

Discounting and uncertainty: a non-economist's view

Steven Sherwood

Received: 25 April 2005 / Accepted: 20 May 2006 / Published online: 11 January 2007
© Springer Science + Business Media B.V. 2007

When asked to choose between immediate and delayed rewards or costs, both human and (other) animal respondents will equate a small reward sooner with a larger one later. This is the well-known phenomenon of individual “discounting”. The capital interest rates that prevail in market economies are a related phenomenon.

Most economists agree that “social” discounting should be applied to compare costs and benefits that accrue to society at different times due to, for example, actions to reduce environmental degradation. They cannot agree, however, on appropriate rates to use. This is important since the policy implications tend to hinge on the assumed rate, as benefits often lag investment costs by times that can exceed $1/\delta$ for plausible discount rates δ . A high discounting rate makes investments (for example, intervention to reduce future climate change) appear unattractive if rewards do not accrue very quickly, even if they eventually far exceed costs. This may seem unjust or unethical. Omitting discounting, on the other hand, can indicate that we should be engaged in radically large investments. Much discourse has ensued over discounting, often pitting economists against others. However, as long as comprehensive cost–benefit analysis is to be attempted, discounting or something equivalent cannot be avoided. A recent review of the economics literature (Spackman 2004) discusses discounting justifications and the many controversies surrounding them. Previous editorials in this publication (Heal 1997; Nordhaus 1997) have also touched upon important aspects.

Two key points will be discussed here. First, individual and social discounting differ in fundamental ways, so it is important to understand the former deeply before any inferences

S. Sherwood (✉)
Department of Geology and Geophysics, Yale University, New Haven, CT 06520, USA
e-mail: ssherwood@alum.mit.edu

can be drawn regarding the latter. The distinction is appreciated by many but not acknowledged in foundational models of macroeconomics.¹ Second, uncertainty in future events qualitatively affects how discounting should be done. While each of these points has been made elsewhere, poor cross-disciplinary appreciation of some key nuances may be hindering the type of dialog needed for everyone to agree on appropriate social discounting methods.

1 Observed discounting

The easiest way to illustrate both points is to begin by examining individual discounting more closely. One reason for delaying expenses derives from opportunity costs or the “cost of money” due to loss of competing investment opportunities. Such economic discounting probably cannot explain all observed behavior (for example that in non-human subjects, or non-economic behavior in humans). Another phenomenon known to behavioral scientists, called “time preference,” is the seemingly innate preference for rewards sooner (or costs later) even when there is no evident disadvantage to waiting. This is sometimes taken as a fundamental preference that should therefore be factored into collective, as well as individual, decisions (Baron 2000).

Observers have noticed and debated some unexpected aspects of observed individual discounting, a particularly interesting example relating to time-dependence of the discount rate. A constant discount rate would cause the future-to-present value ratio V to decay exponentially with into the future t :

$$V(t) = V(0)(1 - \delta)^t \quad (1)$$

where δ is the discount rate compounded over a unit time interval. Relative valuations made by individuals, however, tend to decay more slowly farther into the future, roughly as (Green and Myerson 2004)

$$V(t) = V(0) \frac{1}{1 + \delta t}. \quad (2)$$

This implies a discount rate that decreases into the future as a hyperbolic function; such discounting has come to be called “hyperbolic”.

This phenomenon has presented a conundrum to economists and psychologists alike, as described by Ainslie and Monterosso (2004). It means that subjects, in effect, change their answer as to which of two options they would prefer (a smaller reward at time t or larger reward at time $t + \delta t$) depending on how far into the future t is. This behavior has been taken as irrational or “inconsistent” (Baron 2000; Van Hout 1998; Weinstein and Stason 1977) and its explanation has become a topic of interest to psychologists (McClure et al. 2004). In any case it seems to require that time preference dominate over economic discounting as an explanation, since the latter would imply constant discounting assuming that competing investment opportunities remain similar over time.

¹ In the benchmark Ramsey model for maximizing inter-temporal welfare through savings decisions, no distinction is possible between individual and social discounting; overall discounting is subjective and directly connected to the market interest rate, the appearance of which therefore demonstrates discounting (Blanchard and Fischer 1989).

2 The risk avoidance model and hyperbolic discounting

But is the behavior so strange? We must ask why individuals exhibit time preference at all, if not for return on investment. A reasonable answer, given readily by non-experts as well as many experts (see, e.g., Fehr (2002)), is essentially to hedge against the risk that an asset's practical value to that individual (its "utility") may decline or disappear in the future. One could be hit by a bus tomorrow, could win the lottery so that a reward that is valuable today becomes uninteresting, could develop an ulcer and no longer have any use for that fine Scotch, etc. I will call this common-sense explanation of discounting the "risk avoidance" model. Note the key role here of the distinction (traditional in macroeconomics) between goods or assets, and utility or welfare. The risk-avoidance model makes discounting a strategy (probably an instinctive or unconscious one) for maximizing the expected utility of a personal asset, not an arbitrary preference for immediate gratification.

We may quantify this model by supposing that there is a probability density P per unit time that some unexpected event will cause a asset's utility to vanish.² According to this model, the optimal discounting rate δ is simply equal to P . This model would, at first, seem to support the traditional view that a time-dependent δ is irrational, since it is reasonable to suppose that the probabilities of various events, hence the risk P , are stationary in time.

But this ignores a key point: Not only does the subject not know if or when events will occur, she *does not know the value of* P . We must regard P as a random variable, an example of "deep uncertainty". Initially, an individual will maintain some prior probability distribution $\sigma_P(P)$ on the variable P , which will then evolve as evidence accumulates; this evolution is governed by Bayesian mathematics. All that is known with certainty is that P cannot be negative.³ The initial expectation $E(V(t))$ of future relative utility is given by

$$E(V(t)) = \int_0^{\infty} \sigma_P(P)(1 - P)^t dP \quad (3)$$

In other words, the individual must consider all possibilities as to the true risk P and average over the future utility implied by each possibility.

We may expect that, when P is uncertain, the optimal δ would equal the expectation of P :

$$\delta = \int_0^{\infty} \sigma_P(P)PdP \quad (4)$$

If the uncertainty in P goes to zero, $\sigma_P(P)$ collapses to a Dirac delta function centered on δ , and Equation (3) reduces to Equation (1), reproducing the traditional result in which uncertainty in P is not considered. However, *any other normalized prior distribution* $\sigma_P(P)$ will, according to Equation (3), produce a future expectation value of utility whose log rate of decrease declines with time. This is because $E(V)$ becomes increasingly dominated farther into the future by the contribution from the lowest possible values of P .

² One may easily generalize this to a wider class of events with variable probabilities and fractional losses of utility, in which case the situation is unchanged except P must be interpreted as the net risk per unit time, or sum over all event types of the event probability per unit time multiplied by the fraction of utility lost.

³ One may contemplate events that *increase* an asset's utility (say, losing one's job); in this case, P would still be positive but δ would be negative. Empirically, our instincts must be telling us that value-increasing events are relatively unlikely.

What is a “realistic” prior distribution $\sigma_P(P)$? If one has initially no knowledge of the relevant events (or even a clear image of what kind of events to be worried about), then one must initially consider essentially any positive value of P to be possible. An unrestricted distribution $\sigma_P(P)$, however, clearly must decrease with P rapidly enough for at least its first moment δ to be finite. The simplest distribution to do this is arguably the exponential distribution $\sigma_P = \delta^{-1}\exp(-P/\delta)$.

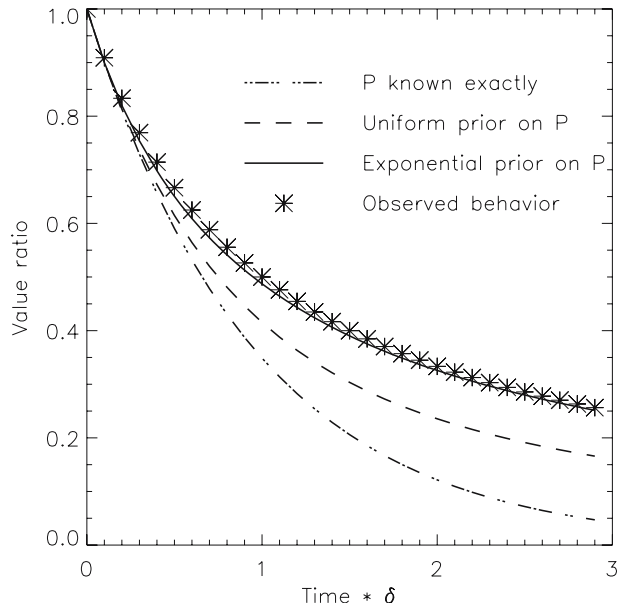
If we insert this distribution into Equation (3) and perform some simple mathematical manipulations (Sozou 1998), we obtain exactly the observed behavior Equation (2). Figure 1 shows a plot of $E(t)$ for a perfectly known risk, a uniform distribution of risk P between 0 and 2δ , and an exponential risk distribution, plotted together with Equation (2). Note that “hyperbolic” behavior also arises from a more restricted prior distribution of risk, though it is less pronounced.

3 But isn’t hyperbolic discounting irrational?

The essence of the above point – that the optimal discounting strategy when the true rate is uncertain is to employ a time-decreasing rate – has been noted in the economics literature (Sozou and Seymour 2003; Weitzman 2001). Weitzman (2001) documents in detail the divergence of expert opinion on appropriate social discount rates, arguing that time-decreasing rates should therefore be adopted, a position that is beginning to win adherents. However, previous explanations have not fully addressed the perception of inconsistency or irrationality that beleaguers hyperbolic discounting.

This perception rests on the assumption that the same choice posed now or in the future should be made the same way. But options posed hypothetically for future election will be irrelevant unless at that time the asset still retains utility; that is unlikely unless P is sufficiently small. By contrast, a decision made now must acknowledge the full range of possible P , including the possibility that P is very high – which a later, more experienced decision maker will be able to rule out. The posterior distribution of P in the future, conditional on no event

Fig. 1 Discounted value plotted against time into the future; (lines) expectation based on risk avoidance model for three priors on the risk P ; (asterisks) observed behavior fitted by Equation (2)



having occurred between now and then, will shift monotonically downward; optimal risk-avoidance discount rates will commensurately decrease. In effect, one is worried that a life-changing event could happen tomorrow, but has no idea how likely that is – and realizes that if it still hasn't happened after 10 years, it probably won't happen in the 11th year either.

We arrive at the conclusion that the observed hyperbolic behavior, far from being irrational, is optimal if we accept the risk-avoidance model of individual discounting. This would seem to support both the model and the use of hyperbolic discounting.

4 Social discounting and risk

I now turn to the perplexing question of how to discount social assets. Social discounting is sometimes expressed as the sum of a pure time preference and “consumption discounting”. These are loosely analogous to the individual time preference and economic discounting rates respectively. An empiricist seeking to quantify the two social discount rates might therefore be tempted to substitute rates inferred from individual behavior, but this would be misguided.

The rate of social pure time preference captures the idea that current inhabitants of Earth are more entitled to utility than future ones. Most commentators would set this to zero on moral grounds, although some would follow the empiricist and argue for a positive value that reflects evident individual preferences. I propose that, even if there were no moral objection, this empirical argument should be rejected because there is no evidence for pure time preference: the individual data are more elegantly explained by the risk-avoidance model, in which utility is maximized without any inter-temporal favoritism. Further, Layton and Levine (2003) analyzed the results of a survey explicitly asking respondents to compare different investment and environmental change scenarios reaching into the far future, a highly relevant study design. The authors report that results implied a discount rate of less than 1%, indicating that despite their strong instinct to discount in everyday situations, most people reject long-term discounting of human welfare.

It follows that we should abandon “time preference” and consider only discounting associated with actual utility, either through (1) any losses of expected asset utility over time or (2) the opportunity cost of tying up an asset that could be used for another purpose. Fortunately these are both tangible economic motivations rather than subjective preferences. The problem is how to apply them to the social case.

The utility of social assets can decrease for similar reasons as those of individuals, and this can again be modeled in terms of risk. The relevant “social risk” must depend on future events that would affect the utility of common assets rather than just those of an individual. The social risk must be smaller than the individual, since any event affecting society necessarily affects all members while the reverse is not true. Individuals discounting for risk-avoidance reasons would be expected to discount at a rate at least equal to the reciprocal of their expected lifespan, and indeed do so (Sozou and Seymour 2003). Past societies have greatly outlived their individual members, even if one defines societies narrowly; it would be wrong to discount common assets at high, individual rate if one hopes that the society will last many human lifetimes. Not all utility-reducing events need cause societal upheaval: Technological change can reduce utility through obsolescence. This type of risk too will be less for society than for the individual, though perhaps not by the same margin.

An important example of this “risk” is increase in wealth, which causes a given asset to be less meaningful; a poor person or nation derives greater satisfaction from the same goods than does a wealthy one. This idea is accepted by economists and dubbed “elasticity of marginal utility”. Expectations or hopes of greater future wealth (founded or not) are probably a strong

driver of individual discounting, as evidenced by rampant high-interest consumer borrowing in industrialized nations. A similar idea applies to social discounting but, again, at a much lower rate, since per-capita economic growth is usually much slower than the real (or hoped-for) growth of an individual's income during his or her own adulthood.

Note that social utility risk too is characterized by deep uncertainty. When the atomic bomb was first developed, for example, reasonable estimates of the likelihood per year of a nuclear war, or of the likely duration of civilization, could easily have spanned orders of magnitude. With the benefit of hindsight we can now rule out very high probabilities per year of nuclear war, but face new uncertainties. Risk of asset obsolescence is also deeply uncertain, as both the pace of change and its qualitative consequences are notoriously unpredictable (Tenner 1997). Finally, I would argue that future economic growth is very uncertain. Past trends in wealth have probably derived to some degree from a hitherto plentiful, but ultimately finite, natural resource base; this may be more important than we prefer to think. History bears multiple examples of human societies that flourished for a few hundred years and then suffered a scarcity-induced collapse (Polynesian settlements of Easter Island and New Zealand are particularly well-documented cases). While optimists may point to the far greater diversity of the modern global economy, it seems unwise to dismiss at least the significant possibility of a plateauing of growth on time scales relevant to environmental policy. This would, of course, imply a small or even negative contribution to the discount rate from elasticity of marginal utility.

The justification most often heard to support social discounting, however, is the opportunity cost of lost alternatives. Why invest in costly climate-change prevention measures if investing elsewhere would yield future returns that exceed the future costs of inaction? This is a central question, since the key decision is always one of allocation. It may seem straightforward that such discounting must be reckoned at the prevailing interest rate, but this is clear only for the case of individual discounting. The social problem is much more subtle, for two reasons. First, there is diminishing return on investment: extra investment will attract lower rates of return than the going average, because it must move into progressively less promising opportunities (during the investment bubble of the late 1990s marginal rates of return were surely negative). Second, any competing physical capital investments have finite lifetimes. A new building may yield social value and lucrative rents for a few decades, but will eventually be bulldozed or require substantial new investment. Its former owner (or his heirs) can then reinvest the gains elsewhere and continue to grow a personal fortune exponentially, but from that point on they would be competing with others for a fixed array of investment possibilities rather than adding to it. It would therefore be inappropriate to reckon opportunity costs – or the economic growth associated with that first house – in perpetuity; otherwise, traveling back in time and giving the ancient Greeks a few more trees or goats would, through the miracle of compound interest, turn today's Greece into the richest nation on Earth. An investment appreciation time limit of order one human lifetime might be more appropriate, although this is a naked guess. The one type of investment that unquestionably outlives all others is an investment to avoid environmental damage, if (as in the case of climate change) recovery cannot occur for tens of thousands of years.

5 Final remarks

Each one of the potential contributors to social discounting that I have so crudely discussed here shares two important characteristics. First, even in the near term, it occurs at a rate that

is arguably much lower than that of the analogous contributor to individual discounting. In particular, individual and market discount rates seem to float above the reciprocal of one human lifetime for various reasons that are irrelevant to social discounting. Second, it occurs at a rate that is either highly uncertain (in the case of utility loss) or must decrease with time (in the case of opportunity cost). And as has been shown here, uncertainty also mathematically implies time-decreasing rates. We thus conclude that, no matter what source of social discounting one prefers, it should be done hyperbolically and with an initial rate well below that appropriate for individuals. Many of the high discount rates suggested by economists in Weitzman's (2001) survey are hard to reconcile with this.

This picture stands in contrast to the traditional approach of choosing best-guess rate parameters (empirically or otherwise) and assuming that they are time-invariant. Such an approach will yield disastrously wrong predictions of utility more than a few decades into the future. Uncertainty in future events or their likelihood, and the limited lifespan of physical investments, each force future discount rates toward zero. While costs and benefits are also uncertain, this uncertainty does not have the same biasing effect on inter-temporal differences in valuation.

It is hard based on these simple arguments to quantify the discounting rate. I will however bravely (or naively) venture that the above discussion implies discounting at something close to the real, per capita, world economic growth rate (about 1% per annum) initially, decreasing into the future due to uncertainty. Interestingly, one obtains a similar result if one assumes only that the discounting rate decreases hyperbolically according to Equation (2) and with a decay time of roughly one human lifetime.

Prickly ethical questions are intertwined with cost–benefit analyses, particularly those related to how costs and benefits are distributed vs. income and vs. agency in causing the harm. These tend to thwart attempts at quantitative analysis. It may be best, as discussed by Spackman (2004), to dispense with discounted, present-value comparisons altogether for assessing environmental options. But the urge to do cost-benefit analysis will probably not recede and discounting will therefore remain central. It is important that the assumptions (stated or unstated) underlying these rates be considered carefully.

References

- Ainslie G, Monterosso J (2004) A marketplace in the brain? *Science* 306:421–423
- Baron J (2000) Can we use human judgments to determine the discount rate? *Risk Anal* 20:861–868
- Blanchard OJ, Fischer S (1989) *Lectures on macroeconomics*, Chapt. 2. MIT Press, p 650
- Fehr E (2002) Behavioural science – the economics of impatience. *Nature* 415:269
- Green L, Myerson J (2004) A discounting framework for choice with delayed and probabilistic rewards. *Psychol Bull* 130:769–792
- Heal G (1997) Discounting and climate change; an editorial comment. *Clim Change* 37:335–343
- Layton DF, Levine RA (2003) How much does the far future matter? A hierarchical Bayesian analysis of the public's willingness to mitigate ecological impacts of climate change. *J Am Stat Assoc* 98:533–544
- McClure SM, Laibson DI, Loewenstein G, Cohen JD (2004) Separate neural systems value immediate and delayed monetary rewards. *Science* 306:503–507
- Nordhaus W (1997) Discounting in economics and climate change; an editorial comment. *Clim Change* 37:315–328
- Sozou PD (1998) On hyperbolic discounting and uncertain hazard rates. *Proc R Soc Lond B Biol Sci* 265:2015–2020
- Sozou PD, Seymour RM (2003) Augmented discounting: interaction between ageing and time-preference behaviour. *Proc R Soc Lond B Biol Sci* 270:1047–1053
- Spackman M (2004) Time discounting and of the cost of capital in government. *Fisc Stud* 25:467–518

- Tenner E (1997) Why things bite back: technology and the revenge of unintended consequences. Random House, p 431
- Van Hout BA (1998) Discounting costs and effects: a reconsideration. *Health Econ* 7:581–594
- Weinstein MC, Stason WB (1977) Foundations of cost-effectiveness analysis for health and medical practices. *N Engl J Med* 296:716–721
- Weitzman ML (2001) Gamma discounting. *Am Econ Rev* 91:260–271