Sustainability Science: An Introduction[a]

Chapter 4*

Time, Risk, and the Generations

(Prepared by Partha Dasgupta: 16 April 2010)

[a] Tentative title of a book being prepared by John Bongaarts, Stephen Carpenter, William Clark, Partha Dasgupta, Robert Kates, Pamela Matson, Elinor Ostrom, John Schulnhuber, and Billy Turner III.

Chapter 2.1a, draft of 28 April 2010 in Sustainability Science: An Introduction by Partha Dasgupta et al.

Contents

An entity's characteristics are sustained if they don't decline or deteriorate over time. Chapter 4 was about ethics and economic evaluation in a timeless economy. We now extend the the discussion by introducing time explicitly. Complications are introduced sequentially. We begin by developing intergenerational ethics in a deterministic world. Small uncertainties are introduced in Section 4*.4, large uncertainties in Section 4*.5.

4*.1 Basics

Time is denoted variously by s and t, and is taken to be discrete (s, $t = ..., 0,1,2,...$). $t = 0$ is the present. As today's decisions have a bearing on future well-beings, the social evaluator of Chapter 4 needs to study the economy from the perspective of not only the present date, but also each future date.[1] Whenever the vantage adopted by her is t, we denote the dates following t by the index s, so that $s \geq t$. On occasion social evaluators are called upon to undertake retrospective studies. To allow for that, let T be a positive integer. s and t are assumed to range over the values $-T$, ..., 0, 1, 2,.... In numerical exercises, the unit of time will be taken to be a year.

Let $N(t)$ be population size at t. $N(t)$ is assumed to be an exogenous variable, meaning that it is unaffected by policy. At a deep level this is no doubt a bad assumption, but economic policies are most commonly discussed and debated without any mention of their demographic impact. We follow that practice.[2]

Assume, as so many who write on the economics of the environment do, that social well-being at each t is the sum of the $N(t)$ individuals' well-beings at t. In Chapter 4 it was argued that, as a measure of social well-being summation of individual well-beings is appealing. Equation (4.4) there reflected that viewpoint. For expositional ease, we now simplify equation (4.4) by imagining that the determinants of well-being can be aggregated into a single numerical index, which we call consumption.[3] At each t individuals are numbered from 1 to $N(t)$. Let $C_i(t)$ and $U_i(C_i(t))$ be individual i's consumption level and the flow of well-being at t, respectively. We denote the flow of social well-being at t by V_t, where t appears as a subscript so as to distinguish social well-being at t from intergenerational well-being at t (the latter to be defined presently). Following equation (4.4), social well-being at t is

[1] Economists frequently use the term "society" to refer to our social evaluator. It's a misuse. Societal decisions, at least in ideal democracies, are reached via the ballot box, where each citizen casts his or her vote on the basis of their personal conception of social well-being. In the welfare economics of climate change, such authors as William Cline and Nicholas Stern have meant the "global community" when speaking of society (Cline, 1992; Stern, 2006).

[2] Population ethics remains an unsettled subject. See, for example, Fishkin and Goodin (2010).

[3] This rules out the influence of habitual consumption on personal well-being, but of course all the concepts we develop and the formulae that are derived can be extended to cover the case where habits matter. The implications of habitual consumption on savings decisions have been studied by Ryder and Heal (1973).

$V_t = {}_i\Sigma[U_i(C_i(t))]$.[4] (4*.1)

Appearances to the contrary, the formulation does not rule out i caring about his friends and relatives; nor does it rule out the "demonstration effect", mentioned in Chapter 4. $C_i(t)$ is an aggregate, constructed not only out of i's consumption of goods and services, but also out of those others whose consumptions affect his personal well-being, whether benignly (as with the consumptions of those he cares about) or adversely (as in the case of status goods).[5]

We now aggregate the V_ts over time as as to arrive at intergenerational well-being. If the social evaluator is to ensure consistency in her reasoning, she requires an ethical theory that can be a guide no matter whether the date is now or in the future. For example, it could be that the evaluation is being conducted today, but the social evaluator wants to view the economy from the perspective of some future date so as to check the way she would assess matters or prescribe policies at the later date. We therefore adopt the convention that the evaluation is conducted at an arbitrary date t.

Sustainability science involves the study of economic possibilities over the long run. But how long is the long run? Consider the household. In making consumption and saving decisions, parents include their children in their reckoning. But because they care about their children's well-being, they cannot ignore the thought that their children will care about their children; that their grandchildren will care about their great grandchildren, and so on, down the generations. Implicitly, then, thoughtful parents include the well-beings of all their descendents in their reasoning.

At the level of the State, matters differ. Five-year Plans were customary in the former Soviet Union. They remain so in India. In the 1960s, the Perspective Plans of the Indian Government's Planning Commission were constructed with a time horizon of 25 years. The planners didn't suppose India would cease to exist at the end of 25 years. What they did instead was to set a target for the capital stocks that were to be accumulated by the end of the period and proceeded to determine a consumption and investment plan for the 25 years that would be consistent with their target. But to set a target for the end of the horizon is to presume something about the well-beings of people beyond the horizon, otherwise there would be no basis for the choice of target. Our social evaluator faces a similar quandary. To settle on a finite horizon is problematic, because no matter how far into the future she draws it, there is a chance that the economy will exist beyond that date. An infinite horizon suggests itself.

---

[4] In equation (4*.1), i assumes values from 1 to $N(t)$.

[5] Suppose people experience the "demonstration effect" (Chapter 4). Then $C_i(t)$ in equation (4*.1) is a deflated value of i's actual consumption at t. To illustrate, if everyone actually consumes \$50,000 worth of goods and services, the $C_i(t)$ that appears in the equation would be less - say, \$30,000 worth of goods and services. That such an aggregate can be constructed is a heroic assumption (see Arrow and Dasgupta, 2009, for ways in which such a construction could be attempted). We make it here nonetheless so to simply the exposition.

How is the social evaluator to know what the well-being functions of people in the distant future will be like? She doesn't, of course. But just as parents choose their saving policy without knowing their children's character when they will turn into adults, the social evaluator chooses the best she can, even though she has little idea of the character of people in the distant future.

So, we are to develop a conception of intergenerational well-being in a world where the future is indefinitely long and the well-beings of people in the deep future are unknown. To cut through the haze, we begin by considering a world where the horizon is infinite but where there is no uncertainty about future well-being functions (but see Sections 4*.4 and 4*.5). Such a stripped down model will allow us to get to the essentials without fuss.

Denote intergenerational well-being at t by $V(t)$. As with V_t, far and away the most widely used expression for $V(t)$ in the literature is additive, but unlike V_t the summation is conducted with possibly declining weights. Thus,

$$V(t) = V_t + V_{t+1}/(1+\delta) + V_{t+2}/(1+\delta)^2 + ... = {}_t\Sigma^{\infty}[V_s/(1+\delta)^{(s-t)}],$$

which, on using equation (4*.1) is

$$V(t) = {}_t\Sigma^{\infty}[{}_i\Sigma U_i(C_i(s))/(1+\delta)^{(s-t)}], \qquad \delta \geq 0.^6 \qquad (4^*.2)$$

In expression (4*.2), intergenerational well-being is the weighted sum of the flow of social well-being at each t, where the weights decline at a constant percentage rate δ. In contrast to what we have named social well-being (V_t), which is a flow, intergenerational well-being ($V(t)$) is a stock. In the economics literature, δ is frequently called the "social discount rate". But because δ is the rate at which future well-beings are attenuated, it is more accurate to call it the well-being discount rate, which is what we will do here. As the economy we are studying is deterministic, δ can be interpreted as a measure of "impatience": future well-beings are discounted simply because they are future well-beings.

Moral philosophers insist that future well-beings should not be discounted. If we follow their stricture, $\delta = 0$ in expression (4*.2).[7] However, Koopmans (1960, 1972) remarkably showed that δ would have to be positive if we accept a set of otherwise intuitively appealing ethical axioms on the distribution of well-being across an infinity of generations. One can re-state Koopmans' theorem as saying that $\delta = 0$ in expression (4*.2) is logically inconsistent with several other ethical principles we may care about and would want $V(t)$ to satisfy. As is the case so often elsewhere, ethical principles clash here and something has to give. In Chapter 4 it was argued that

[6] We assume for the moment that the infinite sum converges. The question of convergence was much studied by economic theorists in the 1960s. Arrow and Kurz (1970) contains an account. The problem is studied in Section 4*.5.

[7] Ramsey (1928: 261) famously wrote that to discount future well-beings is "ethically indefensible and arises merely from the weakness of the imagination." That, of course, is not an argument; merely an expression of one's beliefs. Broome (1992) contains a summary of the arguments that support Ramsey's position.

ethical pluralism encourages us to allow for a certain amount of give and take among principles. We are able to do so now, because Koopmans' axioms do not yield a specific value of δ; they only imply that $\delta > 0$. So we assume in what follows that δ is a small positive number, as close to zero as the social evaluator desires.[8]

If the social evaluator's vantage point is the present ($t = 0$), expression (4*.2) reduces to

$$V(0) = V_0 + ... + V_t/(1+\delta)^t + ... = {}_0\Sigma^\infty[V_t/(1+\delta)^t]$$
$$= {}_{t=0}\Sigma^\infty[{}_i\Sigma U_i(C_i(t))/(1+\delta)^t], \qquad \delta > 0. \tag{4*.3}$$

A useful way to interpret δ in expression (4*.3) is to imagine that from the ethical point of view objective intervals of time shrink at rate δ as the social evaluator peers further and further into the future. There are psychological explanations along these lines for why people display impatience. What Koopmans showed was that widely accepted ethical principles justify impatience, they don't merely have to accommodate it.

4*.1.1 Constant Population

Many environmental concerns involve the long run, global climate change being the most prominent example. Demographers expect world population to stablize in the long run. So we begin by assuming that population size is a constant, N.

In order to focus on an economy's movement over time, we simplify even more by abstracting from intra-temporal ethics. We do that by adopting a standard formulation in economics, in which social well-being at t is a function of an aggregate index of consumption at t. Writing the latter as $C(t)$, social well-being at t is denoted as $U(C(t))$. One way to justify the move is to imagine that everyone has the same well-being function at every date and that institutions have been so designed that at each date consumption is distributed equally among all. The latter assumption can be given an ethical underpining if, as in Chapter 4, we assume U to increase at a diminishing rate with C_i. In that case the desired aggregate is per capita consumption. In this reduced version, V_t (i.e., ${}_i\Sigma U_i(C_i(t))$) in equation (4*.3) is replaced by $NU(C(t)/N)$. As N is a constant, we may as well drop it by pretending $N = 1$ and write intergenerational well-being at t as

$$V(t) = U(C(t)) + U(C(t+1))/(1+\delta) + ...$$
$$= {}_t\Sigma^\infty[U(C(s))/(1+\delta)^{(s-t)}], \qquad \delta > 0.^9 \tag{4*.4}$$

Very commonly, U is taken to have the simple form,

$$U(C) = C^{(1-\eta)}/(1-\eta), \qquad \text{for } \eta > 0,$$

[8] Subsequently we will augment δ by the hazard rate for Humanity's extinction through natural causes.

[9] Expression (4*.4) formed the basis for a classic analysis of optimum national saving by the economist/philosopher/mathematician, Frank Ramsey. Ramsey's article (Ramsey, 1928) has been the starting point of nearly all studies on intergenerational ethics by economists. Expression (4*.4) has been commonly used in the economics of climate change. See Cline (1992), Nordhaus (1994), and Stern (2006), which we discuss below.

and $\quad U(C) = \log(C)$, $\quad$ corresponding to $\eta = 1$. $\quad$ (4*.5)

Figures 4*.1-4*.3 depict $U(C)$ for alternative values of η. $U(C)$ is bounded above but unbounded below if $\eta > 1$; it is bounded below but unbounded above if $\eta < 1$, and is unbounded both below and above in the logarithmic case. Those features will play an important role when we come to discuss the social evaluator's attitude to risk (Section 4*.4).[10]

Formula (4*.5) is useful in sustainability science because the well-being funtion $U(C)$ is defined by a single parameter, η. We can obtain insights by varying η so as to study the way it affects our conception of sustainable development. But simplicity comes at a price. Unless we are careful not to use formula (4*.5) willy-nilly in economic models (e.g., models of the economics of climate change), it can be a source of paradoxes (Section 4*.5). Formula (4*.5) is useful for pedagogical purposes, but should be used sparingly and with care.

Notice that the larger is η, the greater is the curvature of $U(C)$. In Chapter 4 (Section 4.5) it was noted that if the social well-being function in a timeless economy is expression (4.4), the curvature of U is a measure of the social evaluator's aversion to consumption inquality among people. A comparison of expressions (4.4) and (4*.4) tells us that larger values of η imply greater aversion to consumption inequality across people at different dates. Later we confirm that η is also a measure of the aversion toward future risk - the larger is η, the greater is the aversion to risk. η assumes several roles in a compact way. That is its appeal.

4*.1.2 Variable Population

Sustainability analysis frequently involves studying how matters are likely to be (or have been) over a brief period of time. For example, Arrow et al. (2004) studied macroeconomic data that covered the period 1970-2000. In such fine grained analysis population cannot be assumed to remain constant. We now adapt expression (4*.4) for variable populations.

Population can be regarded as a form of capital asset, albeit an asset that is the seat of human well-being, not just a means for protecting and promoting human well-being.

---

[10] When we perform economic calculations later in this chapter, readers will note that neither the level nor the scale of U has any operational content. If instead of U, the social evaluator was to suppose that the social well-being function is $7U+30$, say, none of her economic calculations would alter (see for example, equations (4*.9)-(4*.10)). Formally, this is to say that if U is chosen as the index of social well-being, so could $aU+b$ be chosen, where a and b are constants, and $a > 0$. Thus, the sentence, "social well-being is twice as high in social state x than it would be in y" is operationally meaningless. What does have meaning is a comparison of differences in social well-being (e.g., "the difference in social well-being in social states x and y is thrice as large as the difference in well-being between social states w and z). Mathematicians would say that U is unique upto positive affine transformations. In this aspect, social well-being is similar to temperature, where temperature measured in the Centigrade scale (T_c) is related to temperature measured in the Fahrenheit scale (T_f) by the formula, $T_f = (9/5)T_c + 32$. If someone was to say that the high temperature today is 1.2 times what it was yesterday, we would need to know what scale he was deploying in order to understand him. However, here is a scale-free sentence: "The difference between yesterday's high and low temperatures was twice the difference between the high and low temperatures the day before."

Demographic change introduces complications in intergenerational ethics because the size of each cohort in the population changes. For simplicity, let us continue to assume that cohorts are identical. Then population size, $N(t)$, would be a sufficient demographic statistic at t. Following current demographic thinking, we imagine that $N(t)$ will stabilize at some level in the long run, but that in the immediate future it will continue to grow.

How should variable population enter intergenerational well-being? Should the criterion for sustainability and policy analysis be total well-being or should it be average well-being? If the latter, what might an average over a time-varying population mean?

In expression (4*.4) intergenerational well-being is the present discounted value of the flow of social well-being at each date. In expression (4*.1) social well-being at a given date is the sum of the well-beings at that date. So, intergenerational well-being in expression (4*.4) is an attenuated form of the total well-being of all who will ever exist.

But the formulation isn't uncontroversial. Even though moral philosophers have insisted that it is total well-being with no attenuation ($\delta = 0$) that matters (Sidgwick, 1907), welfare economists have sought to justify average well-being for policy analysis (Gottlieb, 1945). In his classic work on the theory of optimum economic development, Koopmans (1965) took well-being at s to be the average well-being per person at s. Write $Z(s) = C(s)/N(s)$. In Koopmans' formulation, intergenerational well-being is the present discounted sum of the flow of average well-being per person; that is,

$$
\begin{aligned}
V(t) &= U(Z(t)) + U(Z(t+1))/(1+\delta) + \ldots \\
&= {}_t\Sigma^{\infty}[U(Z(s))/(1+\delta)^{(s-t)}], \qquad \delta > 0. \qquad (4^*.6)
\end{aligned}
$$

But there is a problem with expression (4*.6), to which Meade (1955) had drawn attention: If population grows, the criterion discriminates against future people merely on the grounds that they will be members of generations of larger size. That does not feel right. People should matter at each moment in time.

The ethical framework that was constructed in Chapter 4 takes individuals to be the locus of analysis, not the average individual. Applying expression (4*.3) to the case here, intergenerational well-being takes the form

$$
\begin{aligned}
V(t) &= N(t)U(Z(t)) + N(t+1)U(Z(t+1))/(1+\delta) + \ldots \\
&= {}_t\Sigma^{\infty}[N(s)U(Z(s))/(1+\delta)^{(s-t)}], \qquad \delta > 0.^{11} \qquad (4^*.7)
\end{aligned}
$$

Despite its intuitive appeal, expression (4*.7) suffers from a weakness. Suppose it is applied to cases where future population size can be controlled. It can be shown that in certain circumstances the formula advocates pro-natalism even when additional numbers would reduce the well-being of poor people who already exist. Parfit (1984, 1990) called it the Repugnant

[11] Expression (4*.7) was proposed by Meade (1966) and adopted by Dasgupta (1969) and Arrow and Kurz (1970) for policy analysis, and by Arrow, Dasgupta, and Mäler (2003) for sustainability analysis.

Conclusion. Sensing that, economists have traditionally been drawn to a formulation that reflects the average well-being of people (Gottlieb, 1945). We have seen however that the version reflected in expression (4*.6) won't do. The problem with expression (4*.6) is that it doesn't really represent average well-being. The true average would be a population average of expression (4*.7), namely,

$$V(t) = {}_t\Sigma^{\infty}[N(s)U(Z(s))(1+\delta)^{-(s-t)}]/[N(s)(1+\delta)^{-(s-t)}], \qquad \delta > 0. \tag{4*.8}$$

In words, expression (4*.8) is the present discounted value of the flow of social well-being at each date, divided by the present discounted value of population size at each date. $V(t)$ in expression (4*.8) is intertemporal average well-being at t.

If the social evaluator is conducting policy analysis, she has to compare alternative policies at a given moment in time, say t. But because by assumption policies don't affect future numbers, the denominator in expression (4*.8) is the same no matter which policy is considered. That means for policy analysis it's irrelevant whether the social evaluator uses expression (4*.7) or expression (4*.8): she will reach the same conclusion no matter which of the pair she happens to use.[12] But the same does not hold for sustainability analysis (Dasgupta, 2001 [2004]). When in Chapters 5 and 5* we come to formulate the concept of sustainable development, we will see that expression (4*.8) is the ethically appealing formulation of intergenerational well-being.

The ethical basis for policy analysis was settled among economists long before the concept of sustainable development was introduced in the literature. As our ethical intuitions on intergenerational justice have been formed by repeated exercises in policy analysis, social scientists haven't had to ask which of expressions (4*.7) and (4*.8) should be used. That may be why expression (4*.8) does not appear in any textbook on public policy or theoretical welfare economics or moral philosophy. But in sustainability analysis expression (4*.8) is the one to use.

In recent years, our obligations to future generations have been much studied in the context of the economics of global climate change. We use that literature to develop some key concepts in intergenerational ethics. For simplicity of exposition we revert to a world with constant population. Such concepts as consumption discount rates (see below) can easily be adapted for use in worlds where population varies with time.

4*.3 Discounting Climate Change

The concentration of carbon dioxide in the atmosphere stood at approximately 280 parts per million (ppm) for some 11,000 years until the early 18th century, but is now nearly 390 ppm. (We ignore the concentration of methane, which is another greenhouse gas.) If current trends in carbon emissions continue, its concentration in the atmosphere is expected to reach 500 ppm by

[12] To see why, if at time t policy A is superior to policy B when they are compared in terms of, say, $V(t)$ in expression (4*.7), A would remain superior to B if instead the criterion function were to be $aV(t)$, where a is a positive constant. Expressions (4*7) and (4*8) would be related the same way if we set $a = 1/[N(s)(1+\delta)^{-(s-t)}]$.

the middle of this century, and could reach as high a figure as 750 ppm (which is nearly thrice the pre-industrial level) by year 2100. A doubling of present day carbon concentration is expected to give rise to an increase in the mean global atmospheric temperature by 3 to 7 degrees Celsius. With a trebling of concentration, it could rise by 6 to 11 degrees. The temperature that would result even if the rise were limited to 3 degrees is beyond anything that has been experienced on Earth in the past million years.

As noted in Chapter 3, Earth's system is driven by a myriad of interlocking non-linear processes that run at differing speeds and operate at different spatial scales. The speed of changes in the global climate we should expect from further increases in carbon concentrations is particularly significant, because rapid change would make a good portion of our capital assets less than useful long before their planned obsolescence. Some of our infrastructure is expected to disappear under the rising seas. In order to restructure our assets, Humanity will need to invest in mitigation and adaptation, diverting resources from consumption. If we add the impact of rapid climate change on ecosystems (such as changes in the disease environment to which human populations are not immune; and degradation in the composition, geographic distribution, and productivity of ecosystems), the potential costs begin to look huge. Nevertheless, when in 2004 eight eminent economists were invited to Copenhagen to offer advice on how the world community could most usefully spend \$50 billion over a 5-year period, they placed climate change at the bottom of their list of ten alternatives.

4*.3.1 Consumption Discount Rates

Why did the economists do that? One reason was that they discounted future consumption costs and benefits at a positive rate. To understand why the social evaluator may also find it reasonable to discount future benefits and costs, we need a formal definition of (social) discount rates. For simplicity of exposition, assume that population is constant. So we return to the formulation in Section 4*.2.1 (expression 4*.4).

Imagine that the social evaluator has made a forecast of future consumption, which we write as $\{C(0), C(1), ..., C(t), ...\}$. She now conducts a thought experiment round that forecast by asking how much consumption she would demand to be added to $C(t+1)$, in compensation for a reduction in $C(t)$ by one unit, other things remaining unchanged. Denote that additional consumption, less unity, by $\rho(t)$. We call $\rho(t)$ the consumption discount rate. The name is appropriate because the social evaluator would demand $(1+\rho(t))$ units of additional consumption at $t+1$ as a price for giving up one unit of consumption at t; which amounts to saying that she regards an additional unit of consumption at $t+1$ to be worth $1/(1+\rho(t))$ units of additional consumption at t.

In economics, $\rho(t)$ is sometimes called the "social discount rate" at t. We avoid doing that because the term doesn't reveal what is being discounted. The name we are adopting here is appropriate because it draws attention to the fact that $\rho(t)$ represents the trade-off (expressed in percentage terms) between consumption between dates t and t+1.[13] Our analysis has shown that the term "social rate of discount" doesn't convey any meaning unless the unit of account is specified. If well-being is the unit in which all economic quantities are measured, then δ is the social rate of discount; but if consumption is the unit of account, then $\rho(t)$ is the social rate of discount. Notice that although δ has been assumed to be constant, $\rho(t)$ is not constant unless consumption is the same at t and t+1. The difference between the well-being discount rate and consumption discount rates does not pose any problems of consistency in evaluation exercises. Economic evaluation is unaffected by the choice of the unit of account. In economics the unit of account is called numeraire. Chapter 5 explains why any economic object could be made to serve as numeraire.

When mention is made of discount rates, people frequently assume they must be positive numbers. The assumption is wrong. However, it is useful to study the effect positive discounting has on policy evaluation before proceeding to show that under certain scenarios consumption discount rates can even be negative. It proves best to work with a specific environmental problem. Let us apply discounting to global climate change.

Reducing global carbon emissions or investing in technologies for carbon sequestrtation would involve huge costs now, but the benefits from averting economic disruptions would be enjoyed only 50 to 100 years from now. Long-term interest rates on government bonds in the United States have been 3-5% a year. When economists there evaluate public projects, they typically use such a figure to discount future benefits and costs. They regard the figure as the "opportunity cost of capital", the term being applied to the rate of interest that could be earned by investing in government bonds rather than in the project whose benefits and costs are being evaluated. At discount rates of 3-5%, though, consumption benefits in the distant future look minute today. If you discount at 4% a year, a dollar's worth of additional consumption benefits 100 years from now would be worth less than 3 cents today; which is another way of saying that as a price for giving up $1 worth of consumption today, you would demand that more than $30 worth of consumption be made available 100 years from now. A number of economic models of climate change have revealed that if you use an annual discount rate of, say, 4%, the costs (which are negative benefits) are greater than the sum of the discounted benefits from curbing net carbon emissions. Doing something about climate change now, the calculations imply, would be to throw money away in a comparatively bad project.

4*.3.2 Why Positive Discounting?

---

[13] If, as we would wish to do in practical exercises, were to consider both intra- and inter-generational well-being, expression (4*.3) would be the appropriate formula round which to develop the concept of consumption discount rates. In order to create a set of consumption discount rates, we would have to specify a person whose consumption is to serve as the numeraire. The social evaluator would then have a set of person-specific consumption discount rates.

Should the social evaluator discount future consumption benefits and costs at a positive rate?

There are two reasons why it may be reasonable for her to do so. First, a future benefit would be of less value than that same benefit today, if early consumption is favoured over delayed consumption, simply because a delay is a delay. Justification was offered in Section 4*.2 for choosing $\delta > 0$, albeit it was argued that δ should be small. Secondly, considerations of justice and equality demand that consumption should be smoothed across the generations. So, if future generations are likely to be richer than us, there is a case for valuing an extra dollar's worth of their consumption less than an extra dollar's worth of our consumption, other things being equal. Rising consumption (more accurately, rising consumption per capita in a world with changing population size; expression (4*.8)) provides a second justification for discounting future costs and benefits at a positive rate.

To provide a quantitative feel, we use the raw definition of $\rho(t)$ given above to determine the consumption discount rate at t. A simple manipulation of expressions (4*.4) and (4*.5) yields

$$1+\rho(t) = (1+\delta)(1+g(C(t)))^{\eta}, \qquad (4*.9)$$

where $g(C(t))$ is the percentage rate of growth in consumption along the economic forecast (i.e., $1+g(C(t)) = C(t+1)/C(t)$).[14]

In expression (4*.9), δ reflects the first reason we gave for positive discounting, while the term $(1+g(C(t)))^{\eta}$ reflects the second reason. Notice the way δ, η, and the forecast, $g(C(t))$, together determine $\rho(t)$. Notice also that $\rho(t)$ is an increasing function of δ and $g(C_t)$, respectively, but is an increasing function of η if and only if $g(C(t)) > 0$.

We have highlighted the qualifier "if and only if" for a good reason. In studying long run economic development, it has become a habit among economists to confine attention to forecasts in which consumption increases indefinitely. Equation (4*.9) says that if $g(C(t)) > 0$, δ and η play similar roles in the determination of $\rho(t)$, in that, a higher value of either parameter would reflect a greater aversion toward consumption inequality. To confirm, note that even a small increase in $g(C(t))$ would raise $\rho(t)$ so as to stiffen the criterion for supporting any further increases in $g(C(\mathrm{t}))$. That may explain why it has been commonly assumed that, as in the case of η, higher values of δ reflect a greater ethical concern for consumption equality. But if $g(C(t)) < 0$ (i.e., consumption is expected to decline between t and $t+1$), δ and η assume diametrically opposite features. Higher values of δ raise $\rho(t)$, implying an ethical preference for even greater inequality

---

[14] Proof: Recall the thought experiment that led to the definition of $\rho(t)$. Let $\Delta C(t+1)$ and $\Delta C(t)$ be a pair of small variations in consumption at dates t and $t+1$ that leave the magnitude of $V(t)$ in expression (4*.4) unaltered. Then

$$[dU(C(t))/dC(t)]\Delta C(t) + [dU(C(t+1))/dC(t+1)]\Delta C(t+1)/(1+\delta) = 0. \qquad (4*a)$$

By definition,

$$\rho(t) = -\Delta C(t+1)/\Delta C(t) - 1, \qquad (4*b)$$

where $\Delta C(t+1)$ and $\Delta C(t)$ satisfy equation (4*a). Now use equations (4*.5) and (4*a)-(4*b) to obtain equation (4*.9) in the text.

in consumption across the generations; whereas higher values of η reduce $\rho(t)$, implying an ethical preference for lessening that inequality.

A useful approximation to equation (4*.9) can be obtained if δ and $g(C(t))$ are both small. So, suppose they are small. Then equation (4*.9) approximates to the form,

$$\rho(t) = \delta + \eta g(C(t)).^{15} \tag{4*.10}$$

As noted earlier, moral philosophers insist δ should be zero. They argue that to choose a positive value for δ is to favour policies that discriminate against future generations merely on the grounds that they are not present today. Philosophers also say that values frequently in use among economists, ranging as they do between 2-4% a year, are way too high. Let us choose δ to be so small as to be be negligible. In that case we are left with only the second reason for discounting future costs and benefits (which is reflected in the second term on the right hand side of equation (4*.10)). But if rising consumption provides the social evaluator with a reason for discounting future consumption benefits at a positive rate, declining consumption would provide her with a reason for discounting future consumption benefits at a negative rate.

Economists use positive values for consumption discount rates in their models of climate change. They do so because the models assume that global consumption (per head) will continue to grow over the next 150 years and more, even if net emissions of greenhouse gases follow current trends. That however is to assume that climate change is expected to pose no serious threat to the global economy. But an increase in the mean global temperature by 3-5 degrees Celsius would take the biosphere into a climatic zone not visited in millions of years on Earth. The possible consequences of such changes to our productive base are so huge, that it isn't to be an alarmist to question forecasts of continual economic growth even after Earth enters that zone. Suppose you fear that if nothing substantial is done today to discover ways to sequester carbon or to find alternatives to fossil fuels as sources of energy, there is a sizeable chance that global consumption per head, suitably weighted across regions and income groups, will decline - owing, say, to a big increase in the frequency of extreme weather events, more severe droughts in the tropics, the emergence of new pathogens, and degradation of vital ecosystems. In that case, as we confirm below, you should use negative rates to discount future consumption benefits. Notice though that applying negative rates amplify benefits in the distant future when viewed from the present, it doesn't attenuate them.

Let us perform a quick calculation to get a feel for orders of magnitude by using equation (4*.10). Assume there is no future uncertainty. Empirical evidence from societal and personal choices suggests that η is in the range 1 to 3. For concreteness, let us work with 3. Using equation (4*.10), that means $\rho(t) = \delta + 3g(C(t)$. Following the advice of moral philosophers, let

[15] Proof: Take the logarithm of both sides of equation (4*.9). Using the fact that if a scalar number m is small in absolute value, $\ln(1+m)$ is approximately equal to m, equation (4*.10) follows.

$\delta = 0$. Now imagine that carbon emissions follow the trends that are expected under "business as usual". Consider a scenario in which global consumption (per capita) increases at an annual rate of 0.5% for the next 50 years but declines at 1% a year for the following 100 years. Under that scenario, the social evaluator ought to discount future consumption benefits at 1.5% a year for the next 50 years (3 times 0.5) and at minus 3% for the subsequent 100 years (3 times minus 1). A simple calculation now shows that a dollar's worth of additional consumption 150 years from now is worth $9 of additional consumption today. To put it another way, the social evaluator (by whom we now mean the global community) should be willing to forego $9 worth of additional consumption today for an extra dollar's worth of consumption benefits 150 years in the future. The calculation reverses the message that has been conveyed by economic models of climate change.

There should be little doubt that private investors would be using a positive rate to discount their personal earnings even under the above scenario. They would be doing so, because the interest rate offered by commercial banks on deposits would most likely remain positive. But there is no contradiction here. Under "business as usual", the atmosphere is an open access resource (Chapters 5 and 6). So long as people are free to emit carbon dioxide, there will be a wedge between private rates of return on investment and the rates the global community ought to use to discount collective costs and benefits. The former could be positive even while the latter is negative. That wedge is a reason for controlling carbon emissions into the atmosphere and bringing the two rates closer to each other; it isn't a reason for claiming that the problem of global climate change should be shelved for the future.

4*.3.3 Canonical Examples

The most-preferred values of δ and η in Cline (1992), Nordhaus (1994), and Stern (2006) are:

Cline: $\delta = 0$; $\eta = 1.5$

Nordhaus: $\delta = 3\%$ a year; $\eta = 1$

Stern: $\delta = 0.1\%$ a year; $\eta = 1$

Notice how close the authors are in their choice of η. Notice also that Cline and Stern are close in their specifications of δ, but that Nordhaus is an outlier in his choice of δ. To say that $\eta = 1$ is to say that any proportionate increase in someone's consumption ought to be of equal social worth to that same proportionate increase in the consumption of anyone else who is a contemporary, no matter how rich or poor that contemporary happens to be. It is also to insist that, if in addition $\delta = 0$, any given proportionate increase in consumption today ought to be of equal social worth to that same proportionate increase in consumption at any future date, no matter how rich or poor people will be at that future date. Taken at face value, though, it isn't immediate whether such tradeoffs are ethically reasonable. The only way to tell is to run numerical tests on simple models of economic development. It can be shown (Dasgupta, 2008) that the pair (δ=0, η=1) can recommend bizarre policies in simple models of consumption and saving.

The point estimate of consumption growth under business as usual in Stern (2006) is $g(C(t))$ = 1.3% a year. Using this in equation (4*.10) implies

$\rho(t)$ = 2.05% a year for Cline

$\rho(t)$ = 4.30% a year for Nordhaus

$\rho(t)$ = 1.40% a year for Stern

4.3% a year may not seem very different from 1.4% a year, but is in fact a lot higher when it is put to work on the economics of the long run. Just how much higher can be seen from the fact that the present-value of a given loss in consumption, owing, say, to climate change 100 years from now, if discounted at 4.3% a year is seventeen times smaller than the present-value of that same consumption loss if the discount rate used is 1.4% a year. The moral is banal: If the time horizon is long, even small differences in consumption discount rates can mean large differences in the message cost-benefit analysis gives us. Cline (1992) and Stern (2006) recommended that the world spends substantial sums today to tame climate change, while Nordhaus (1994) recommended a gradualist investment policy. Their differences can be traced to the difference in their choice of δ. Nordhaus (2007) confirms that by using Stern's specifications for δ and η in the climate-change model he has developed over the past two decades.

4*.3.3 Commentary

We have seen that contrary to general belief, consumption discount rates are not a primary ethical concept; they are derived jointly from an overall conception of intergenerational well-being and the economic forecast. Discount rates can't be plucked from air, they have to be derived from such considerations as those that are formalized in equation (4*.8). We have noted that just as growing consumption provides a reason why discount rates in use in social cost-benefit analysis should be positive, declining consumption would be a reason why they might be negative. In an imperfect economy consumption discount rates are not equal to private rates of return on investment. They differ because of imperfect capital markets and corporate income taxes. If government policies are imperfect, consumption discount rates are not even equal to social rates of return on investment (Lind, 1982; Arrow et al., 1996). And as we have already noted, consumption discount rates could be negative even while private rates of return on investment are positive.[16]

Consumption discount rates are not necessarily constant over time. Suppose long-run economic forecasts indicate that growth in consumption is not sustainable, but rather, that its growth is expected to decline at a constant rate of 1 percent per year, from the current figure of

[16] This feature parallels a point familiar in analyses of the "tragedy of the commons", that if the damage to others arising from someone's use of a polluting commodity is large enough, the commodity's shadow price would be negative even though its market price is positive.

2 percent per year to zero. Assume that $\delta = 0$ and $\eta = 2$. In that case the consumption discount rate will decline over time at 1 percent per year, from a current-high 4 percent per year to zero. That means costs and benefits over the very near future, measured in consumption units, should be discounted at 4 percent per year, but those, say, 70 years hence should be discounted at 2 percent per year. And so forth. In economics, discounting future benefits and costs at a declining rate is known as hyperbolic discounting. Our analysis shows that if the forecast is that the rate of growth in consumption will decline over time, the social evaluator should apply hyperbolic discounting.

4*.4 Consumption Discount Rates under Uncertainty

Carbon concentration in the atmosphere is a very crude summary measure of the global climate system. There is enormous uncertainty about the changes in the spatial and temporal character of our climate that would result from further increases in the concentration level. Uncertainties are compounded when we try to fathom the economic consequences of those increases. But the basic implications of that uncertainty for consumption discount rates are clear: uncertainty is a reason for reducing those rates. We now build on the previous analysis to show why.

Imagine that the social evaluator is able to articulate uncertainties in the form of a probability distribution over all possible consumption paths. We assume she has an expectation of future consumptions, say $\{C(0),C(1),...,C(t),...\}$, but knows that the realised path will almost surely be different from her expectation. She regards intergenerational well-being at t under uncertainty to be the expected value of $V(t)$ in expression (4*.2). As she is peering into the deep future, our social evaluator also regards it a possibility that Humanity as we know it will cease to exist at some unknown future date. A simple way to reflect that uncertainty is to build it into δ. For concreteness, we imagine that the social evaluator continues to be persuaded by moral philosophers that it is unethical to discount future well-beings. So, δ in expression (4*.2) is the probability that Humanity will be extinct at t, conditional on it having survived until t-1. δ should now be called the "hazard rate".[17]

Now it is a deep truth in rational choice theory that if a social evaluator is averse to inequality in consumption among people, she ought to be averse toward uncertainty in future consumption as well. We have previously shown that if consumption is certain to increase under business as usual, inequality aversion (i.e., a strictly concave U) is a reason for discounting future consumption costs and benefits at a positive rate. If, however, the increase in consumption is not certain, an aversion to risk kicks in, meaning that a 50% chance that the increase will be less than expected should weigh heavier in the social evaluator's mind than a 50% chance that the increase

[17] The interpretation is due to Yaari (1965). Stern (2006) used Yaari's argument to justify setting $\delta = 0.1\%$ a year. Later we shall argue that the hazard rate should be an increasing function of time in the deep future.

will be greater. Put another way, because U is strictly concave, the downside risk would be given a higher weight than the corresponding good fortune; which implies that consumption discount rates ought to be smaller than what they would be if there were no uncertainty - other things being equal, of course. But this is the same as saying that the social evaluator would prescribe insurance against risks. In the context of global climate change, purchasing insurance amounts to additional reductions in carbon emissions, greater R&D expenditure toward clean energy technologies, methods for capturing carbon from the atmosphere, and so forth.

To develop the argument formally with the help of a concrete example, let $\tilde{C}(t)$ be the uncertain consumption rate at t (the "tilde" being the symbol indicating that consumption at t is uncertain when viewed for any previous date). Imagine the social evaluator is studying the economy at t. $C(t)$ is known at t, but $C(t+1)$ is not. Suppose that for all t and $C(t)$, $\log[\tilde{C}(t+1)/C(t)]$ has the same normal distribution.[18] As previously, let $g(\tilde{C}(t))$ be the (uncertain) rate of change in consumption between t and $(t+1)$; that is, $g(\tilde{C}(t)) = [\tilde{C}(t+1)/C(t)-1]$. Let $E[g(\tilde{C}(t))]$ denote the expected rate of growth in consumption, and $\text{var}[g(\tilde{C}(t))]$ the variance in the growth rate. By assumption, each is independent of t.

Working with models containing random variables is hard work. So, mathematicians have devised ways to convert decision problems under uncertainty into problems containing no uncertainty in such a way that there would be no error in solving the wrong problem. We call the converted problem the "certainty-equivalent problem". Let us see how that might work in the economics of climate change.

Suppose the social evaluator pretends that there is no uncertainty and that consumption will grow with certainty at the rate $E[g(\tilde{C}(t))]$. Is there a consumption discount rate she could use if she is to ensure that she will make no error in her recommendations despite the error in her assumption? In the finance literature, that discount rate is called the risk-free rate. In the economics literature it is called the certainty-equivalent rate. Let that rate be $\rho(t)$. It can be shown that, provided δ, $E[g(\tilde{C}(t))]$, and $\text{var}(g(\tilde{C}(t)))$ are small, the risk free rate is

$$\rho(t) = \rho = \delta + \eta E[g(\tilde{C}(t))] - \eta^2\text{var}(g(\tilde{C}(t)))/2.^{19} \qquad (4^*.11)$$

Compare equations (4*.10) and (4*.11). The latter contains an extra term, reflecting the effect of uncertainty on the consumption discount rate. As $E[g(\tilde{C}(t))]$ and $\text{var}(g(\tilde{C}(t)))$ are constants, the risk-free rate is a constant. Earlier we noted that the consumption discount rate ought to be negative at any date at which consumption is forecast to decline. Equation (4*.11) builds on that insight. The third term on the right hand side of the equation shows that an increase in uncertainty

[18] The uncertainty in the exact value of $C(t+1)$ would be due to uncertainty in the economy's productivity between t and $t+1$. We are not modelling that more basic uncertainty. For the latter, see Dasgupta (2008) and Arrow (2009).

[19] The equation is familiar in the theory of finance (Cochrane, 2005: p. 10). Notice that if $\text{var}(g(\tilde{C}(t))) = 0$, equation (4*.8) reduces to equation (4*.7).

reduces the (certainty equivalent) consumption discount rate, other things being equal. The term reflects a form of the "precautionary principle", that uncertainty is a reason for showing extra concern about future prospects.

4*.5 Large Uncertainties

Equation (4*.11) holds only if the uncertainties are small. What if they are large?

What one means by "large" of course depends on the other parameters, which are δ, η, and $E[g(C(t))]$. All those who have written on the economics of climate change have assumed that $\eta \geq 1$. Let us also do so. $U(C)$ is therefore unbounded below (Figures 4*.2 and 4*.3), which is to say that if C is small, $U(C)$ is a large negative number.

Imagine that the economic evaluation is being conducted at the present time. Consider a date t far into the distant future. Equation (4*.4) says that the social evaluator would regard the flow of social well-being at t to be $U(C(t))/(1+\delta)^t$. The denominator is a large positive number; but if C is small, $U(C(t))$ is a large negative number. It follows that for any t, no matter how large, there are Cs, sufficiently small, for which $U(C(t))/(1+\delta)^t$ is a sizeable negative number. It can be shown that if var$[g(C(t))]$ is large relative to the expected value of $g(C(t))$, and you choose a positive number as close to zero as you care, C will be less than that number sufficiently frequently to make the expected value of $U(C(t))/(1+\delta)^t$ a sizeable negative number no matter how large t happens to be. We conclude that if the uncertainty in $g(C(t)))$ is large, the expected value of expression (4*.4) is, loosely speaking, minus infinity, which is to say that every possible economic policy is infinitely awful.[20] As minus infinities are unrankable, consumption discount rates are undefinable. That in turn means the social evaluator is able to conduct neither policy analysis nor sustainability analysis, which were the starting motivations in Chapter 4. The combination of assumptions we have made is incoherent.

What are we to make of this? Well, not much. Such paradoxes of infinity are artifices, manufactured by bad assumptions.[21] Large uncertainties should certainly be considered. But the other assumptions on which we have built our account are, taken together, wholly unreasonable. Consider that the integrated assessment models that inform the Intergovernmental Panel on Climate Change entertain a finite number of consumption paths, called "scenarios", the consumption level in none of which is ever even close to zero. In those models the downside

---

[20] See Dasgupta (2008) and Arrow (2009).

[21] Weitzman (2007a,b) has shown that if $E[g(C(t))]$ is itself uncertain, the distribution of $C(t+1)/C(t)$ can plausibly have a "thick" lower tail, implying that over the infinite future, long runs of low realizations of $C(t)$ would not be improbable. In the text, we have been assuming that $E[g(C(t))]$ is known, and that it is distributed normally. The normal distribution, however, has a thin tail. So, the paradox of infinity that has been much discussed in the recent literature on climate change is not a feature unique to thick-tailed distributions, they arise even if the distributions are thin tailed, provided of course that the uncertainties are large in a suitable sense. On this, see Dasgupta (2008) and Arrow (2009).

risks associated with climate change are bounded. The way to introduce such a bound in our example here would be to truncate the assumed normal distributions of $\log[C(t+1)/C(t)]$ on the left. That would eliminate the paradox.

Another route for avoiding the paradox of infinity would be to abandon the assumption that $U(C)$ is unbounded below. Assume instead that no matter how greatly the economy were to be hit by bad luck, the loss in well-being people would suffer from would be bounded. The paradox of infinity would again disappear.

There are two further assumptions we have made that are surely artifacts: a horizon that is infinite and hazard rate (δ) for Humanity's extinction that is constant. We consider them in turn.

One way to ensure that the ethical framework we invoke doesn't have contradictions would be to to abandon the infinite time horizon. But the choice of a terminal date would at best be arbitrary. That is why economists have avoided working with finite time horizon models. Another way out would be to continue to postulate an infinite horizon, but formalise Humanity's extinction process in terms of a hazard rate that increases in an unbounded fashion over time at a sufficiently high rate. Modern cosmology advises to do that. Even assuming Humanity doesn't destroy itself, cosmologists tell us that the sun will not be able to sustain life on Earth for more than a few billion years more.

The paradoxes of infinity in the recent literature in sustainability science are artifices. They are a creation of a combination of simplifying assumptions economists have made over the years in their modelling of the possibilities of long run economic growth. Each assumption in that modelling has been known to be questionable, but because simplicity is attractive, each has been entertained. Taken together, though, they have now been shown to lead to incoherence. Since we know how that incoherence can be avoided, the thing to do is to abandon the worst of those assumptions.

References

Arrow, K.J. (2009), "A Note on Uncertainty and Discounting in Models of Economic Growth", Journal of Risk and Uncertainty, 38(2), 87-94.

Arrow, K.J. and P.S. Dasgupta (2009), "Conspicuous Consumption, Inconspicuous Leisure," Economic Journal, 119 (541), F497-F516.

Arrow, K.J., W.R. Cline, K.-G. Mäler, M. Munasinghe, R. Squitieri, and J.E. Stiglitz (1996), "Intertemporal Equity, Discounting, and Economic Efficiency", in J.P. Bruce, H. Lee and E.F. Haites, eds., Climate Change 1995: Economic and Social Dimensions of Climate Change, Contribution of Working Group III to the Second Assessment Report of the Intergovernmental Panel on Climate Change (Cambridge: Cambridge University Press).

Arrow, K.J., P. Dasgupta, L. Goulder, G. Daily, P.R. Ehrlich, G.M. Heal, S. Levin, K.-G. Mäler, S. Schneider, D.A. Starrett, and B. Walker (2004), "Are We Consuming Too Much?", Journal of Economic Perspectives, 18(1), 147-172.

Arrow, K.J., P. Dasgupta, and K.-G. Mäler (2003), "The Genuine Savings Criterion and the Value of Population", Economic Theory, 21(2), 217-225.

Arrow, K.J. and M.Kurz (1970), Public Investment, the Rate of Return and Optimal Fiscal Policy (Baltimore: Johns Hopkins University Press).

Broome, J. (1992), Counting the Cost of Global Warming (London: White Horse Press).

Cline, W.R. (1992), The Economics of Global Warming (Washington, D.C.: Institute for International Economics).

Cline, W.R. (1993), "Give Greenhouse Abatement a Fair Chance", Finance and Development, March, 3-5.

Cochrane, J.H. (2005), Asset Pricing, 2nd Ed. (Princeton, NJ: Princeton University Press).

Dasgupta, P. (1969), "On the Concept of Optimum Population", Review of Economic Studies, 36(3), 295-318.

Dasgupta, P. (2008), "Discounting Climate Change", Journal of Risk and Uncertainty, 37(4), 141-169.

Fishkin, J. and R. Goodin, eds. (2010), Population and Political Theory - Philosophy, Politics and Society 8 (Malden, MA: Wiley-Blackwell).

Gottlieb, M. (1945), "The Theory of Optimum Population for a Closed Economy", Journal of Political Economy, 53(2), 289-316.

Koopmans, T.C. (1960), "Stationary Ordinal Utility and Impatience", Econometrica, 28(2), 287-309.

Koopmans, T.C. (1965), "On the Concept of Optimal Economic Growth", Pontificiae Academiae Scientiarum Scripta Varia, 28. Reprinted in The Econometric Approach to Development Planning, 1966 (Amsterdam: North Holland).

Koopmans, T.C. (1972), "Representation of Preference Orderings over Time", in C.Bart McGuire and Roy Radner (eds.), Decision and Organization (Amsterdam: North Holland).

Lind, R.C. ed. (1982), Discounting for Time and Risk in Energy Policy (Baltimore: Johns Hopkins University Press).

Meade, J.E. (1955), Trade and Welfare (Oxford: Oxford University Press).

Mirrlees, J.A. (1967), "Optimum Growth when Technology is Changing", Review of Economic Studies, 34(1), 95-124.

Nordhaus, W.D. (1994), Managing the Global Commons: The Economics of Climate Change (Cambridge, MA: MIT Press).

Nordhaus, W.D. (2007), "The Stern Review on the Economics of Climate Change", Journal of Economic Literature, 45(3), 686-702.

Parfit, D. (1984), Reasons and Persons (Oxford: Oxford University Press).

Parfit, D. (1990), "Overpopulation and the Quality of Life", in J. Glover ed. Utilitarianism and its Critics (London: Macmillan).

Ramsey, F.P. (1928), "A Mathematical Theory of Saving", Economic Journal, 38(4), 543-49.

Ryder, H.E. and G.M. Heal (1973), "Optimal growth with intertemporally dependent preferences", Review of Economic Studies, 40(1), 1-31.

Stern, N.H. (2006), The Stern Review of the Economics of Climate Change (Cambridge: Cambridge University Press).

Weitzman, M.L. (2007a), "The Stern Review of the Economics of Climate Change," Journal of Economic Literature 45(3), 703-724.

Weitzman, M.L. (2007b), "The Role of Uncertainty in the Economics of Catastrophic Climate Change," Discussion Paper, Department of Economics, Harvard University.

Yaari, M. (1965), "Uncertain Lifetime, Life Insurance, and the Theory of Consumer", Review of Economic Studies, 32(2), 137-158.

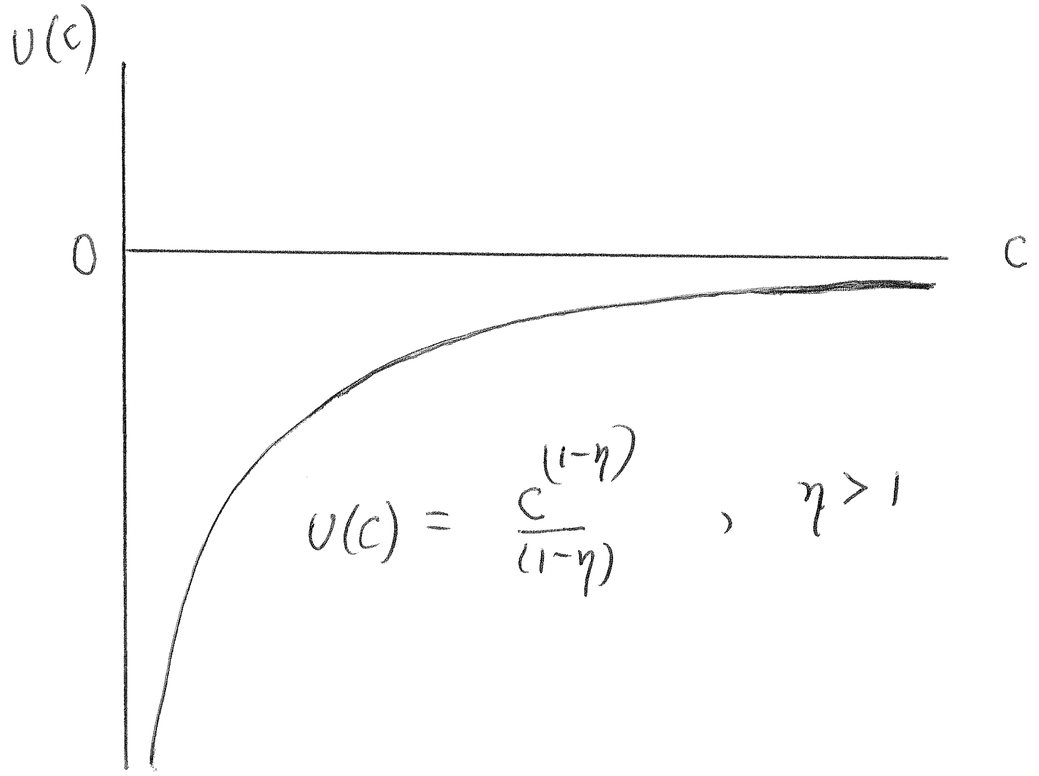

Figure 4*.1

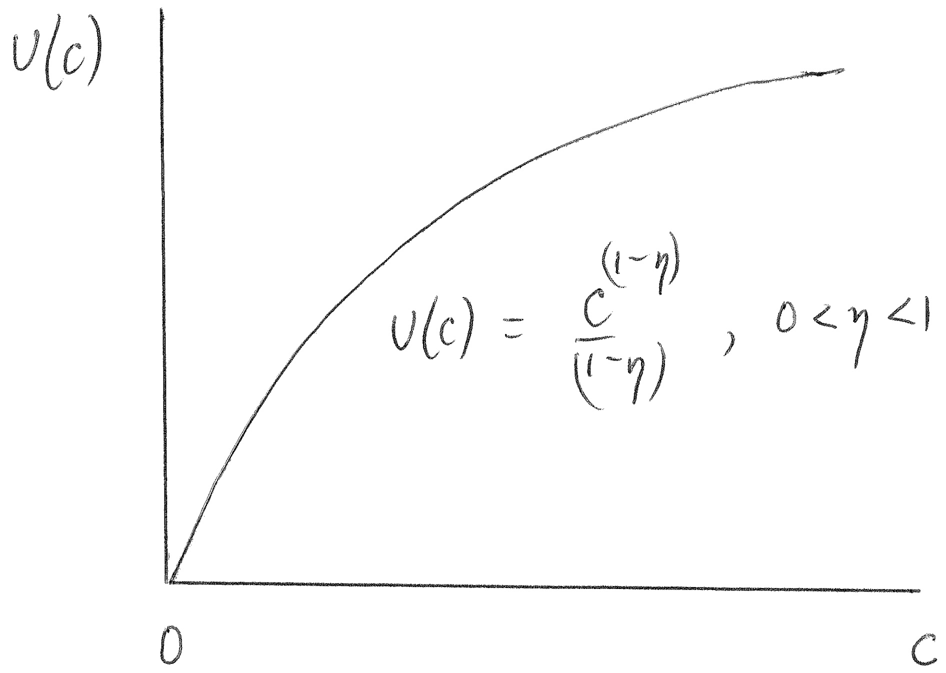

figure 4*10.2

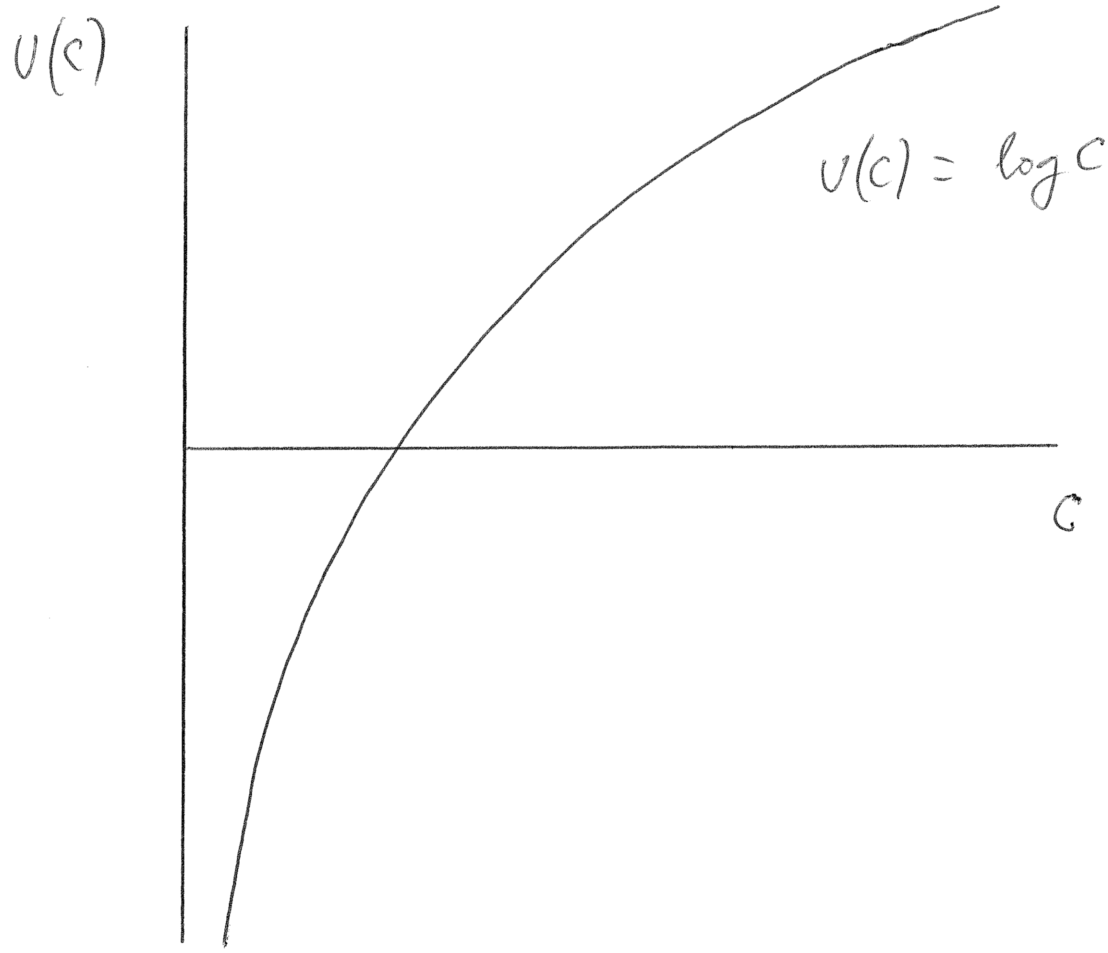

Figure 4*.3