-200.

Stern, Nicholas. 2007. *The Economics of Climate Change: The Stern Review*. <http://mudancasclimaticas.cptec.inpe.br/~rmclima/pdfs/destaques/sternreview_report_complete.pdf>

Wagner, Gernot, and Martin Weitzman. 2015. *Climate Shock: The Economic Consequences of a Hotter Planet.* Princeton: Princeton University Press.