

NBER WORKING PAPER SERIES

RISK-ADJUSTED GAMMA DISCOUNTING

Martin L. Weitzman

Working Paper 15588

<http://www.nber.org/papers/w15588>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

December 2009

The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2009 by Martin L. Weitzman. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Risk-Adjusted Gamma Discounting
Martin L. Weitzman
NBER Working Paper No. 15588
December 2009
JEL No. Q54

ABSTRACT

It is widely recognized that the economics of distant-future events, like climate change, is critically dependent upon the choice of a discount rate. Unfortunately, it is unclear how to discount distant-future events when the future discount rate itself is unknown. In previous work, an analytically-tractable approach called "gamma discounting" was proposed, which gave a declining discount rate schedule as a simple closed-form function of time. This paper extends the previous gamma approach by using a Ramsey optimal growth model, combined with uncertainty about future productivity, in order to "risk adjust" all probabilities by marginal utility weights. Some basic numerical examples are given, which suggest that the overall effect of risk-adjusted gamma discounting on lowering distant-future discount rates may be significant. The driving force is a "fear factor" from risk aversion to permanent productivity shocks representing catastrophic future states of the world.

Martin L. Weitzman
Department of Economics
Harvard University
Littauer 313
Cambridge, MA 02138
and NBER
mweitzman@harvard.edu

Risk-Adjusted Gamma Discounting

Martin L. Weitzman*

December 8, 2009. Preliminary. Comments much appreciated.

Abstract

It is widely recognized that the economics of distant-future events, like climate change, is critically dependent upon the choice of a discount rate. Unfortunately, it is unclear how to discount distant-future events when the future discount rate itself is unknown. In previous work, an analytically-tractable approach called “gamma discounting” was proposed, which gave a declining discount rate schedule as a simple closed-form function of time. This paper extends the previous gamma approach by using a Ramsey optimal growth model, combined with uncertainty about future productivity, in order to “risk adjust” all probabilities by marginal utility weights. Some basic numerical examples are given, which suggest that the overall effect of risk-adjusted gamma discounting on lowering distant-future discount rates may be significant. The driving force is a “fear factor” from risk aversion to permanent productivity shocks representing catastrophic future states of the world.

1 Introduction: Discounting Climate Change

The concept of discounting is central to economics, since it allows effects occurring at different future times to be compared by converting each future dollar into the common currency of equivalent present dollars. Because of this centrality, the choice of an appropriate discount rate is one of the most critical issues in economics. It represents an especially acute dilemma for projects involving long time horizons, because in such situations the results of cost-benefit analysis (CBA) can depend enormously on the choice of a discount rate.

The problem of an unsure discount rate has long bedeviled CBA, but it has acquired renewed relevance lately because economists are increasingly being asked to analyze environmental projects or activities whose effects will be felt very far out into the future.

*Department of Economics, Harvard University (mweitzman@harvard.edu). Without necessarily tying them to the contents of this paper, I am extremely grateful for critical comments by Christian Gollier.

Examples include nuclear waste management, loss of biodiversity, groundwater alterations, minerals depletion, and many others. The most prominent example by far is the economics of climate change, which will be used throughout this paper as the prototype application.

The effects of global warming and climate change will be spread out over what might be called the “distant future” – centuries and even millennia from now. The logic of compounding a constant positive interest rate forces us to say that what one might conceptualize as monumental – even earth-shaking – events do not much matter when they occur in the distant future. Perhaps yet more disconcerting, when exponential discounting is extended over very long time periods there is a truly extraordinary dependence of CBA on the choice of a discount rate. Seemingly insignificant differences in discount rates can make an enormous difference in the present discounted value of distant-future payoffs. In many long-run situations, including climate change, almost any answer to a CBA question can be defended by one particular choice or another of a discount rate.

Answers to questions thrown up by climate-change CBA therefore hinge critically on the core issue of how to discount the distant future. There is a high degree of uncertainty about what should be taken as the appropriate real rate of return on capital in the long run, accompanied by much controversy about its implications for long-run discounting. For specificity, the investigation of this paper is focused sharply on CBA of small investments that might incrementally impact distant-future events. However, there are broader welfare ramifications as well. Choosing the “right” distant-future discount rate is emblematic of a whole series of thorny issues concerning how to conceptualize and how to evaluate actions spanning across many generations.

I think it is important to begin by recognizing that there is no deep reason of principle that allows us to extrapolate past rates of return on capital into the distant future. The industrial revolution itself began some two centuries ago, and only slowly thereafter permeated throughout most of the globe. The seeming trendlessness of some past rates of return is a purely empirical reduced-form observation, which is not based on any underlying theory that would confidently allow projecting the past far into the future. There are a great many fundamental non-extrapolatable factors, just one example of which is the unknown future rate of technological progress. Even leaving aside the question of how to project future interest rates, additional issues for climate change involve which interest rate to choose out of a multitude of different rates of return that exist in the real world. Furthermore, there is a strong normative element having to do with what is the “right” rate, which includes an ethical dimension to discounting climate change across many future generations that is difficult to evaluate and incorporate into standard CBA.

A large number of additional examples of economic and non-economic features could be

given that would be highly relevant to determining the distant-future discount rate, including the effects of climate change itself on productivity via what is sometimes called a “damages function.” (Having the damages to output be an uncertain function of temperature is loosely analogous to having an uncertain real rate of return on capital.) The fundamental point is that there is enormous uncertainty and controversy about choosing the appropriate real rate of return on capital for discounting distant-future events, like climate change. Moreover, the great uncertainty about how to discount the distant future is not just an academic curiosity, but it has critically important implications for climate change policy. This disturbing ambiguity has given rise to a great deal of controversy and a variety of proposed solutions. My purpose here is to focus sharply on clarifying the long-run discounting issue by using a super-simple super-crisp formulation that abstracts away from all other elements of CBA, many of which may also be important in practice.

Climate change is so complicated, and it involves so many sides of so many different disciplines and viewpoints, that no analytically-tractable model or paper can aspire to illuminate more than but a facet of this problem. Because the climate-change problem is so complex, there is frequent reliance on numerical computer simulations. These can be indispensable, but sometimes they do not provide a simple intuition for the processes they are modeling. In this paper, I go to the opposite extreme by focusing exclusively on analytically tractable solutions to the pure discounting problem – in order to get to the core essence of this particular issue as directly as possible. There is no good substitute for seeing clearly before one’s eyes the basic structure of a model laid bare. What I am presenting here is a kind of fable or parable about how to approach discounting the distant future when discount rates are uncertain, which is several steps removed from practical implementation. This paper is more about conceptualizing the distant-future discounting problem and conveying some rough sense of the magnitudes involved. It is less about giving decisive numerical values for actual practical discount rate schedules. The beauty of this approach is that the formulas I will derive are so simple and transparent that readers can easily plug in different parameter values, or attach the model to other frameworks – and draw their own conclusions.

The core mathematical foundation, which undergirds the entire paper, is easy enough to state formally.¹ Suppose that discount rate $r_i > 0$ will occur with “probability-like weight” $w_i > 0$, where $\sum w_i = 1$. Mechanically define the average or expected discount *factor* at time t to be

$$\Phi(t) = \sum w_i \exp(-r_i t), \tag{1}$$

whose corresponding discount *rate* $R(t)$ satisfies $\exp(-R(t)t) = \Phi(t)$, which can be rewritten

¹See Weitzman (1998) for further details.

as

$$R(t) = \frac{-\ln(\sum w_i \exp(-r_i t))}{t}. \quad (2)$$

It is not difficult to show that

$$\Phi(0) = 1, \quad \dot{\Phi}(t) < 0, \quad \Phi(\infty) = 0, \quad (3)$$

and that

$$R(0) = \sum w_i r_i, \quad \dot{R}(t) < 0, \quad R(\infty) = \min\{r_i\}. \quad (4)$$

The discount rates in (4) decline over time starting from their average value and going down over time to approach their lowest possible value. Over time, the impact of the higher discount rates in formulas (1) and (2) diminishes because the higher rates effectively discount themselves exponentially out of existence, leaving the field to the lower discount rates (and, eventually, to the lowest).

The formulas (1), (2), (3), (4) are the building blocks for this paper. To them I attach a simple Ramsey optimal growth model, whose future productivity of capital is unknown. I then use this apparatus for analyzing how to discount the distant future for an investment decision that must be made now, when the future productivity of capital is as yet an uncertain random variable having some probability density function (PDF). The future productivity of capital comes in the form of a permanent productivity shock that is drawn from its PDF immediately *after* the investment decision is made. Then the deterministic economy adjusts instantaneously to the then-known realization of the real rate of return on capital. I do not defend this model for its realism and immediate applicability to such long-term issues as CBA of climate change. Rather, I defend this particularly simple example of decision making under uncertainty for its ability to isolate cleanly and clarify sharply a set of controversial issues that have bedeviled the discounting of distant-future events like climate change.

I derive the relevant future discount rate schedule as a simple closed-form function of time. One of the key simplifying assumptions leading to this analytically tractable formula is that uncertain future productivity is a random variable whose PDF is a gamma distribution. In previous work, the gamma-distributed probabilities of future discount rates were simply taken as given.² The approach of this paper probes perhaps a little deeper by “risk adjusting” these raw gamma probabilities with the marginal utility weights that are endogenously determined from the solution to an optimal growth problem. It then emerges

²Weitzman (2001).

that even moderate risk aversion can enhance significantly the effect of gamma discounting by increasing the risk-adjusted probabilities of low-productivity states, thereby putting more probability weight on scenarios with low discount rates and low endogenously-chosen consumption levels.

The paper works with a constant relative risk aversion (CRRA) utility function whose coefficient of relative risk aversion will be denoted η . The standard normative interpretation of the famous deterministic Ramsey formula is that an exogenously-given growth rate g determines the endogenous future discount rate r . Under this usual non-optimizing interpretation, higher values of η cause *higher* discount rates. This kind of partial-equilibrium argument has underpinned much of the thinking and the debate concerning what discount rates should be used in climate-change CBA. I believe that this standard partial-equilibrium interpretation may not be the best way of looking at the normative discounting problem. In the normative general-equilibrium model of this paper, the direction of causality is reversed. I think that it may make more economic sense to have the primitive driver be the exogenously-given future productivity of capital r . Then any given value of r subsequently determines the endogenous future growth rate g and endogenous consumption level C as the solution to a Ramsey optimal growth problem (given that value of r). The paper shows that when future productivity r is uncertain, then higher values of η are associated with *lower* future discount rates, thereby reversing the conventional wisdom. A base-case numerical example will be given, using what I believe are conventional parameter values. This example – which is based on numbers that conform more or less with real-world observations – appears to indicate that the risk-aversion effects of uncertain future productivity on lowering distant-future discount rates might be quite powerful. The driving force is a “fear factor” associated with the possibility of low-probability but catastrophically-high permanent damages to future productivity.

2 Ramsey Optimal Deterministic Growth with Linear Production and CRRA Utility

The solution to the deterministic Ramsey optimal growth problem with a linear production function and a constant relative risk aversion (CRRA) utility function is well known. Following what Ramsey proposed, I take the rate of pure time preference (or the so-called “utility discount rate”) in this paper to be zero. As Ramsey famously put the issue, “it is assumed that we do not discount later enjoyments in comparison with earlier ones, a practice which is ethically indefensible and arises merely from the weakness of the imagination.” Many

other eminent economists concur with this Ramsey interpretation of intergenerational equity.³ Taken together, quotations from these “eminent economists” sound to me much more like a normative judgement about intergenerational ethics than a description of short-run individual behavior. I think that the Ramsey case of zero discounting of future utilities is the appropriate abstraction for a normative analysis of climate change. Ethically or morally, the Ramsey abstraction treats the utility of different generations equally, while taking full account of the fact that economic growth will make future generations richer and less needy than the present generation.

The notion of capital at time t , denoted $K(t)$, is intended here to be all-inclusive, incorporating human capital, knowledge capital, R&D capital, skill capital, health capital, ecosystem capital, and so forth. The economy produces net output $Y(t)$ at time t according to the linear technology $Y(t) = rK(t)$, where, for the deterministic model of this section, r is a given positive constant. (In the literature, such kind of linear production function is known as a so-called “AK” technology, but for the purposes of this paper I will use the symbol r in place of A to remind us that the linear coefficient multiplying K is always the relevant real rate of return on capital, irrespective of the utility or welfare side.) Capital here is essentially a vehicle for transferring consumption between the present and the future at interest rate r . Later in the paper r will be considered a random variable representing the unknown future productivity of capital, but for now r is treated as a given known constant.

Net output $Y(t)$ can be divided between consumption $C(t)$ and net investment $\dot{K}(t) = dK/dt$ according to the schedule

$$C(t) + \dot{K}(t) = rK(t). \quad (5)$$

The utility of consumption is taken to be of the isoelastic or CRRA form

$$U(C) = \frac{C^{1-\eta}}{1-\eta}, \quad (6)$$

where $\eta > 1$ is the coefficient of relative risk aversion or, equivalently, the elasticity of marginal utility. Ramsey’s “bliss level” for utility function (6), with $\eta > 1$, is $B = 0$. The marginal utility of consumption is

$$U'(C) = C^{-\eta}. \quad (7)$$

³Pigou: [pure time preference] “implies ... our telescopic faculty is defective.” Harrod: “pure time preference [is] a polite expression for rapacity and the conquest of reason by passion.” Koopmans: “[I have] an ethical preference for neutrality as between the welfare of different generations.” Solow: “in solemn conclave assembled, so to speak, we ought to act as if the social rate of pure time preference were zero.” (All quotes are taken from Arrow (1999).) I think it should be clear that the above citations refer to a normative or prescriptive, rather than a positive or descriptive, view of the world.

The deterministic Ramsey optimal growth problem is to choose as control variable a consumption trajectory $\{C(t)\}$ to maximize

$$V(\{C(t)\}) = \int_0^\infty \frac{C(t)^{1-\eta}}{1-\eta} dt, \quad (8)$$

subject to the production constraint (5) and the initial condition that $K(0)$ is given. Let the maximized value of V in (8) be denoted V_r^* .

It is well known⁴ that the solution to the above Ramsey optimal growth problem has consumption, and also investment and capital, all growing at the same constant steady-state rate g satisfying the fundamental Ramsey equation

$$r = \eta g. \quad (9)$$

It is readily shown that the solution corresponding to equation (9) obeys the simple linear optimal-consumption rule

$$C_r^*(t) = \left[\frac{\eta - 1}{\eta} \right] r K(t). \quad (10)$$

Due to the linear production possibilities frontier (5), reducing $C_r^*(0)$ marginally by one small unit would result in $\exp(rt)$ extra units available for increased consumption at future time t (without altering the rest of the optimal trajectory $\{C_r^*(t)\}$). From this simple perturbation argument, the optimal trajectory must satisfy the first-order marginal condition

$$U'(C_r^*(0)) = \exp(rt) U'(C_r^*(t)) \quad (11)$$

for all t and for all r . With the CRRA utility function (7), equation (11) translates into

$$(C_r^*(t))^{-\eta} = (C_r^*(0))^{-\eta} \exp(-rt). \quad (12)$$

As was previously mentioned, the standard normative interpretation of the Ramsey formula (9) is that the future growth rate g is more or less exogenously given and, for postulated η , it determines the appropriate r for CBA via (9). In this paper the causality is reversed. For any exogenously given productivity r , CRRA coefficient η implies a corresponding endogenously-chosen value of the growth rate $g = r/\eta$ and initial consumption level $C_r^*(0)$. As will be shown later, under uncertainty it turns out that the direction of causality between r and g matters.

Real rates of return on investments vary from $r \approx 1\%$ for relatively riskfree U.S. T-Bills

⁴See, e.g., Weitzman (2003).

up to $r \approx 10 + \%$ for public health or education in underdeveloped regions of the world. The base case for this paper is $r=6\%$ per year. This is approximately the average after-tax real rate of return on capital throughout much of the world. For a CBA of the economics of climate change, $r=6\%$ would be considered a very high discount rate. I personally would be inclined toward a much lower climate-change discount rate than 6% per annum, but the ultimate goal of this paper, to which this deterministic section is leading, will be to show that *under uncertainty*, even with *expected* discount rates as high as 6% , the “effective” discount rate, which “ought” to be used, can be much lower than 6% . Back to the deterministic case of this section, with projected $r=6\%$, the relevant discount factor $\Phi(t)$, which would be used to convert money values from future year t into money values today (at time zero), is given by the formula

$$\Phi(t) = \exp(-.06 t). \tag{13}$$

Some distant future values of $\Phi(t)$ from formula (13) are given in the following table.

future year $t =$	50	100	150	200	300	500	1,000
discount factor $\Phi(t) =$.05	.002	.0001	6×10^{-6}	2×10^{-8}	9×10^{-14}	9×10^{-27}

Table 1: Discount Factor $\Phi(t)=\exp(-.06 t)$ for Future Year t

It should be emphasized strongly that the numbers in Table 1 would be construed by most people as severely biasing CBA toward minimizing into near-nothingness the present discounted value of distant-future events, like climate change. This kind of exponential discounting, perhaps more than anything else, makes scientists and the general public suspicious of the standard economist’s CBA of climate change, since it trivializes even truly enormous distant-future impacts. To be honest, I think that among economists there are but few who do not feel uneasy about evaluating distant-future climate-change impacts this way. As will be shown later, the introduction of uncertainty in the form of future r being a random variable from a mean-preserving spread (so that the measure of central tendency remains $E[r]=6\%$) can radically alter the discounting picture by causing $\Phi(t)$ to be significantly higher than the values given in Table 1.

To set forth a deterministic numerical example here, which will later be extended to a situation of uncertainty, as just mentioned my base-case point-estimate real rate of return on capital is $r=6\%$ per year. My base-case CRRA coefficient is $\eta=3$. My base-case point-estimate future growth rate of per capita consumption is $g=2\%$ per year.⁵ Notice that these

⁵These values are close to those that were proposed by Dasgupta (2008), and were considered fully acceptable by Nordhaus (2008, pp. 61 and 187).

base-case values are consistent with the fundamental Ramsey equation (9). Were η to be changed substantially, then r and g would not mesh quite so nicely with past reality. If $\eta=2$ and $r=6\%$, then (9) implies $g=3\%$ – probably too high. If $\eta=4$ and $g=2\%$, then (9) implies $r=8\%$, also probably too high. So I think it is fair to say that this proposed “package” of base-case point-estimate values ($r=6\%$, $\eta=3$, $g=2\%$) looks more or less realistic and is internally consistent. The implied net savings rate is $\frac{1}{3}$, which I would consider to be very roughly synchronous with a concept of generalized all-inclusive capital, whose investments include not only traditional forms of savings, but also expenditures on education, training, R&D, health, and so forth.

When the rate of pure time preference (or “utility discount rate”) is $\rho > 0$, then the deterministic Ramsey formula (9) generalizes to $r = \rho + \eta g$. It is then often argued that combinations of ρ and η calibrated to the same values of r and g are more or less equally legitimate. According to this logic, the package ($\rho=0$, $\eta=3$, $r=6\%$, $g=2\%$) is not essentially different from the package ($\rho=2\%$, $\eta=2$, $r=6\%$, $g=2\%$). However, this paper shows that, under uncertain permanent productivity shocks, the Ramsey-normative case $\rho=0$ has a special status because it does not allow the negative consequences of low- r states to be pushed forward onto distant future generation, as would happen, say, with $\rho=2\%$. Instead, zero pure time preference forces current generations to look more favorably upon CBA opportunities that might transfer current costs into distant-future benefits. In the uncertainty formulation of this paper, the normative choice $\rho=0$ can have substantive consequences.

Henceforth, the paper considers only the “ethical” case of a zero rate of pure time preference $\rho=0$. I now introduce exogenous uncertainty in r and, by (9), (10), endogenous uncertainty in g and in C (for given η).

3 Uncertain Future Rates of Return on Capital

To state loosely the issue at hand in this super-simple expository model⁶, a decision must be taken now, just before time zero (call it time 0^-), about whether or not to make an investment that costs δ now but will yield a benefit payoff ϵ at future time t . Right now, at time 0^- , it is unknown what will be the real rate of return on capital in the economy because the future productivity of capital r is uncertain. A decision must be made now (at time 0^- , *just before* the “true” state of the world r is revealed at time $t = 0^+$) about whether or not to invest in the proposed project. To pose the problem sharply, it is assumed that immediately *after* the investment decision is made, at time 0^+ , the true rate of return r becomes known

⁶The model of this section is a particular application of the more general framework of Gollier and Weitzman (2009), which itself was based on earlier work by Gollier (2009).

and the productivity of capital will thenceforth permanently be r , from time $t = 0^+$ to time $t = \infty$. The idea that productivity shocks are permanent (and materialize overnight) is very extreme, but it seems to me like an allowable initial abstraction for buying simplicity in analyzing the appropriate distant-future discount rate. In the case of climate change a recent study⁷ found that a significant amount of the CO₂ above pre-industrial-revolution levels remains in the atmosphere for centuries, and even millennia – long after *all* CO₂ emissions cease entirely. This might be construed as partial justification for damages to productivity acting like permanent shocks in an abstract super-simple model. In any event, and most importantly, the simplistic abstract model of this paper will give a structural framework with some useful insights that can be used as a point of departure for further refinements.

It is critical in this model to understand the exact timing sequence concerning what information is available at what time, and when decisions are made. To restate the setup more formally, at time 0^- the inherited capital stock is given as $K(0^-)$ and the permanent future value of r is a random variable whose realization is unknown. At time 0^+ , the productivity of capital will be known precisely to be some $r > 0$, which has been drawn from a given PDF $f(r)$. Suppose that an investment opportunity suddenly arises at time $t = 0^-$, *just before* the “true” value of r is drawn from the PDF $f(r)$ and revealed at time $t = 0^+$. This take-it-or-leave-it investment opportunity expends marginal cost δ now in order to yield a marginal benefit ϵ at future time t . Of course the representative agent wishes that the investment decision could be made with the precise information available at time $t = 0^+$, *just after* the “true” state of the world is revealed and the relevant future productivity of capital (from time $t = 0^+$ to time $t = \infty$) is known with certainty. But the abstract essence of the problem of doing CBA with an uncertain future discount rate is that the investment decision must be made at a time when the future productivity of capital is unknown.

If the investment opportunity is accepted, it means that the inherited initial capital stock $K(0^+)$ will be diminished by δ . The corresponding loss of utility from lowering the initial capital stock $K(0^+)$ by δ is equivalent to lowering $C(0^+)$ by δ (this well-known equivalence follows from the envelope condition). Therefore, the *expected loss* of welfare from investing δ now, probability averaged over all r , is

$$\delta \int_0^{\infty} U'(C_r^*(0)) f(r) dr \left[= \delta \int_0^{\infty} (C_r^*(0))^{-\eta} f(r) dr \right]. \quad (14)$$

Using the same logic, the *expected gain* of overall welfare from having ϵ extra units of

⁷Solomon et al (2009).

consumption available at time t , probability averaged over all r , is

$$\epsilon \int_0^{\infty} U'(C_r^*(t)) f(r) dr \left[= \epsilon \int_0^{\infty} (C_r^*(t))^{-\eta} f(r) dr \right]. \quad (15)$$

Therefore, the proposed investment will increase overall expected welfare if and only if

$$\epsilon \int_0^{\infty} (C_r^*(t))^{-\eta} f(r) dr > \delta \int_0^{\infty} (C_r^*(0))^{-\eta} f(r) dr. \quad (16)$$

Combining (12) with (16), the investment should be made if and only if

$$\epsilon \int_0^{\infty} (C_r^*(0))^{-\eta} \exp(-rt) f(r) dr > \delta \int_0^{\infty} (C_r^*(0))^{-\eta} f(r) dr. \quad (17)$$

Define the risk-adjusted discount factor $\Phi(t)$ to be

$$\Phi(t) \equiv \frac{\int_0^{\infty} (C_r^*(0))^{-\eta} \exp(-rt) f(r) dr}{\int_0^{\infty} (C_r^*(0))^{-\eta} f(r) dr}. \quad (18)$$

Combining (18) with (17), the proposed investment (incurring a marginal cost δ at time 0^- , in order to yield a marginal benefit ϵ at future time t) should be made if and only if

$$\epsilon \Phi(t) > \delta. \quad (19)$$

Condition (19) establishes in this setup that $\Phi(t)$ defined by (18) is the relevant CBA discount *factor* for time t . The corresponding discount *rate* $R(t)$ satisfies $\Phi(t) = \exp(-R(t)t)$, which can be rewritten as

$$R(t) = \frac{-\ln(\Phi(t))}{t}. \quad (20)$$

Equations (18) and (20) must obey the corresponding continuous versions of (3) and (4). What now remains is to assume some reasonable functional form of $f(r)$ that yields a simple closed-form expression for $R(t)$ and $\Phi(t)$, and to try out some numerical experiments.

4 Risk-Adjusted Gamma Discounting

In previous work I used a gamma distribution of discount rates⁸, which, for the purpose at hand, seemed like an ideal compromise among flexibility, generality, transparency, and analytical tractability. This gamma PDF was originally applied directly to the distribution of future discount rates. A possible justification for this direct approach is a variant of the so-called “Lucas fruit-tree model.” In this story, applied to the present setup, both the initial value of consumption $C(0)$ and its growth rate g are exogenously given by the process

$$C(t) = C(0)e^{gt}, \tag{21}$$

where g is the realization of a random variable representing a permanent growth-rate shock. Suppose that the growth rate g has a gamma distribution with mean $E[g]$ and variance $V[g]$. Then from the Ramsey equation (9), the discount rate r also has a gamma PDF, but with mean $E[r] = \eta E[g]$ and variance $V[r] = \eta^2 V[g]$. This is a rigorous story that could be used to justify using directly an as-if-given gamma distribution of discount rates with given mean $E[r]$ and given variance $V[r]$.

However, this partial-equilibrium “Lucas fruit-tree model” is not really in the normative spirit of the Ramsey general-equilibrium model, because there is no place for optimizing behavior. Instead, once the uncertainty is resolved, the consumption growth trajectory (21) is exogenous – with given g and given $C(0)$ – rather than being the solution of an optimal growth problem à la Ramsey. A drawback of this fruit-tree exogenous-growth story is that it appears to reverse the appropriate causality. I think that it may make more sense to consider future productivity r as the given primitive (in the form of a random variable), whose realization then induces *both* g and $C_r^*(0)$ in equation (21). The key distinction is that a permanent productivity shock to r will cause $C_r^*(0)$ (as well as g) to be endogenously determined.⁹ When realized productivity r is low, then, by (10), $C_r^*(0)$ will also be *chosen* to be low, in proportion to r . Low- r states force current generations to repress their consumption in favor of higher savings to compensate future generations for loss of output. The normative model of this paper takes full account of such an endogenous effect on the chosen level of consumption $C_r^*(0)$. I thereby extend the previous closed-form gamma PDF approach to include risk-adjusted marginal-utility-weighted probability densities from a Ramsey optimal growth solution, as they appear in (18).

⁸Weitzman (2002).

⁹As was emphasized early on by Gollier (2009), and elaborated further in Gollier and Weitzman (2009), in this kind of endogenous-growth setup it is incorrect to use the original PDF of r without a risk adjustment for the different marginal utilities associated with different values of r .

The gamma PDF for the random variable r is of the form

$$f(r) = \frac{\beta^\alpha}{\Gamma(\alpha)} r^{\alpha-1} e^{-\beta r}, \quad (22)$$

with positive parameters α and β . The mean and variance of the PDF (22) are

$$\mu = \frac{\alpha}{\beta}, \quad \sigma^2 = \frac{\alpha}{\beta^2}. \quad (23)$$

Now plug (10) (for $t = 0$) and (22) into (18). After canceling from the numerator and denominator all identical terms, (18) becomes

$$\Phi(t) = \frac{\int_0^\infty r^{\alpha-\eta-1} e^{-(\beta+t)r} dr}{\int_0^\infty r^{\alpha-\eta-1} e^{-\beta r} dr}. \quad (24)$$

(Equation (24) holds for $\eta < \alpha$. If $\eta \geq \alpha$, then a limiting argument for the lower part of the ratio of indefinite integrals as $r \rightarrow 0$ shows that $\Phi(t) = 1$ for all t .)

Now it is a fact of calculus that

$$\int_0^\infty x^{a-1} e^{-bx} dx = \frac{\Gamma(a)}{b^a}. \quad (25)$$

In (25), substitute the values $a = \alpha - \eta$ (assumed positive) and $b = \beta + t$, which makes (25) equal to the numerator of (24). Likewise substitute the values $a = \alpha - \eta$ and $b = \beta$ into (25), which makes (25) equal to the denominator of (24). Next make use of the right hand side of (25) to evaluate the numerator of (24) (for values $a = \alpha - \eta$ and $b = \beta + t$) and the denominator of (24) (for values $a = \alpha - \eta$ and $b = \beta$). After simplification, equation (24) then gets transformed into

$$\Phi(t) = \left(\frac{\beta}{\beta + t} \right)^{\alpha-\eta}. \quad (26)$$

Because μ and σ have economically intuitive meanings, while α and β have no economic significance per se, it will be more appealing in this paper to work with the inverse of (23), namely

$$\alpha = \left(\frac{\mu}{\sigma} \right)^2, \quad \beta = \frac{\mu}{\sigma^2}. \quad (27)$$

After substituting from (27), the relevant discount *factor* (26) can be rewritten as

$$\Phi(t) = \left(\frac{1}{1 + \frac{\sigma^2 t}{\mu}} \right)^{\left(\frac{\mu}{\sigma}\right)^2 - \eta}. \quad (28)$$

From (20), the relevant discount *rate* corresponding to (28) is

$$R(t) = \frac{\left[\left(\frac{\mu}{\sigma}\right)^2 - \eta \right] \ln \left(1 + \frac{\sigma^2 t}{\mu} \right)}{t}. \quad (29)$$

Equations (29) and (28) give simple closed-form expressions for discounting a future whose discount rates are themselves uncertain. Note from formulas (28) and (29) that $\Phi(t) \rightarrow \exp(-\mu t)$ and $R(t) \rightarrow \mu$ in the limit as $\sigma \rightarrow 0$. Note also that $\Phi(t) \rightarrow 0$ and $R(t) \rightarrow 0$ in the limit as $t \rightarrow \infty$. Finally, and most importantly, notice the sensitivity of both formulas (28) and (29) to η . Higher risk aversion causes higher discount factors and lower discount rates. This effect can be quite pronounced, as a numerical example in the next section will demonstrate. Actually, as previously indicated, a limiting argument applied directly to (24) shows that sufficiently high risk aversion $\eta \geq \mu^2/\sigma^2$ implies that $\Phi(t) = 1$ and $R(t) = 0$ for all t .

What remains is to calibrate “reasonable” values of μ , σ , and η , followed by a plausible numerical exercise that conveys some implications of risk-adjusted gamma discounting for the distant future.

5 A Simple Numerical Gamma Example

The primitive exogenous driver in this paper is the uncertain future productivity of capital r , which is assumed to be a random variable distributed as a gamma PDF with mean μ and standard deviation σ . As previously mentioned, my base-case point-estimate deterministic values were ($r=6\%$, $\eta=3$, $g=2\%$). Now I investigate what happens when r is turned into a random variable by a mean-preserving spread in the form of a gamma distribution with $\sigma > 0$. Thus, under uncertainty, my base-case average real rate of return on capital is $E[r]=6\%$ per year. (Again, I personally view $E[r]=6\%$ as being much too high, but the “shock value” of the numerical example is to show that even with $E[r]$ as high as 6% there might be strong effects.) My base-case CRRA coefficient remains $\eta=3$. My base-case average growth rate is $E[g]=2\%$ per year. Notice that these base-case expected values are consistent with the fundamental Ramsey equation (9). Were η to be changed substantially, then $E[r]$ and $E[g]$ would not mesh quite so nicely with past reality. If $\eta=2$ and $E[r] = 6\%$,

then (9) implies $E[g] = 3\%$, which is perhaps too high. If $\eta=4$ and $E[g] = 2\%$, then (9) implies $E[r]=8\%$, which is also perhaps too high. So I think it is fair to say that this proposed package of base-case values ($E[r]=6\%$, $\eta=3$, $E[g]=2\%$) is internally consistent and looks more or less “realistic” – even though the zero rate of pure time preference is a tip-off that the analysis is intended to be normative.

What remains to fully specify my base case under uncertainty is to choose σ . I choose for my base case $\sigma = 3\%$ per year. Aside from having some desirable calibration properties (as will be shown), $\sigma=3\%$ also has the useful quality of being a round number that will simplify numerical calculations. The Stern Review used a deterministic discount rate of r not very much above 1% per year¹⁰, which was widely criticized by other economists as being much too low.¹¹ The value $\sigma = 3\%$ (along with $\mu = 6\%$) implies from gamma cumulative distribution functions for (22), (27) that $P[r < 1\%] \approx .005 = 1/200$, which I would regard as an extraordinarily small probability that the future rate of return r will be less than 1%. Speaking loosely, I am assigning a subjective confidence level of 99.5% to r being above 1%, thereby, I hope, forestalling criticisms that the results of this paper are coming from assuming too-high probabilities of too-low rates of return in the left tail of the gamma PDF.

The numerical results of this paper are definitely sensitive to the choice of σ . If $\sigma > 3\%$, then the striking outcomes of the numerical examples to be presented will become yet more pronounced, thus making my case stronger. (Even just choosing $\sigma = 3.5\%$ makes a big difference in strengthening my numerical results.) Therefore, I do not explore further in the direction of $\sigma > 3\%$. In the other direction, if σ is chosen to be even a little bit lower than 3% then the implied probabilities of low- r values become almost ridiculously small. For example, if $\sigma = 2.5\%$, then $P[r < 1\%] \approx 1/1300$, which I would regard as an incredibly low probability. (Even my base case $P[r < 1\%] \approx 1/200$ seems to me to be bending over backwards to assign tiny probabilities to low values of r .)

With $\mu=6\%$ and $\sigma=3\%$, the implied gamma PDF is shown as the curve in Figure 1. The mean is 6%, the median is 5.5%, and the mode is 4.5%.

¹⁰Stern (2006) actually chose as base case $r \approx 1.4\%$ per year, which, taking literary licence here, I am interpreting as “not very much above 1% per year.”

¹¹Among many other such criticisms, see, e.g., Dasgupta (2008), Nordhaus (2007) or Weitzman (2007).




Figure 1. Gamma PDF of $y=f(r)$ and $x=r$ for base-case example.

With $\mu=6\%$ and $\sigma=3\%$, some numbers describing the cumulative distribution function of r at the lower end of its range are given in the following table.

$\hat{r} =$.5%	1%	1.5%	2%	3%
$P[r \leq \hat{r}] =$.04%	.49%	1.9%	4.6%	14.3%

Table 2: Lower-End Probabilities of Future Capital Productivity

I think it is fair to say that the numbers in Table 2 portray a situation where especially low values of the future productivity of capital are extremely unlikely. Nevertheless, it turns out that, even with these arguably much-too-small probabilities, there is sufficient fear of catastrophic low-productivity states to drive the strong numerical results that follow.

Plugging the base case $\eta = 3$, $\mu=6\%$, and $\sigma=3\%$ into (28) and (29) gives the super-simple formulas

$$\Phi(t) = \left(\frac{1}{1 + .015t} \right). \quad (30)$$

and

$$R(t) = \frac{\ln(1 + .015t)}{t}. \quad (31)$$

Some numerical values of $R(t)$ and $\Phi(t)$ from (30) and (31) are given in the following table.

$t =$	50	100	150	200	300	500	1,000
$R(t) =$	1%	.9%	.8%	.7%	.6%	.4%	.3%
$\Phi(t) =$	57%	40%	31%	25%	18%	12%	6%

Table 3: $R(t)$ and $\Phi(t)$ as functions of future year t

Even considering the enormous magnitude of the simplifications that have gone into this model, I find the numbers given in Table 3 to be striking. The reader should confirm that the differences between the discount factors $\Phi(t)$ from Table 1 and $\Phi(t)$ from Table 3 are enormous. I will not clutter up the paper further by considering too many alternative numerical variants of Table 3, but suffice it here to say that results are still notable when parameter values are reasonably close to my base case. As just one example, consider the effect of different coefficients of relative risk aversion on the discount factor $\Phi(t)$ from choosing $\eta=2.5$ or $\eta=3.5$ (instead of the base case $\eta=3$). In the interest of saving space, I will not reproduce two supplementary versions of Table 3 for alternative values $\eta=2.5$ or $\eta=3.5$, but content myself here with showing what happens for a “representative” time of two centuries hence ($t=200$). From Table 3, $\Phi(200 | \eta = 3) = 25\%$. Simple calculations from (28) indicate that $\Phi(200 | \eta = 2.5) = 12.5\%$, while $\Phi(200 | \eta = 3.5) = 50\%$. So there is genuine sensitivity of Φ to η , with higher values of η giving higher values of Φ . But I think it is fair to say that the “big picture” is not dramatically altered. Note that it is very easy to perform sensitivity analysis here. Given the simple closed-form solutions (29) and (28), readers can readily plug in their own favorite parameter values and calculate what happens.

Why do such striking results as shown in Table 3 emerge from the small amount of probability calibrated (via choosing $\sigma=3\%$) to $P[r \leq 1\%] \approx 1/200$? The answer is that the moderate degree of relative risk aversion represented by $\eta=3$ is sufficient (along with the empirically-compatible zero rate of pure time preference) to make a representative agent fear greatly catastrophic states of low- r productivity, even though their probabilities of occurrence (from Table 2) are tiny. As was already pointed out in connection with (1)-(4), over time the higher rates of return discount themselves exponentially out of existence, enhancing over time the impact of lower and lower rates of return. This type of effect is enormously magnified after risk-adjusted weighting of probabilities by marginal utilities. What drives the striking results of Table 3 is the great fear (via $\eta=3$) of catastrophic damages to productivity contained in the far left tail of the gamma PDF. This example seems to be hinting that the “fear factor” associated with low-probability high-impact catastrophic states may play a significant role in long-term discounting. In a sense, low effective discount rates and high effective discount factors are being applied to CBA as a kind of insurance policy against the possibility of really bad future outcomes.

A simple numerical illustration of the effect of this “fear factor” is the following. Consider the value $r < 0.5\%$, which from Table 2 occurs with probability $\approx 1/2500$. Because the $r < 0.5\%$ state occurs only with a seemingly negligible probability, it might be reasonable to expect a negligible impact on the discount factor. However, this $r < 0.5\%$ state is catastrophic because, in optimal Ramsey growth with no time preference, it forces current generations to severely repress consumption to save capital for future generations. A zero “utility discount rate” prohibits shifting the bad consequences onto the far-distant future. It turns out that $E[r \mid r \leq .5\%] = .4\%$. In what follows, for an illustrative exercise I use this .4% number as a certainty equivalent. Compared with the base case $r = 6\%$, from equation (10) the ratio of the two initial consumptions is $6/.4 = 15$, meaning that $C_{6\%}^*(0) = 15 \times C_{.4\%}^*(0)$. With $\eta = 3$, this factor of 15 for consumption ratios translates into a factor of $15^3 = 3375$ for the ratio of initial marginal utilities. This huge factor of 3375 for the ratio of initial marginal utilities makes itself felt strongly even though $P[r \leq .5\%] \approx 1/2500$. Apparently, a relative risk aversion coefficient of $\eta = 3$ is powerful enough (via its cubic leverage effect on marginal utility) to make a sizable impact on distant-future discounting.

The risk-adjusted gamma discounting examples of this paper have $C_r^*(0)$ being strongly reactive to catastrophically-low values of r because $C_r^*(0) \propto r$. It is here that Ramsey’s zero discounting of future utilities makes itself most felt. With no time preference, there is no incentive to shift the burden of limited consumption prospects forward to distant future generations. Rather, low- r states cause present generations to sacrifice consumption in the present. This feature, when greatly leveraged by risk aversion into marginal-utility-adjusted probabilities, largely drives the numerical results of Table 3. As for interpretation, I believe that the results of this paper should be understood in the normative spirit of the Ramsey model itself (with its “ethical” zero rate of pure time preference), rather than in any literal descriptive sense.

An especially insightful way to understand what is happening is to ask what is the *risk adjusted* PDF that corresponds to (22). Let this risk-adjusted PDF, where the probabilities are weighted by marginal utilities, be denoted $h(r)$. From combining (22) and (27) with (7) and (10) (for $t=0$), the relevant *marginal utility weighted* PDF corresponding to (22) is also a gamma distribution of form

$$h(r) = \frac{b^a}{\Gamma(a)} r^{a-1} e^{-br}, \quad (32)$$

where, here,

$$a = \frac{\mu^2}{\sigma^2} - \eta, \quad b = \frac{\mu}{\sigma^2}. \quad (33)$$

Fortuitously for giving strong insights, in the base case here ($\eta = 3$, $\mu = 6\%$, $\sigma = 3\%$), the PDF $h(r)$ takes an especially simple form. When these particular base-case parameter

values ($\eta=3$, $\mu=6\%$, $\sigma=3\%$) are plugged into (33), then $a = 1$ and $b = 1/.015$. The PDF (32) then becomes the especially simple exponential form

$$h(r) = \lambda \exp(-\lambda t) \tag{34}$$

with $\lambda=66\frac{2}{3}$, which is depicted in Figure 2.




Figure 2. Risk Adjusted PDF of $y=f(r)$ and $x=r$ for base-case example.

What is happening here is quite remarkable. The risk-adjusted version of the gamma distribution for the base case ($\eta=3$, $\mu=6\%$, $\sigma=3\%$) becomes the exponential distribution (34), whose mean is 1.5% and median is $\approx 1\%$. Compared with the gamma distribution of Figure 1, the exponential distribution in Figure 2 shifts the marginal-utility-adjusted PDF extremely strongly toward putting large risk-adjusted probability weights on very low values of r . To see what it is doing in the marginal-utility-weighted PDF (34), some numbers describing the risk-adjusted cumulative distribution function of r at the lower end of its range are given in the following table.

$\hat{r} =$.5%	1%	1.5%	2%	3%
$P[r \leq \hat{r}] =$	33%	55%	70%	80%	91%

Table 4: Lower-End Risk Adjusted Probabilities of Future Capital Productivity

The risk-adjusted results of Table 4 should be compared with the non-risk-adjusted results of Table 2. Because the lower-tail probabilities of r in Table 4 are so much higher than the lower-tail probabilities of r in Table 2, the risk adjustment of probabilities by marginal utility weights makes an extraordinary difference on raising significantly distant-future discount factors.

6 Discussion

Needless to say, an overwhelming number of caveats apply to the super-simple model of this paper. When layers of reality are piled on, there is no way that such strong results as those shown in Table 3 will continue to apply.

For one thing, the idea that the future productivity of capital is set in stone tomorrow, and consumption is instantaneously adjusted immediately after making the CBA decision today, is very extreme, to put it mildly. Nevertheless, I believe that some of the flavor of the model of this paper should remain even after waiting a reasonable amount of time for the future rate of return to reveal itself. If the permanent productivity shock occurs after T years, then one should multiply the $\Phi(t)$ values in Table 3 by some appropriate discount factor of form $\Phi_T = \exp(-r_0T) < 1$, thereby lowering the applicable discount factors in Table 3. Even so, I believe that the effects of this paper will likely show themselves, albeit in dampened form.

The suggestive results of Table 3 seem as if they might be sufficiently powerful that an appropriately muffled version would likely survive some other modest modifications to the basic model. For example, there will presumably be a muffled version that describes situations of weak mean reversion where low productivity can persist for a very long time, but not forever. All in all, I suspect that it may be difficult to dislodge the finding that, at least for some reasonable specifications, risk-adjusted gamma discounting might have some very different implications for evaluating the distant future than discounting at a constant rate.

As was mentioned earlier, under uncertainty the effect of higher or lower risk aversion embodied in the coefficient η has the opposite effect on discounting than what is conventionally believed. The conventional normative interpretation, whether implicitly or explicitly, begins by postulating some growth rate g , and then infers from the fundamental Ramsey equation (9) what is the implied discount rate r , given the assumed degree of relative risk aversion η .¹² Seen in this light, the conventional wisdom is that, other things being equal (i.e., for given projected g), higher values of risk aversion η imply higher discount rates r .

¹²See, e.g., Arrow et al (1995)

Thus, proponents of low discount rates for climate change CBA typically favor low values of relative risk aversion (as well as near-zero rates of pure time preference). For example, the Stern Review¹³ proposed $g \approx 1.3\%$ per year, $\eta \approx 1$, and $\rho \approx .1\%$ per year (as a symbolic proxy for $\rho = 0$), which implied that $r \approx 1.4\%$ per year. Since discounting at $r \approx 1.4\%$ can give a very different CBA from discounting at a rate which is closer to real-world average annual rates of return (say, for example, $r \approx 6\%$), a firestorm of controversy swirled around whether or not Stern’s results were due primarily to discounting at such a low rate.¹⁴

This paper is hinting that it might matter whether a general-equilibrium or a partial-equilibrium approach is taken in interpreting the normative Ramsey discounting equation under uncertainty. It seems to me that the fundamental unknown for the distant future is productivity. The realized future productivity of capital r then induces (via an assumed value of η) the endogenously-chosen future growth rate *and* the endogenously-chosen consumption level, more than the other way around. In the climate-change application, the uncertainty about the distant-future productivity of capital is likely exacerbated by the unknown effects of climate change itself, because the impact of the relevant “damages function” on distant-future productivity seems so fundamentally uncertain, especially at high temperatures.¹⁵

When an uncertain future real rate of return on capital r drives an uncertain future growth process of the form $C(t) = C(0) \exp(gt)$, where both $C(0)$ and g are endogenously chosen to maximize Ramsey welfare, the effect of the risk aversion coefficient η on discounting is reversed from the situation where both $C(0)$ and g are viewed as exogenous and there is no optimization. Within the uncertainty framework of this paper, higher value of η cause *lower* “effective” discount rates. As was shown by numerical example, the impact of uncertain r on lowering the “effective” distant-future discount rate can be quite pronounced, even for modest degrees of relative risk aversion corresponding to $\eta \approx 3$. Under this interpretation, the Stern Review might have made a stronger case for low “effective” discount rates if, in addition to a near-zero rate of pure time preference, it had postulated $\eta \approx 3$ along with uncertainty about future productivity (including damages under climate change). Such an approach might have been more difficult to criticize on grounds of not matching reality-based projections, because it would be consistent with $E[r] \approx 6\%$ per year and $E[g] \approx 2\%$ per year.

¹³Stern (2006).

¹⁴Among many other such criticisms, see, e.g., Dasgupta (2008), Nordhaus (2007) or Weitzman (2007).

¹⁵This argument, however, requires a tricky fixed-point reinterpretation of the reduced form PDF of r , because higher growth rates can cause higher emissions, which make future temperatures higher, which drives down future productivity.

7 Concluding Remarks

When future discount rates are uncertain but have a permanent component, then the normative discount rate declines over time toward its lowest possible value. The present paper has applied this basic idea to risk-adjusted gamma discounting. The core model is so crude that its implications are nowhere near to being definitive, but rather should be seen as suggestive at best. Some simple numerical examples hint that adjustments for risk, in the form of marginal-utility-weighted probabilities in a Ramsey optimal growth setting, may exert a possibly strong effect on raising the normative distant-future discount factor. Empirically, this feature might have ramifications for climate-change CBA – by weighting the distant future more heavily than is done by standard exponential discounting at a constant rate.¹⁶ In the context of climate change, where reduced-form productivity is itself an uncertain function of future temperature increases, the findings of this paper constitute yet another warning that the results of CBA can be largely driven by the “fear factor” associated with low-probability high-impact catastrophes.¹⁷

References

- [1] Arrow, Kenneth J. (1999), in Portney and Weyant (eds). *Discounting and Intergenerational Equity*. Resources for the Future.
- [2] Arrow, K. J., W. R. Cline, K.-G. Maler, M. Munasinghe, R. Squitieri, and J. E. Stiglitz (1995). “Intertemporal Equity, Discounting, and Economic Efficiency.” In Bruce, Lee and Haites (eds.), *Climate Change 1995: Economic and Social Dimensions of Climate Change*. Cambridge University Press.
- [3] Dasgupta, Partha. (2008). “Discounting climate change.” *Journal of Risk and Uncertainty*, 37: 141-169.
- [4] Gollier, Christian. (2009). “Should I discount the far-distant future at its lowest possible rate?.” *Economics: the Open-Access, Open-Assessment E-Journal*, vol. 3: 2009-25.
- [5] Gollier, Christian, and Martin L. Weitzman (2009). “How Should the Distant Future be Discounted When Discount Rates are Uncertain?.” Working paper submitted to *Economics Letters*, November 2009.

¹⁶As examples in the (less strong) case of non-risk-adjusted probabilities, see, e.g., Newell and Pizer (2003), or Groom, Koundouri, Panopoulou and Pantelidis (2007).

¹⁷For some other examples, see Weitzman (2009a) and Weitzman (2009b).

- [6] Groom, B., P. Koundouri, E. Panopoulou and T. Pantelidis (2007). “Discounting the distant future: how much does model selection affect the certainty equivalent rate?”. *Journal of Applied Econometrics*. 22 (3): 641-656.
- [7] Hepburn, Cameron. and Ben Groom (2007). “Gamma discounting and expected net future value.” *Journal of Environmental Economics and Management*. 53 (1): 99-109.
- [8] Newell, Richard G. and William A. Pizer (2004). “Uncertain discount rates in climate policy analysis. ” *Energy Policy* 32 (4): 519-529.
- [9] Nordhaus, William D. (2007) “The Stern Review on the Economics of Climate Change.” *Journal of Economic Literature*, 45 (3), pp. 686-702.
- [10] Nordhaus, William D. (2008) *A Question of Balance*. Yale University Press.
- [11] Solomon, Susan, Gian-Kasper Plattner, Reto Knutti, and Pierre Friedlingstein (2009). “Irreversible climate change due to carbon dioxide emissions.” *Proceedings of the National Academy of Sciences* 106 (6): 1704-1709.
- [12] Stern, Nicholas (2006). *The Economics of Climate Change: The Stern Review*. Cambridge University Press.
- [13] Weitzman, Martin L. (1998) “Why the Far-Distant Future Should Be Discounted at Its Lowest Possible Rate.” *Journal of Environmental Economics and Management* 36 (3): 201-208.
- [14] Weitzman, Martin L. (2001). “Gamma Discounting.” *American Economic Review*, 2001 (March), 91 (1): 260-271.
- [15] Weitzman, Martin L. (2003). *Income, Wealth, and the Maximum Principle*. Harvard University Press.
- [16] Weitzman, Martin L. (2009a) “On Modeling and Interpreting the Economics of Catastrophic Climate Change.” *Review of Economics and Statistics*, 91 (1) [February]: 1-19.
- [17] Weitzman, Martin L. (2009b) “Additive Damages, Fat-Tailed Climate Dynamics, and Uncertain Discounting.” *Economics e-journal*, 3, 2009-39 (October 22).