On Discounting

Tom Fiddaman
Balaton 2007

The challenge

Suppose one has a model that makes predictions about future income, based on some assumptions about resource availability and investment behavior. The model generates trajectories for per capita consumption like the following:

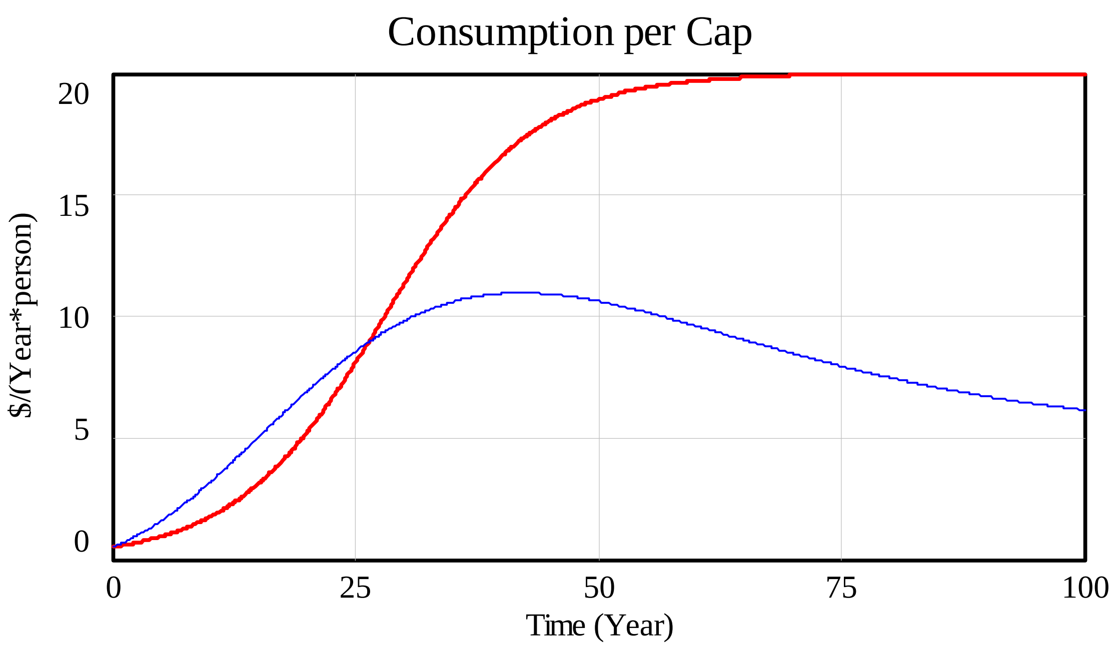

The question is, which trajectory is better? In choosing one or the other, one implicitly prefers having a little more now and a little less later, or vice versa. More generally, when weighing policy outcomes, one must make ethical choices about concern for the welfare of others who may be remote in time and space.

The process essentially involves choosing a weight associated with each dot in this familiar figure from *The Limits to Growth:*

Figure 1 HUMAN PERSPECTIVES

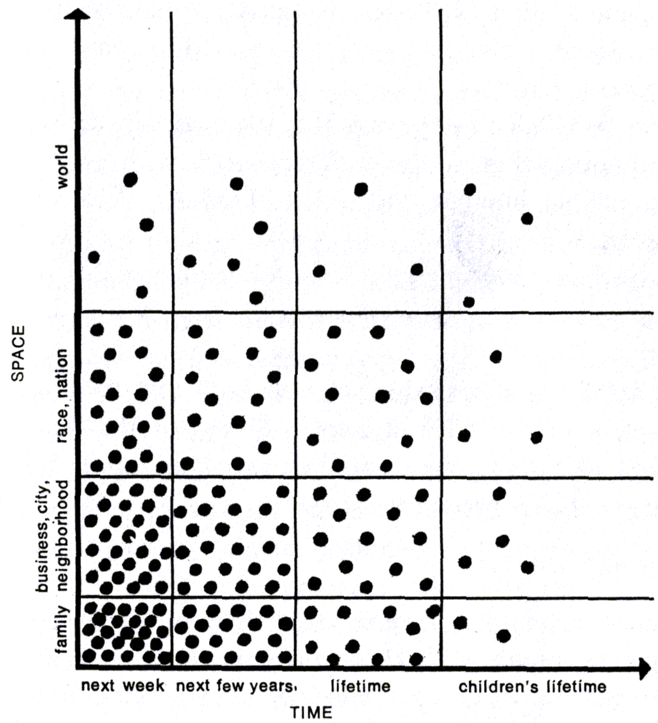

Although the perspectives of the world's people vary in space and in time, every human concern falls somewhere on the space-time graph. The majority of the world's people are concerned with matters that affect only family or friends over a short period of time. Others look farther ahead in time or over a larger area—a city or a nation. Only a very few people have a global perspective that extends far into the future.

Discounting is a way of formalizing the choice of weights, so that many scenarios can be compared quickly by a single metric, often for the purpose of letting an optimization routine find the best.

Time value of money

At its simplest, discounting simply takes account of the so-called time value of money or opportunity cost of capital, using market interest rates (perhaps adjusted for risk). Would it be better to receive $1.00 today, $1.10 next year, or 50¢ now and 55¢ next year? Theoretically, if one could invest at 10%/yr interest, one ought to be indifferent among these options. One could take the $1.00 today and invest it to yield $1.10 next year, or borrow $1.00 today and repay it with the $1.10 to be received next year, and so on.

Complications

For small decisions over short time spans, discounting is easy, because an individual or firm can take the rate of interest available in the market (e.g., from government bonds) as a given. For large projects with long time horizons, that doesn't work, because the interest rate may change over time, because population, technology, and other factors driving economic growth may influence it, and because the investments involved in the project (e.g., combating climate change) are large enough to influence the it.

In such cases, modelers typically adopt a more ambitious approach, which explicitly translates $ into welfare or utility. Almost invariably, welfare is a function of current consumption of goods and services only. Under such a framework, the discount on a $ at a point in time arises from two sources:

(a) from the shape of the utility function, which normally has diminishing returns (i.e. as consumption goes up, the happiness created by each additional unit of consumption goes down). Intuitively, if you think you will be richer in the future, you'd rather have money now because it will contribute more to your happiness while you're poor. This is often termed *inequality aversion* because it gives greater weight to the poor.

(b) from a declining weight applied to utility over time. Intuitively, this means you'd rather have a bit of happiness today, regardless. This is often termed *pure time preference* because it implies a preference for present happiness irrespective of consequences.

Inequality aversion is usually described as an elasticity, that is, in terms of the percentage change in happiness one receives from a given percentage change in consumption (assuming happiness is a function only of current consumption). Usually that elasticity is one, which means that an individual's utility = log(consumption). That also implies that Bill Gates would receive the same increment in utility from doubling his income as anyone else would.

Pure time preference is normally an exponential decay function, which declines by a constant fraction each year. Considering heterogeneity or uncertainty about individual preferences leads to the idea of a declining time preference rate, which has been implemented in several models. Still, declining discount rates typically imply that concern for just three generations forward is less than 10% of concern for today – hardly equitable.

It can be shown that, assuming a lot of things that aren't really true, an infinitely-lived benevolent agent should invest such that inequality aversion and pure time preference are consistent with the risk-free interest rate and per capita consumption growth rate observed in the market, as follows:

interest_rate = pure_time_preference + inequality_aversion*consumption_growth_rate

This relationship was the basis of several economic critiques of the Stern Review. Stern was criticized for choosing low time preference (0.1%) and “ordinary” inequality aversion (1), consistent with an interest rate lower than that observed in the market. This combination of choices inflates the social cost of carbon because at the same time it takes full account of the harm to generations in the distant future, and does not take much account of the fact that they will (hopefully) be richer than us and thus presumably better able to tolerate climate impacts. The Stern critiques failed to mention that most other studies use discount rates that are ethically indefensible, artificially deflating the value of emissions abatement, and that equitable choices result in extremely high values in cases where the economy grows slowly.

Discount behavior

The discount from pure time preference looks like the following:

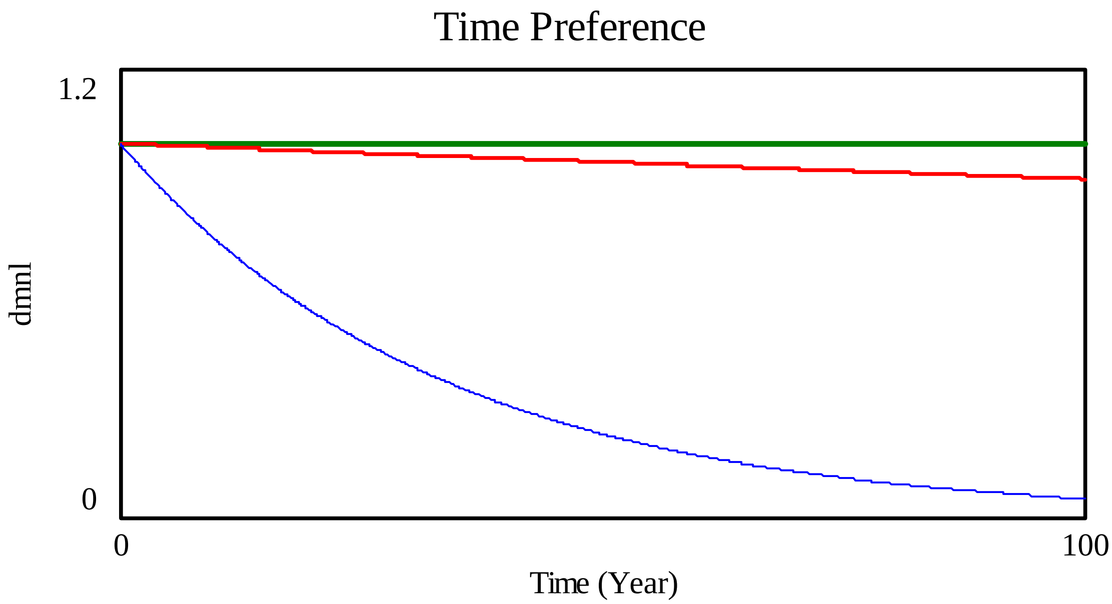

The top, level line indicates equal preference for all times, that is, equal treatment of all generations. The middle line is a time discount declining at 0.1% per year, as used in the Stern Review. The bottom line shows time preference declining at 3% per year, as used by Nordhaus in the DICE model, and in many other climate policy assessments.

The functions that translate consumption into welfare or utility are typically drawn from a small family that looks like those in the figure below, where inequality aversion describes the curvature of the shape:

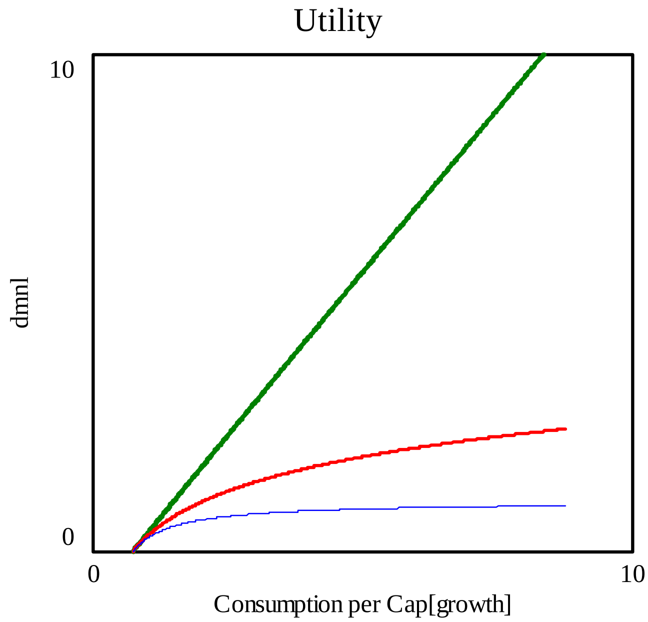

The top line shows linear utility, which is to say, a buck is a buck no matter who gets it. The middle line is logarithmic utility, as used in many models, where the incremental happiness created by additional consumption declines as consumption rises, so that 10% more consumption always adds the same increment to happiness, regardless of consumption level. The bottom line uses higher inequality aversion, so that utility saturates quickly, making both an additional $ and an additional % of consumption more useful to a poorer person.

The effective discount that arises from inequality aversion is contingent on the rate of growth of the economy (specifically, of consumption). When the economy grows rapidly, future consumption is discounted heavily, because future income contributes little at the margin to the happiness of wealthy people. On the other hand, if the economy shrinks, the discount rate in effect becomes negative, and concern actually increases over time as future generations become poorer.

The same logic can be extended to considerations of the welfare of people in regions with different income levels. Here, there is tension between idealized concepts of utility, which suggest diminishing returns to additional wealth and thus that inequity is suboptimal, and the observed gross inequities in income distribution around the world. As a technical solution, some models use Negishi weights, which are in effect an additional discount, used to adjust the welfare of different regions to equalize the utility of consumption. This procedure is the ultimate implementation of the old joke that economics is the search for a utility function that makes observed behavior optimal. The consequence, of course, is that in such a model, policy is made as if existing inequities were ordained by the laws of nature.

Conceptual problems

The very use of discounting presumes that one has a stream of costs and benefits that all have the same units of measure, i.e. money. Thus the problems with discounting are largely the same as the problems with cost benefit analysis more broadly. These include:

- Prices used to value outcomes are imperfect reflections of the underlying preferences of humanity; they include errors and biases as well as inequities in the existing allocation of wealth.
- Costs and benefits are treated as interchangeable, though in fact they may fall upon different groups. In principle the winners under some policy could compensate the losers, but in practice this seldom occurs.
- Values assigned reflect stable preferences, but preferences are not laws of the universe, and are neither stable nor uniformly known to individuals. Instead, preferences evolve over time as a social phenomenon.
- Built capital and natural capital are regarded as interchangeable. In fact future generations may find money a poor substitute for services of nature.
- Uncertainties in model structure, particularly concerning extreme and irreversible events like extinction, are seldom sufficiently represented or explored in models.
- Analyses normally take a risk-neutral stance, where a 50-50 chance of winning $2 is as good as a sure chance of receiving $1. However, risks to planet Earth cannot in fact be diversified away by spreading investments across multiple planets.
- Models often identify as attractive tradeoffs that appear bizarre to a human reviewer; in many cases it would be better to view actual time trajectories for scenarios before constructing an arbitrary score, but such subtleties seldom reach the policy arena.
- Utility functions used typically consider only flows, ignoring services from long lived capital stocks like housing.
- Utility functions are indifferent to time ordering of welfare – a feast-or-famine cycle scores as well as stable welfare, as long as there’s enough feast.

One could argue that the theoretical problems of discounting pale in comparison to the implementation problems when it is put into practice. Reviews of development projects and firm investments indicate that real projects seldom achieve the (discounted) results promised, in part because model results are not taken seriously; they are manipulated or cherry-picked for use as propaganda.

Some examples

I used Nordhaus' DICE model to conduct some quick tests of the effect of discounting choices on policy conclusions. The absolute results are subject to the biases of that model, but the relative outcomes should be directionally correct.

Value of a ton of carbon emissions avoided in year 2055

		Economic Growth		
		Slow	Moderate	Fast
Discounting	Nordhaus (3%/yr time preference, inequality aversion = 1)	$9	$13	$17
	Stern (0.1%/yr time preference, inequality aversion = 1)	$89	$114	$165
	Fair (0%/yr time preference, inequality aversion = 2.5)	$74	$63	$17

Three things are worth noting:

- Nordhaus' preferred discounting, with pure time preference declining at 3% per year, is simply too myopic to ever favor controlling carbon emissions significantly.
- The Stern Review approach yields carbon prices that are closer to those under fair discounting (as advocated by Dasgupta and others) than Nordhaus' approach.
- Fair discounting suggests that one needn't bother to control carbon if future generations are rich (instead we should worry about ourselves).

What this suggests to me, more than anything else, is that mindless reliance on discounting to choose an attractive future is dangerous. Discounting can be a useful tool for quickly sorting through many options, but ultimately one must also look at outcomes directly, in order to determine where they are driven by limitations of the model rather than reality, and how they affect interested parties at different points in time and space.

Further Reading

Richard Howarth and Richard Norgaard have shown that, when rights are given to future generations, like the right to a reasonable climate conditions, the discount rate changes. See:

Howarth, Richard B. and Richard B. Norgaard. 1992. Environmental Valuation Under Sustainable Development. American Economic Review 82(2):473-477.

Howarth, Richard B. and Richard B. Norgaard. 1993. Intergenerational Transfers and the Social Discount Rate. Journal of Environmental and Resource Economics 3:337-358.

Their articles and others proved so controversial within economics that Resources for the Future assembled a group of economists to show, not that they were wrong, but that their argument does not matter "empirically," a sad story indeed, for with respect to climate change it clearly does matter. See:

Portney, Paul R. and John Weyant. 1999. Discounting and Intergenerational Equity. Resources for the Future. Washington, D.C.

Gardiner & Ford, "Which Policy Run is Best, and Who Says So?" In Legasto, A. A., Jr., J. W. Forrester & J. M. Lyneis, eds. 1980. *System Dynamics*. TIMS Studies in the Management Sciences. Vol. 14. Amsterdam: North-Holland.

The Stern Review
www.hm-treasury.gov.uk
/independent_reviews/stern_review_economics_climate_change/sternreview_index.cfm

Sir Partha Dasgupta critique
www.econ.cam.ac.uk/faculty/dasgupta/STERN.pdf

Nordhaus critique
nordhaus.econ.yale.edu/nordhaus_stern_science.pdf