

Valuing the Future: Intergenerational Discounting, Its Problems, and a Modest Proposal

by Stephen Marks

Stephen Marks is a Professor of Law
and former Associate Dean,
Boston University School of Law.

Editors' Summary

Competing theories exist for how intergenerational investment projects, such as investments related to global warming, natural resources, energy, etc., should be undertaken. In particular, there are two popular prescriptions: (1) In making intergenerational investments, policymakers should use a zero discount rate; and (2) In making intergenerational investments, policymakers should use the market rate. Neither of these prescriptions is correct. Indeed, using present-value discounting at all is extremely problematic. Instead, the best we can probably do is to adopt a simple algorithm: set certain minimal goals for future generations: clean air; potable water; sufficient energy supplies; a nontoxic environment; etc.; and then analyze the most cost-effective way of achieving those goals.

In 2007, the *University of Chicago Law Review* published papers from its symposium on Intergenerational Equity and Discounting.¹ At issue was how we should evaluate investments in the distant future. Typical investments of this sort are those in global warming and nuclear energy. The former could affect climatic changes hundreds, if not thousands, of years in the future. Likewise, nuclear waste produces waste materials with half-lives measured in hundreds of thousands and even millions of years. How shall we evaluate current decisions that have significant effects on distant future generations? The 2007 symposium issue provided a variety of competing answers. In those answers were two competing prescriptions:

1. In making intergenerational investments, policymakers should use a zero discount rate.
2. In making intergenerational investments, policymakers should use the market rate.

In this Article, I show that, from the standpoint of a policymaker who seeks to maximize intergenerational social welfare, neither prescription is correct. In making this claim, I lay out a general framework for thinking about intergenerational discounting. At the outset, please note the following: (1) There is nothing radical in this approach, as it builds on standard financial and welfare theories. (2) The approach does not yield a silver bullet, a grand algorithm for resolving intergenerational issues. (Thus, I cannot title this Article "Toward a General Theory of Intergenerational Investing," or anything of the sort.) (3) Rather, the approach counsels humility. This is an extremely difficult nut to crack. (4) In the end, we are most likely left with nothing more than an empathic gut (perhaps Kantian) response: let's leave future generations with a relatively unpolluted environment and with enough resources to make a go of it.

I. The Issue

It is a fundamental principle of finance that one dollar today is worth more than one dollar tomorrow, simply because one dollar today will grow into more than one dollar tomorrow. More specifically, one dollar today will grow into $(1 + r)^n$ dollars after n periods, where r is the period interest rate. Equivalently, to get one dollar n periods from today, you would need to invest $1/(1 + r)^n$ today. We say that the present value of one dollar received n periods from today is $1/(1 + r)^n$.

In a typical investment, resources are invested in the near-term in exchange for returns later on. Modern finance

1. Mathew D. Adler et al., *Symposium: Intergenerational Equity and Discounting* 74 U. CHI. L. Rev. 1 (2007).

theory counsels us to find the net present value of cash flows and to make investments when the net present value is greater than zero. This is done by finding the present value of all cash flows and adding them together. The problem, as some see it, is that, when applied to intergenerational investment, the present value of one dollar 500 years from now (assuming an annual real interest rate of, say, 3%), is less than 0.0004 cents. Thus, using the present value criterion, it would not be worth investing \$1 million today to avoid a \$2 trillion environmental disaster 500 years from now.

The Chicago symposium highlighted the state of disagreement on the question of the appropriateness of intergenerational discounting using market rates. This symposium brought together contributions by some of the best minds in law and economics. More than anything, however, the diverse offerings and conclusions highlighted the state of disagreement as to appropriateness of intergenerational discounting. Some of the contributors argued that we should apply standard time-value-of-money discounting to intergenerational projects using market rates. Others suggested using a zero percent discount rate for intergenerational projects. In such a scheme, one dollar now is worth the same as one dollar 500 years from now. If sacrificing \$1,000 of consumption today results in a gain of \$1,001 in the next generation, then this investment should be undertaken.

The ultimate issue is how we as voters and citizens should direct our government to act in making intergenerational investment decisions. Should we instruct our government to employ present-value discounting using market rates in choosing whether to invest in nuclear power generation or global warming abatement and the like? Or should we instruct them to use a zero rate? In this Article, I side with the critics who argue against the use of market rates for intergenerational investments. Nevertheless, I also show that a zero discount rate is problematic.

In terms of organization, I will lay out the arguments as follows: Part II discusses what optimal intergenerational distribution looks like theoretically. Part III demonstrates that discounting using a zero rate does not usually lead to good distributional outcomes. Part IV demonstrates that discounting using the market rate does not usually lead to good distributional outcomes. Part V discusses, in light of these two conclusions, how we might approach intergenerational investments.

II. What Is a Socially Optimal Intergenerational Distribution?

A. The Social Welfare Function

To some, the use of present value discounting and market rates disrespects the equal claims of future generations. Some analysts have expressed the respect for future generations using the following stylized (and highly simplified) intergenerational social welfare function:

$$W = u_0 + \beta u_1 \quad (1)$$

where W is social welfare, u_0 is the utility of the present generation and u_1 is utility of future generations. In this formulation, β represents how much we should weight the utility of future generations in the social welfare function. If $\beta = 1$, then we weight the utility of future generations equal to that of the present. If $\beta < 1$, we weight the future less than the present. Note that we can write β as

$$\beta = \frac{1}{1+d}$$

where d can be called the *utility discount rate* over the relevant periods. Thus, a utility discount rate of zero (meaning that we do not discount the future at all) implies that $\beta = 1$. Conversely, if $d > 0$ then we are discounting future utility relative to present utility. A discount rate of infinity implies that we place no weight on the utility of future generations, that is, $\beta = 0$.

The social welfare function with $\beta = 1$ becomes

$$W = u_0 + u_1 \quad (2)$$

I will refer to this as the *strict utilitarian welfare function*.² The strict utilitarian welfare function does not favor either generation. It gives the two generations equal respect. In general, equal respect means that if we were to reverse the utilities, the social welfare would not change. That is,

$$W(u_0, u_1) = W(u_1, u_0) \quad (\text{Equal Respect Condition})$$

Due to the commutative property of addition, the strict utilitarian welfare function satisfies the Equal Respect Condition. The strict utilitarian function is not the only function that satisfies the Equal Respect Condition. We encounter others later in this Article.

In spite of satisfying the Equal Respect Condition, the strict utilitarian social welfare function is not very attractive. Consider, for example, the following distributions:


2. This function dates back to at least 1863. See JOHN STUART MILL, ON UTILITARIANISM (1863).

Distribution	u_0	u_1
A	0	1,000
B	1,000	0
C	500	500
D	499	499

Under the strict utilitarian social welfare function, distributions A, B, and C are all equal in terms of social welfare, and each of these is superior to D. Yet most of us, if we were designing our ideal society (and especially not knowing which generation we would find ourselves in) would prefer D to either A or B.

Graphically, the utilitarian social welfare function is represented by the following isowelfare graph. (Read this like you would a topographic map. Points on the same line represent the same social welfare. Social welfare increases as we move up and to the right on the graph.)

Figure 1: Strict Utilitarian Welfare Function




An alternative intergenerational social welfare function is the so-called maximin function. Social welfare is calculated as the minimum of the two generations' utilities.³

$$W = \min(u_1, u_2)$$

The function gets its name because the goal, as with all social utility functions, is maximization, so that in this case, the goal is to maximize the minimum of the two generations' utilities. The maximin social welfare function satisfies the Equal Respect Condition. Graphically,

we can represent this function with a series of isowelfare lines as follows:

Figure 2: Maximin Welfare Function




For many, the maximin function is also problematic. Consider the following distributions:

Distribution	u_0	u_1
A	500	500
B	500	600
C	600	500

Under the maximin function, A and B and C are equivalent. This seems odd. In fact, most of us (but perhaps not all of us), standing behind the veil of ignorance, could very well prefer B or C to A.

In a sense, the utilitarian and the maximin functions are boundaries for a wide variety of convex social welfare functions. A typical member of this set is represented by the following graph:

3. See, e.g., John E. Roemer, *The Ethics of Intertemporal Distribution in a Warming Planet*, 77 LA ÉTICA DE LA DISTRIBUCIÓN INTERTEMPORAL EN UN PLANETA EN CALENTAMIENTO CLIMÁTICO 94-95 (2009), available at http://www.centrodeestudiosandaluces.info/cursos/adjuntos/JohnE_Roemer.pdf.

Figure 3: Concave Welfare Function

In order for a convex function of this sort to satisfy the Equal Respect Condition, it is necessary graphically for the isowelfare lines to be symmetrical around a 45-degree line emanating from the origin.

B. Consumption

When a society makes an intergenerational investment, it is giving up consumption by the present generation for increased consumption by a future generation. Therefore, we must understand social welfare in terms of consumption, not just in terms of utility. We begin by noting that utility depends on consumption. Symbolically, we represent this as follows:

$$u_0(c_0) \text{ and } u_1(c_1)$$

Thus, social welfare can be written as a function of consumption directly:

$$V(c_1, c_2) = W(u_1(c_1), u_2(c_2))$$

The social welfare function V will look a bit different than the social welfare function W , due to the diminishing marginal utility of consumption. That is, each additional unit of consumption confers less utility than the previous unit. Consider the strict utilitarian social welfare function:


$$W(u_1, u_2) = u_0 + u_1$$

As a function of consumption, this becomes:

$$V(c_0, c_1) = u_0(c_0) + u_1(c_1) \quad (3)$$

Recall that graphing the strict utilitarian social welfare as a function of utility yielded linear parallel isowelfare lines.

Graphing the strict utilitarian welfare function as a function of consumption yields iso-utility lines that are convex to the origin:


Figure 4: Concave Welfare Function (Welfare as a Function of Consumption)

In what follows, we will use both types of social welfare functions (functions of utility and functions of consumption) to glean various insights about intergenerational investing. First, however, we must consider the role of consumption possibilities.

C. Consumption Possibilities and the Rate of Return on Investment

We cannot ascertain the optimal distribution of consumption across generations without knowing the consumption possibilities, so we turn to this now. In general, an investment project requires giving up present consumption in exchange for future consumption. We can represent all of the possible combinations of consumption possibilities with a graph such as the following.

**Figure 5: Consumption Possibility Set
(Welfare as a Function of Consumption)**



Any combination of c_0 and c_1 that is within the boundaries of the consumption set is possible. The slope of the boundary of the consumption set is related to the rate of return on the marginal investment project. To see this, start by supposing that all consumption is in the period 0 (the present generation). This is represented by point A. Now, suppose we want to move to point B. This requires giving up consumption equal to amount a in the first period and getting consumption of amount b in the second period. In other words, a is invested in the first period, and b is received in the second period. The rate of return is just:

$$\frac{b-a}{a} \quad \text{or, equivalently,} \quad \frac{b}{a} - 1$$

This expression is the *rate of return on investment*. We can denominate the returns in any currency, for example, in dollars. If one gives up \$100 in consumption in period 0 ($a = 100$) for \$110 additional consumption in period 1 ($b = 110$), then the rate of return on investment is 0.1, or 10%. We can find the rate of return on infinitesimally small investments at a single point on the boundary by taking the negative of the slope of the tangent at that point and subtracting one. As more investments are made (as we move right to left on the boundary), the rate of return on investments falls. This makes sense. Investments come with different returns, and we engage in the ones with the higher returns first.


If you were a company looking at investment opportunities, you would invest in projects with positive net-present values. Suppose that your costs of capital were 10%, that is, you could borrow funds at a 10% interest rate. To make profitable investments, you would evaluate projects by their net-present value using a 10% discount rate. (The

discount rate is the rate used in present value calculations.) Projects with rates of return over 10% will have a positive net-present value.

If you could obtain capital at 5%, then you would use this as your discount rate and invest in projects with returns greater than 5%. You would invest in more projects than if the rate were 10%. The same is true for intergenerational investing. If managers use a 5% rate for evaluating intergenerational investments, they will invest in more projects than if they used a 10% rate. If they use a zero rate, they will invest in yet more projects.

We can transform the consumption possibilities graph into a utility possibilities graph. If we do so, we must note that diminishing marginal utility comes into play once again. The utility possibility set, corresponding to the above consumption possibility set, is shown in Figure 6.

**Figure 6: Utility Possibility Set
(Welfare as a Function of Utility)**




The slope of the boundary of the utility possibility set is the *utility rate of return on investment*. That is, if one gives up 1,000 utils today for 1,200 utils in the future, then the utility rate of return is 0.2 or 20%. As with the ordinary rate of return on investment, one can calculate the utility rate of return on investment for any point on the boundary by taking the negative of the slope of the tangent and subtracting one.

D. The Social Optimum


Representing the social optimum graphically involves combining the social welfare function with the consumption possibility set. Let us start with the easiest case. Suppose that we have a strict utilitarian social welfare function. If

we lay this social welfare function on top of the consumption possibility set, we get the following graphs, one in terms of utilities, and the other in terms of consumption.

**Figure 7a: Welfare Optimization
(As a Function of Consumption)**



**Figure 7b: Welfare Optimization
(As a Function of Utility)**




III. The Zero Rate


Examination of Figure 7a should convince the reader that the optimal rate of investment is unlikely to be zero, even with the strict utilitarian welfare function, a function that satisfies the Equal Respect Condition. If we wanted inter-generational investors to invest up to the social optimum, then we should set the discount rate equal to that indicated at the social optimum in Figure 7a. This optimal discount rate is just the negative of the slope of the tangent at the optimization point, that is, where the consumption possibilities set touches the highest isowelfare curve. The optimal discount rate depends both on the curvature of the isowelfare lines (determined by the degree that marginal utility of consumption decreases with consumption) and the curvature of the consumption possibilities set (determined by the available investment opportunities.) Theoretically, optimal discount rates can be either positive or negative. Figure 8a provides an example where the optimal discount rate would be negative.

As an aside, note that with the strict utilitarian welfare function, the *utility* discount rate is indeed zero at the optimum. But investment decisions are not made in terms of utility, but rather in terms of dollars (or some other currency), so that this interesting fact is irrelevant.

**Figure 8a: Welfare Optimization
(As a Function of Consumption)**



**Figure 8b: Welfare Optimization
(As a Function of Utility)**



If we substitute the strict utilitarian welfare function with some other social welfare function, even with another one that satisfies the Equal Respect Condition, it will still be unlikely that a zero rate will maximize social welfare. The isowelfare curves will look much the same, that is, concentric convex lines, though with different curvature. As long as they have this shape, then the welfare-maximizing discount rate will be an interaction of the curvature of the isowelfare lines and the curvature of the boundary of the consumption possibility set.

And note this: even if the isowelfare lines relative to consumption were straight (which could occur if we used a strict utilitarian welfare function and if there was no diminishing marginal utility of consumption, which, of course, is counterfactual), the result of social optimization could still result in highly inequitable distribution of consumption between generations. This can be seen by looking at 7b and assuming that it represents consumption rather than utility. These graph shows social optimization with very inequitable distributions. Again, this comes from the poor distributional aspects of the strict utilitarian welfare function.

That a zero investment discount rate can lead to a very unequal distribution can be made by a simple example. Suppose there are two generations. The first starts out with \$100. Some of this is consumed and some invested. Suppose that the second generation consumes the result of the investment projects. Suppose that there are 10 possible intergenerational investment projects:

- Project 1: Invest \$10 now, reap \$50 later. (400% return)
- Project 2: Invest \$10 now, reap \$30 later. (200% return)
- Project 3: Invest \$10 now, reap \$20 later. (100% return)

- Project 4: Invest \$10 now, reap \$15 later. (50% return)
- Project 5: Invest \$10 now, reap \$14 later. (40% return)
- Project 6: Invest \$10 now, reap \$13 later. (30% return)
- Project 7: Invest \$10 now, reap \$12 later. (20% return)
- Project 8: Invest \$10 now, reap \$11 later. (10% return)
- Project 9: Invest \$10 now, reap \$9 later. (-10% return)
- Project 10: Invest \$10 now, reap \$8 later. (-20% return)

Keep in mind that these are intergenerational projects, so large rates of return should not surprise us. We now can choose to invest in projects based on a return criterion. For example, we could invest in all projects that give a rate or return greater than or equal to 50%. Note that it makes sense to invest in the projects with the highest returns first. The various possibilities are given in the table below.

Figure 9

Investments	c_1	c_2
None	100	0
1	90	50
1,2	80	80
1,2,3	70	100
1,2,3,4	60	115
1,2,3,4,5	50	129
1,2,3,4,5,6	40	142
1,2,3,4,5,6,7	30	154
1,2,3,4,5,6,7,8	20	165
1,2,3,4,5,6,7,8,9	10	174
1,2,3,4,5,6,7,8,9,10	0	182

Setting the discount rate to zero would result in the first eight projects being adopted and the first generation consuming \$20 and the second \$165. Now, suppose we are committed to equal consumption across generations and that we are charged with the responsibility for setting the interest rate for projects to attain this goal. By setting an intergenerational discount rate of 150%, we cause the first two projects to have positive net-present values and thus to be undertaken. The result is that both generations will have \$80 to consume. (The first two investment projects require an investment of \$20, reducing first-generation consumption to \$80, and have a total payout of \$80.)

I should emphasize once again the unattractiveness of the strict utilitarian welfare function. Both figures 6 and 7 demonstrate that social optimizing using the strict utilitarian welfare function can produce really unequal distributions. (See the distribution of consumption in both figures.)

So, here are our first results in a nutshell:

1. The strict utilitarian intergenerational social welfare function satisfies the equal respect condition.
2. The strict utilitarian intergenerational social welfare function implies a utility discount rate of zero.
3. In spite of 2, the optimal discount rate (in dollars) under the strict utilitarian intergenerational social welfare function is highly unlikely to be zero.

4. The optimal discount rate (in dollars) under any other social welfare function, including those satisfying the Equal Respect Condition, is highly unlikely to be zero.
5. Social optimization using a strict utilitarian intergenerational social welfare function is not likely to produce equitable distributions between generations.
6. Application of a zero discount rate for investment is not likely to produce equitable distributions between generations.
7. Social optimization using other social welfare functions could produce more equitable distributions between generations (but see 4).

Need we say more about zero rates?

IV. Market Rates and Social Optimality

If we eschew the use of a zero rate, should we instead adopt market rates? Proponents of the use of market rates come to it from two directions. One approach is to assert that, in equilibrium, market rates result in a socially fair division between present and future generations. I will show that there is no convincing evidence for this and, in fact, there are theoretical reasons that suggest this cannot be the case. A second approach is to assert that, even if the market rate does not result in a fair distribution between generations, government should nevertheless use this rate in making public investments, because to do otherwise will result in distributions that are not Pareto-optimal. I will show that, even though the use of differential rates (the market rate by the private sector and a different rate by the public sector) may not be Pareto-optimal, such differential rates nevertheless could result in distributions that are socially superior to the Pareto-optimal distributions. Indeed, many Pareto-optimal distributions are highly unattractive in terms of intergenerational fairness. Both these findings suggest that government should probably make intergenerational investment projects using rates that differ from the market rate, or use a different mechanism completely.

A. *One Extreme: The Long-Lived Robinson Crusoe*

To better understand the setting of market rates, let us begin with the simplest of economies, that of a long-lived Robinson Crusoe. Suppose that our Mr. Crusoe lives through two generational time periods. Because the economy is Mr. Crusoe and no one else, his two-period utility function is the social welfare function. Faced with a consumption possibility set, Mr. Crusoe will choose to allocate consumption between the two time periods in a manner that maximizes his utility and, by definition, maximizes the social welfare function.⁴

It should be noted, at least in passing, that Mr. Crusoe's two-period utility function is unlikely to be the strict utilitarian function that we utilized several times above, but rather a function that pays some heed to the distribution of utility, and not just its sum total. That is, its graph is likely to be convex to the origin. But that is neither here nor there. We take Mr. Crusoe's utility function for whatever it is.

So far, we have Mr. Crusoe, but no market. We can create a market in this economy by creating a business. Mr. Crusoe is the sole shareholder, but the business is managed by Ms. M, who makes all investment decisions for the business. Her goal is to maximize profits for the shareholder. Maximizing profits involves investing in all projects with a positive net-present value. Now, let us add a bank to the mix. The bank exists simply to take deposits from Mr. Crusoe and to make loans to Ms. M's business. Here is how it works: the bank announces an interest rate for both deposits and loans. (We are assuming that the bank just exists to clear the markets.) Using this rate of interest, Ms. M then calculates which projects the business will invest in, using the present value criterion. She announces the amount that she will need to borrow and the profits that the business will return in the second period after paying back the loan. (These profits belong to the shareholder, Mr. Crusoe.) Mr. Crusoe, taking into account both the interest rate and profits, decides how much he is willing to invest. He will do this to maximize utility according to his utility function. If the amount of funds that Mr. Crusoe is willing to supply is less than the amount Ms. M wants to borrow, then the bank announces a higher interest rate. At a higher rate, Ms. M will invest in fewer projects. Because she is investing in fewer projects, she will lower her demand for funds. These fewer projects will produce less cash in the second period. Because the bank is offering a higher rate of interest, Mr. Crusoe will increase his deposits in the bank. The bank will continue to increase the rate of interest until supply equals demand. (If the interest rate is initially too high, the opposite adjustment will take place.) At equilibrium, the market will maximize Mr. Crusoe's interperiod utility. That is, the result will be the same as when Mr. Crusoe simply made all of the decisions. Thus, the result of the Robinson Crusoe economy with markets is the generation of a market rate that produces a temporal distribution of consumption that maximizes social (that is, Mr. Crusoe's) welfare.

B. *The Other Extreme: Nonoverlapping Generations*

What happens if we replace Mr. Crusoe with two individuals, one of whom lives and dies in the first period, and the other who lives and dies in the second. It should be clear that, absent any sort of altruism on the part of the first generation, that there will be no investment. That is, any intergenerational investment reduces the utility of the first generation. The discount rate for evaluating intergenera-

4. A variant of this Robinson Crusoe economy is that assumed by optimal-growth models in which the horizon is infinite and the social welfare function is maximized.

tional investments is infinite. No intergenerational investments will be made.

C. *In the Middle*

The Long-Lived Robinson Crusoe and the Nonoverlapping Generations represent two extremes, one in which market rates result in equitable distribution over time (the Robinson Crusoe case), and one in which market rates are infinite and intergenerational distribution is completely skewed toward the current generation. Intuitively, one can guess that the result in a more realistic economy would be something in between, and indeed it is.

Let us begin with the Nonoverlapping Generations model. There are two possible escapes from the result of an infinite market rate and a skewed distribution. One is that, even if generations are nonoverlapping, altruism may operate to provide intergenerational equity. The second possible escape is that perhaps *overlapping* generations will produce a market rate that, even though it is the result of nonaltruistic behavior, produces intergenerational equity. Each of these escapes will be examined in the next two subsections.

I. *Altruism*

Again, taking the nonoverlapping generations model (with one person in each generation) as our starting point, let us ask the following: what if the utility function of the first generation incorporated notions of intergenerational equity? The first generation could have a utility function that incorporated the utility of the second generation. At the extreme, the utility function of the first generation could simply replicate an equitable intergenerational social welfare function. In such a case, the market outcome would replicate the Robinson Crusoe economy, and we would have intergenerational fairness.

The argument falters, however, when we populate each generation (still for the moment nonoverlapping) with numerous consumers. Assume that each consumer gets utility not only from his or her own consumption, but also from the average consumption of the distant generation. It should be clear that this model creates a massive externality. If a single consumer increases investment in the future generation, this increases the utility not only of the future generation (which is nonproblematic), but also the utilities of all of the other altruistic members of the first generation. As with the case of all positive externalities, this will result in insufficient investment. Intuitively, with a large population in both generations, any effect that a single consumer can have on average future consumption is negligible. Such a consumer would be willing to make such an investment in concert with other consumers, but not alone. Because no consumer would individually be willing to make a nonnegligible investment in future generations, the distributional concerns of future generations cannot become incorporated into the market rate, which depends on the aggregate of individual decisionmaking.

2. *Overlapping Generations*

There is the possibility that because generations in fact overlap, the market decisions of individuals who optimize over several periods could result in a market interest rate that results in a fair intergenerational distribution over long periods of time. In a simple overlapping generations model, each generation lives for two time periods and the second time period of each generation coincides with the first time period of the succeeding generation. Imagine for now a model with three time periods. The first generation lives during time periods 0 and 1, and the second generation lives during time periods 1 and 2. Suppose that we start time period 0 with a limited amount of resources available to the first generation. These can be consumed or invested to generate resources in period 2. In this model, absent any bequest motive, there is no reason for the first generation to invest in something that generates returns in period 3 or to share the resources generated in period 2. Thus, the first generation will consume in both periods, and the second generation will starve.

However, this is all too simple. In most overlapping generations models, the story goes something like this. The first generation begins with some resources. Some of these resources are consumed in period 0 by generation 1, and the rest is invested in capital. In time period 1, the first generation retires, but is still living. Output is produced by the combination of generation 1's capital and generation 2's labor, and this output is divided between the two generations. (Generation 2 receives a wage equal to its marginal product of labor.) Generation 1 consumes its share (it will die by the next period). Generation 2 consumes some of its share and invests some of its share in capital that will be used by a third generation. This goes on forever. Thus, even though the first generation begins with the resources, and uses them, there is consumption by all generations throughout their two-period lives.

For our purposes, it is important to note that there is nothing in these overlapping generations models that suggests that the distribution of consumption is distributionally fair across generations. Distribution is a matter of the interplay of technology and individual preferences. Depending on the parameters, there can be multiple equilibria, and even equilibria that are not Pareto-optimal.⁵ Theoretical evidence suggests that market rates will be significantly higher than in a long-lived Robinson Crusoe economy and that there will be less investment. Since the long-lived Robinson Crusoe economy resulted in a fair multiperiod distribution, this theoretical evidence suggests that using market rates in an overlapping generations economy will result in future generations getting short-changed.

5. See Peter Diamond, *National Debt in a Neoclassical Growth Model*, 55 AM. ECON. REV. 1126 (1965); Oded Galor, *A Two-Sector Overlapping-Generations Model: A Global Characterization of the Dynamical System*, 60 ECONOMETRICA 1351 (1992).

D. Market Rates and Pareto-Optimality

So, if market rates do not result in fair intergenerational distributions, should governments use different rates to evaluate intergenerational investment? To some, the answer is still no. Using rates that are different than the market rate, they argue, will result in a distribution that is not Pareto-optimal. Therefore, in spite of the failure to produce a fair intergenerational distribution, governments should still use market rates.

To understand the Paretian argument, we must grasp the full meaning of Pareto-optimality. In a sense, Pareto-optimality is an incomplete welfare function. It divides the set of possible consumption bundles into those that are Pareto-optimal and those that are not. Let us define Pareto-optimality in the context of our simple two-generation model. We will define Pareto-optimal using the notion of Pareto-superiority.

1. An intergenerational distribution A is *Pareto-superior* to intergenerational distribution B if, in distribution A, one generation is better off and the other is not worse off, compared to distribution B.
2. An intergenerational distribution A is *Pareto-optimal* if no other distribution is *Pareto-superior* to it.

As a social welfare function, the Pareto criterion is incomplete. Given two Pareto-optimal distributions, it is agnostic as to which is socially preferred. Likewise, given two non-Pareto-optimal distributions. To some commentators, a Pareto-optimal distribution is to be preferred to a non-Pareto-optimal distribution. This conclusion does not, however, come out of the Pareto criterion itself. The Pareto criteria merely say that a distribution is socially preferred to another if it is Pareto-superior to it. This all can be seen in the following example.

Figure 10		
Distribution	u_0	u_1
A	0	900
B	1,000	0
C	500	500
D	498	498
E	0	600

Suppose that the above five distributions are the only ones possible. Comparing each to each other pairwise, we can say the following:

- A is Pareto-superior to E.
- C is Pareto-superior to D.

Thus, since distribution A does not have any distribution Pareto-superior to it, it is Pareto-optimal. Likewise for B and C. However, the Pareto criterion says nothing about whether A, B, or C is socially preferred to the other two.

Some commentators have suggested that, in terms of social welfare, Pareto-optimal outcomes are socially preferable to non-Pareto-optimal outcomes. This notion, if

examined closely, is not intuitively appealing. *In the above example, most people, in the role of social planner, would prefer distribution D to A, for example, even though A is Pareto-optimal and D is not.*

The Pareto criterion is an incomplete social welfare function because, although it does compare some pairs of distributions, it says nothing about other pairs. The social welfare functions that we have considered above (the strict utilitarian function, the maximin function, etc.) rank all distributions. In the social welfare maximization problems that we have done using the strict utilitarian function, the maximin function, etc. all resulted in outcomes that were Pareto-optimal. Nevertheless, as this discussion indicates, we should be wary of accepting as socially desirable any outcome whose only positive characteristic is that it is Pareto-optimal.

The Paretians argue that the use of differential rates should be avoided since the outcome will not be Pareto-optimal. (In the University of Chicago symposium issue, this view is represented by contributions by Kaplow, Samida and Weisbach, and Viscusi.) In fact, these commentators argue that to select investment projects using any criteria other than present value discounting using the market rate will result in outcomes that are not Pareto-optimal across generations.

To understand this, suppose that there are two periods, that there are a variety of investment projects with differing returns. Suppose that there are private investors who invest using the present value criterion using the market rate. Suppose that the government also invests using this same criterion. The result will be Pareto-optimal across generations. Why is this so? For illustration, suppose that both the government and private investors use a rate of 10%. That is, they adopt projects whose rate of return exceeds 10%. Is there any other combination of projects that will make one generation better off without making the other worse off? The answer is no. All unadopted projects have returns less than 10%. Suppose we decide to adopt one of these projects. This requires additional investment by the first generation, making it worse off unless we eliminate some other investment. But all of the other investments have a return of 10% or better. Substituting a sub-10% project for a 10% or better project will make one of the generations worse off.

Note that by this same reasoning, any common rate, market rate or not, will have this same effect. However, if we concede that the private sector will use the market rate, then use of the market rate by the government will insure Pareto-optimality.

The converse is also true: failure to use market rates will result in non-Pareto-optimal distributions, as long as the potential investment sets for the government and the private sector are different. To see this, suppose that the government uses a discount rate of 5%, while the rate used by the private sector is 10%. This means that all public projects of greater than 5% are being undertaken, and all private projects of greater than 10% return are undertaken.

A Pareto-improvement could be had by the government foregoing a, say, \$1,000 6%-return investment, while the private sector undertakes an additional \$1,000 9% investment. Since a Pareto improvement is possible, the original distribution of investments is not Pareto-optimal.

Again, this result does not depend on either rate being the market rate. But it does apply if, for example, the private sector uses the market rate, while the government uses a rate below the market rate.

To make this more concrete, suppose that it is possible to make a real investment \$1 million today to prevent \$2 trillion worth of environmental damage 500 years from now. Suppose that the government is considering levying a tax of \$1 million and then making this environmental investment. Finally, suppose that the real market rate of return is 3%. A Paretian would counsel us to forgo this \$1 million investment in favor of a different strategy. Rather than invest in the environmental project, the Paretians would advise us to invest in some other project. Since the market rate is 3%, there should be some investment (perhaps non-environmental) that returns close to 3%. If this were rolled over a sufficient number of times, then, after 500 years, we should have assets worth \$2.62 trillion. This \$2.62 trillion in additional consumption will more than compensate for the loss of a \$2 trillion environmental asset.⁶

For Paretians, the bottom line is this: failure by public actors (government) to adopt the market rate (or some other common rate) for all investments leads to anomalous non-Pareto-optimal outcomes.⁷ Note that the Paretians do not argue that use of market rates will lead to good distributional outcomes. Rather, they argue that failure to use market rates produces inefficiencies. Eliminating the inefficiencies would allow more to go around for both present and future generations. As for distributional issues, the Paretians invoke tax-and-transfer arguments. That is, they counsel, adopt a system that is Pareto-efficient (adopt a system using market rates), and then adjust intergenerational distribution through a tax-and-transfer program.

Two critiques to the Paretian program are important:

1. Nothing in the Paretian argument refutes the possibility that by using rates different than the market rate, the government may be able to improve inter-

generational social welfare. Even if the resulting distribution is not Pareto-optimal, it still may achieve results that are socially superior to those if the government used market rates. (Recall that above we demonstrated that non-Pareto-optimal distributions may result in greater social welfare than Pareto-optimal distributions.)⁸

2. A tax-and-transfer policy coupled with the use of market rates can produce a socially superior intergenerational allocation only through altering the market rate. In a closed economy (and the world economy is a closed economy) one can alter the intergenerational distribution only through real investment. If we insist that real investment be evaluated using market rates, as the Paretians do, then real investment can be changed only by changing the market rate. That is, taxing the first generation and putting the money into the bank for the second generation will affect real distribution only if it changes the interest rate. If it does not change the interest rate, then all of the same real investments will be made, and the amount of consumption available to the second generation will not change. Of course, tax-and-transfer policies can affect the market rate. However, as indicated above, figuring out the socially optimal rate is difficult and, even if this were possible, the ability to reset it through fiscal or monetary policy is very limited, given the plethora of other policy goals that these tools are directed at. It is highly unlikely that the government can accomplish intergenerational distribution goals in this manner.

It turns out, then, that the Paretian argument is not very telling. Applying the Paretian criterion, we could end up with a Pareto-efficient, but highly undesirable intergenerational distribution. That is, we could end up with distribution A or B, in Figure 10, rather than the non-Pareto-optimal, but much more socially desirable, distribution D.

V. Policy

If the market is unlikely to produce attractive intergenerational outcomes, then it is worth speculating whether explicit government policy can improve intergenerational distributions. It should be noted that there are three possible policy levers that the government has. First, as discussed before, the government has some control over market rates of interest. If the government can manipulate real rates downward, then there will be more investment and increased consumption in the future.

6. Note that things are different if the public actor and the private actor draw from the same investment pool. If they draw from the same investment pool, then the use of differential rates could still be Pareto-optimal. To understand this seemingly bizarre result, assume that the public actor uses a discount rate of 2%, while the rate used by the private actor is 3%. This means that all projects of greater than 2% are being undertaken. (Only the public actor undertakes projects with a return greater than 2% but less than or equal to 3%. Projects with returns greater than 3% are undertaken by both the public and private actor.) Since all projects with returns greater than 2% are being undertaken, there is no room for a Pareto improvement. Thus, the outcome is Pareto-optimal.

We will not focus on this result for two reasons. First, in capitalist societies at least, the set of public and private investment opportunities are indeed segregated. That is, most private investment opportunities are considered inappropriate for the government. Second, in a market economy, if the government did indeed engage in such a wide-ranging investment strategy, then the government would drive the market rate to its chosen investment rate and we would effectively have uniform rates.

7. There are additional anomalies peculiar to a discount rate of zero.

8. Put another way, if Pareto-optimality is judged (as it should be) over a politically feasible set of policies, rather than over a set of technologically possible outcomes, it may well turn out that a policy of differential rates is Pareto-optimal after all. That is, if we eliminate market-rate adjustments as an effective policy, it may turn out that a policy of differential rates is Pareto-optimal. That is, there may be no other *politically feasible policy* that makes one generation better off without making the other worse off.

The second lever is direct investment by the government. That is, the government can tax and then make investments that will benefit future generations. For example, the government could make investments in wind power that the private sector is currently unwilling to make at current interest rates and energy prices.

The third lever is the creation of other incentives or disincentives for investment. For example, the institution of carbon permits and the development of a carbon market effectively put a cap on the use of carbon-based fuels. More of these fuels will therefore be available for future generations.

I will consider each of these.

A. *Adjusting the Market Rate*

As discussed previously, a possible policy lever suggested by theory is for the government to adjust distribution between generations by setting the market rate at a rate that produces optimal intergenerational distribution. There are several problems, however, that make this impractical. First, let us suppose that there is some sort of social consensus on what an optimal distribution of consumption across generations is. The theory suggests that there is some interest rate that will produce this distribution. Unfortunately, the theory does not tell us what that rate is. That not only depends on the social consensus on the optimal distribution (in terms of a social consensus about the correct intergenerational social welfare function), but also knowledge of the universe of investment projects and the consequent consumption possibility set. This set and the optimization algorithm are actually more complex than the simple model given above, since projects come with all sorts of risk levels, time horizons, and patterns of expected cash flow. Because of its vast complexity, this is knowledge that we do not possess and are unlikely ever to possess.

Use of market-rate adjustments as a policy lever for intergenerational distribution adjustments is further complicated by the use of market-rate adjustments for other policy goals. The market rate affects both intragenerational investing and intergenerational investing. Of course, a social welfare function with sufficient complexity could theoretically balance these two effects. But the additional level of complexity makes the problem of adjustment all the more intractable. In addition, interest rate adjustments are common tools of macroeconomic policy used to stimulate the economy in recessionary times and to cool it in boom times. Finally, the problem of coordinating interest rate policy among many governments, each with their own reasons for controlling interest rates, adds yet another level of complexity.

In short, because of informational complexity and the plethora of other government goals, interest rate adjustments cannot effectively be employed to adjust intergenerational distribution.

B. *Direct Investment by the Government*

Even without the issue of intergenerational distribution, there are well-known arguments for government investment. Economically, governments should make direct investments in those cases where significant externalities impede private firms from making socially desirable investments. These are known as public goods. Investment in public goods is typically evaluated using cost-benefit analysis. Costs and benefits over time are discounted to the present. If the net-present value is positive, then the investment is socially attractive.

If the government makes intergenerational investments, it needs criteria in order to evaluate them. Since cost-benefit analysis based on either market rates or a zero rate is problematic, the question becomes what other criteria might be used. Again, use of any other criterion than the market rate is subject to the Paretian critique, but could also lead to superior intergenerational welfare relative to present-value discounting, using either market rates or a zero rate.

Rather than making the transition from distribution to interest rate, one could simply make the transition from distribution to project selection. A fairly minimalist way to do this is to decide on some standards for what future generations should be entitled to: clean air; a relatively stable climate; a basic amount of energy sources; a diversity of species; a relatively nontoxic environment; etc. Once goals have been established, investments could be made by government to achieve those goals in the most effective manner.

C. *Creating Incentives*

Rather than making direct investments itself, the government, through taxes, subsidies, and regulation, could intervene in the private markets to encourage (or discourage) intergenerational investing. Examples include taxes on carbon fuels, subsidies for alternative energy sources, regulation on the production of greenhouse gases, etc. Again, these interventions are subject to the Paretian critique, but also to the counterarguments. Such interventions, although they could lead to non-Pareto-optimal distributions, could improve intergenerational social welfare.

Such programs still must be evaluated, but we can apply the same arguments we did to evaluating direct government investment. That is, a policy of setting certain minimal goals (clean air, stable climate, species diversity, nontoxic environment, etc.) for a fair intergenerational distribution, and then evaluating projects in terms of the most efficient way to achieve these goals, may be the best we can do.

VI. Conclusion

So, what is the bottom line on intergenerational discounting? What are the lessons? Here are a few:

1. We should not assume that the market and the market-interest rate will produce a socially desirable

distribution of consumption among generations. Indeed, the theory indicates otherwise.

2. We should not assume that equal respect for all generations implies that a zero interest rate is socially optimal. The theory indicates that a zero interest rate is almost certainly not socially optimal.
3. We should not accept the Paretian argument that government investment must use the same investment criterion (present value discounting based on the market rate) as private investment, simply because it will produce a Pareto-optimal distribution. We should not accept this because Pareto-optimal distributions are likely to be socially inferior to the distributions produced by non-Pareto-optimal distributions. In essence, rather than looking at Pareto-optimality, we should be looking at social welfare maximization.
4. Trying to adjust intergenerational distribution by government interventions to change the market rate is informationally and practically infeasible. Even given a consensus about a fair intergenerational distribution, it is overwhelmingly difficult to know what this implies about the desired interest rate for investment. And, practically, the use of market-rate adjustments for other policy purposes renders it unavailable for intergenerational distribution adjustment.
5. Likewise, trying to adjust intergenerational distribution by requiring the government to use a different discount interest rate than the market rate for intergenerational investments is informationally infeasible. Again, even given a consensus about a fair intergenerational distribution, it is overwhelmingly difficult to know what this implies about the desired interest rate for public intergenerational investment.

Put more briefly, we are left with the idea that intergenerational justice is important to our notion of a good society, that the market will not guarantee intergenerational justice, and that broad programs of government intervention focused on interest rates will not solve the problem. We are therefore left, almost as a matter of default, with a policy that steers a middle ground between those who would do nothing (leave it to the market) and those who would adopt broad interest-based policies. This middle ground looks at distribution directly and sets certain minimal goals for future generations: clean air; potable water; sufficient energy supplies; a nontoxic environment; etc., and then analyzes the most cost-effective way of achieving those goals. The program is modest in several ways. First, while it looks at intergenerational distribution, it does not require equality or even equal respect. The distributional goals are minimalist. Furthermore, the algorithm is tractable. It does not require positing a complex social welfare function, and it does not require translating distributional goals into an optimal interest rate.

It should be noted that while this program is modest in its goals and practical in its information requirements, it could nevertheless result in significant intergenerational investments. And such a program will not be without controversy. Costs and benefits attributed to intergenerational investments could be highly contested. Indeed, we are seeing this in currently proposed environmental legislation, including a highly charged debate on cap and trade and on investments in alternative energy sources. This Article suggests that these are the very debates we should be having.