

The Ramsey Discounting Formula for a Hidden-State Stochastic Growth Process

Martin L. Weitzman

Accepted: 7 August 2012 / Published online: 21 August 2012
© Springer Science+Business Media B.V. 2012

Abstract The long term discount rate is critically dependent upon projections of future growth rates that are fuzzier in proportion to the remoteness of the time horizon. This paper models such increasing fuzziness as an evolving hidden-state stochastic process. The underlying trend growth rate is an unobservable random walk hidden by noisy transitory shocks and recoverable only as a probability distribution via Bayesian updating. A simple expression is derived for the time-declining Ramsey discount rate. The components of this hidden-state Ramsey discounting formula are then analyzed, followed by a few remarks about possible implications and applications.

Keywords Discounting · Climate change · Long term · Hidden state

1 Introduction to Long Term Discounting

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. 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.

M. L. Weitzman (✉)
Department of Economics, Harvard University, Cambridge, MA, 02138, USA
e-mail: mweitzman@harvard.edu

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 compound interest forces us to say that what one might conceptualize as monumental 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, it may not be too much of an exaggeration to say that 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 investments for the long run, accompanied by much controversy about the implications for long-run discounting. For specificity, the investigation of this paper is focused sharply on CBA of small investments that only incrementally impact distant-future events. However, many climate change investments may themselves be big and involve large impacts. Yet even after some of these big investments have been made, the issue of choosing the appropriate long-term rate of return for additional marginal investments remains important and is emblematic of a whole series of thorny issues concerning how to conceptualize and how to evaluate actions spanning across many generations. The discount rate represents the current price of future benefits or costs, and as such it is a convenient shorthand vehicle for studying the general nature of investment tradeoffs over time.

The famous Ramsey equation relates the discount rate to the underlying growth rate of consumption. There is no harm in accepting the implicit causality of this route by conceptualizing the growth rate of consumption as being the more primitive entity, which then determines the discount rate via the Ramsey equation. Such an approach has the advantage of rigorously laying bare the exact structure of uncertainty and its resolution, which undergirds time-varying discount rates. This route has already been taken by many researchers, and is the route taken here. For the convenience of having a unified framework, I henceforth pretend that all research on time-varying discount rates was done *as if* it came from within this Ramsey framework where growth rates determine discount rates.

Behind the uncertainty about how to discount the distant future, therefore, lie uncertainties about what will be the growth rate of the global standard of living leading to this distant future. Many kinds of uncertainty are involved. One important uncertainty concerns long term prospects for technological progress and the degree to which this future technological progress will be able to meet, overcome, or fall behind the unknown unfolding challenges of resource, environmental, and historical constraints. Despite decades of research into the subject of economic growth, including the development of so-called endogenous growth theory, we are nowhere near being able to predict worldwide growth rates of the standard of living much beyond anything more fundamental than rough guesstimates of what the future portends. As we all know, individual views of long term future growth prospects can range from ultra-optimistic to ultra-pessimistic.

Thus, I think it is important to begin by recognizing that there is no deep reason of principle that allows us to extrapolate past growth rates into the distant future. The industrial revolution itself began some two centuries ago, and only slowly thereafter permeated throughout the world. Merely leafing through the pages of Angus Maddison’s monumental work *The World*

*Economy: A Millennial Perspective*¹ should convince a reader that the world-wide growth process is, in the book's summing-up, "uneven in space and time" and "in no way linear." It is difficult enough to predict worldwide economic growth in the next decade or two. Making growth projections over the next century or two seems almost unbearably fuzzy. Yet dealing with this kind of fuzzy distant uncertainty is exactly what is required to infer long-term future discount rates.

A line of research over the past few decades has shown that future growth rates of consumption which are uncertain but *persistent* can lead to a schedule of discount rates that decline over time. This line of research on the mechanics and consequences of time-declining discount rates due to various forms of growth-rate persistence is sufficiently comprehensively summarized in Christian Gollier's new book² that I do not feel the necessity to delve into details of the underlying models here, except to note one point.

For most of the underlying models in this literature, the variance of the *average* future growth rate over a given future time horizon declines (towards zero) as the horizon lengthens, due to the law of large numbers. In other words, the prediction interval for the average future growth rate typically narrows rather than widens with the prediction horizon. This seems to me like a counter-intuitive specification. The more distant is the time horizon, the more uncertain should we be now about the average future growth rate. It seems only natural that predicting the average growth rate over the next two centuries is fuzzier and more speculative than predicting the average growth rate over the next two decades. In this paper I examine the discounting consequences of a specification where the variability of the predicted average future growth rate widens with the time horizon.

It seems to me natural, at least on a first pass, to model such kind of increasingly uncertain growth scenario as an evolving hidden-state stochastic process in which the underlying trend growth rate is an unobservable random walk. The underlying random-walking trend growth rate is hidden by noisy transitory growth shocks and can only be recovered as a probability distribution via Bayesian updating in a manner first emphasized by Muth and Kalman. Such an approach highlights our critical inability to separate out what is permanent and predictable about our present growth experience from that which is temporary and transient. In this paper I propose to apply the simplest standard model of a hidden-state stochastic process to the topic of determining the long term discount rate.

My intent, then, is to merge the Muth-Kalman updated-prediction apparatus with the Ramsey discounting apparatus. I believe that such a hidden-state formulation raises important questions about modeling future growth rates and sheds useful new light on the sources of a time-declining future discount rate. One source of a time-declining future discount rate is that we are uncertain about the *present* underlying trend growth rate, which effect would remain to be played out over time even if the *future* trend growth rate were to remain still and do no further meandering. The other source of a declining discount rate is that we are uncertain about the *future* trend growth rate of a random walk, which effect would remain even if we knew exactly the state of the *present* underlying trend growth rate.

The Ramsey formula for deriving a discount rate schedule from a single representative agent, who stands in for an enormously heterogeneous real world having widely dispersed growth rates, degrees of risk aversion, and rates of pure time preference, all of which are then projected into a hazy distant future, is at the outer limits of abstract modeling. We wouldn't be taking this route if the issues were not so important. For all of its oversimplifications, I believe that the Ramsey approach yields fundamental insights—and that the novel combination of

¹ Maddison (2001). See especially Table 8b on page 263.

² Gollier (2012).

the Ramsey equation with a hidden-state model of uncertain growth rates is especially fruitful in unpacking a few key relationships hitherto unexplored.

I thus view the hidden-state Ramsey approach primarily as a conceptual device for understanding an important series of interconnected issues concerning discounting in the presence of a hazy future. But there is no use pretending that such a super-aggregated super-abstract model has direct operational consequences obtainable by merely plugging in numbers and directly obtaining quantitative outcomes. In this paper I emphasize overall simplicity and understandability by leaning heavily on specifications having great analytical tractability. The formulation here yields a remarkably simple expression for a time-declining discount rate, which, I believe, gives some fundamental insights into the nature of long-term discounting that would otherwise not be available. The basic components of this simple discounting formula will be analyzed in terms of underlying parameters, followed by a few speculative remarks about possible implications and applications. The purpose of this paper is purely to outline the role of a previously neglected hidden-state aspect of the thorny problem of distant future discounting, rather than to propose an actual usable schedule of long term discount rates.

A word of warning is in order about any model attempting to derive discount rate schedules for the distant future (including the model of this paper). The models we use in such exercises are all oversimplified approximations. We have relatively more confidence in the approximations when conditions being investigated are not too different from the current familiar situation. As we move further into the future, however, we have relatively less confidence in the approximations because we are attempting to extrapolate the local into the global. Typically, some implication or another of the model is not credible in the limit as future time approaches infinity. For example, as noted above, most models used to derive future discount rates in the existing literature have the unbelievable feature that the variance of the average future growth rate approaches zero in the limit, so that by the law of large numbers we effectively know with certainty the long term growth rate. The model of this paper does not have this particular unbelievable feature, but instead has the different unbelievable feature that discount rates decline indefinitely, approaching minus infinity in the limit. The underlying problem with all such models of discount-rate formation is that the oversimplified local approximations are increasingly breaking down as the time horizon lengthens. I do not think there is any general cure for this problem except to be careful when trying to use such oversimplified models for much more than giving general insights into the process of long term discounting.

2 Growth Rates as Random Walks

This section attempts to succinctly provide the backdrop analytical framework that will be used to develop the hidden-state stochastic growth process used throughout the rest of the paper.

Time is broken up into discrete periods represented by the integer-valued variable t . The present corresponds to time zero, while future times correspond to integer values $t > 0$ and past times correspond to integer values $t < 0$. Consumption in period t is C_t . The growth rate of consumption in period t is

$$\ln C_t - \ln C_{t-1} = Y_t. \quad (1)$$

In turn, Y_t is decomposed into the expression

$$Y_t = X_t + z_t \quad (2)$$

where X_t is the *permanent* component of the growth rate at time t , while z_t is a *transitory* shock to the growth rate at time t . The permanent growth-rate component X_t is a random walk satisfying

$$X_t = X_{t-1} + w_t. \tag{3}$$

The random variables $\{z_t\}$ and $\{w_t\}$ are mutually independent iid Gaussian processes with distributions

$$z_t \sim \mathcal{N}(0, V_y) \tag{4}$$

and

$$w_t \sim \mathcal{N}(0, V_x) \tag{5}$$

The stochastic process (2)–(5) constitutes a standard model well represented in the literature. (The notation V_y is used as a mnemonic in (4) because the variance of z_t impacts Y via (2); likewise the notation V_x is used in (5) because the variance of w_t impacts X via (3).)

Let $t > 0$ be some future time. Define new random variables

$$Z_t \equiv \sum_{s=1}^t z_s \tag{6}$$

and

$$W_t \equiv \sum_{\tau=1}^t \sum_{s=1}^{\tau} w_s \tag{7}$$

Carefully iterating out and summing the Eqs. (1)–(3) yields, after applying (6), (7), the basic result

$$\ln C_t - \ln C_0 = Z_t + t X_0 + W_t. \tag{8}$$

From Eqs. (4) and (6), it is readily apparent that

$$Z_t \sim \mathcal{N}(0, t V_y). \tag{9}$$

As for W_t , reversing the order of summation in (7) gives

$$W_t = \sum_{s=1}^t (t - s + 1) w_s, \tag{10}$$

whose variance is $V_x \sum_{s=1}^t (t - s + 1)^2$, which in continuous time becomes the much neater formula $V_x \int_0^t (t - s)^2 ds = V_x t^3/3$. If we had chosen to express the difference Eq. (3) as a continuous-time stochastic diffusion process, we would then have obtained the neater formula directly. Leaving aside the arbitrary discreteness of the period length, we have shown that the random walk (3) exerts a stochastic influence on the random variable $\ln C_t - \ln C_0$ whose variance is *cubic* in time:

$$W_t \sim \mathcal{N}(0, t^3 V_x / 3). \tag{11}$$

Where X_0 observed directly, formula (8) would be sufficient to push on further with the analysis in the reduced form

$$\ln C_t - \ln C_0 \sim \mathcal{N}(t X_0, t V_y + t^3 V_x / 3). \tag{12}$$

However, when the permanent component of the growth rate is *not* directly observable, then X_0 becomes a random variable that must itself be inferred from the noisy past growth rates that *have* been observed. The next section describes briefly this process of hidden-state inference and updating.

3 Hidden-State Growth Rates

This section develops succinctly the simplest standard model of a Muth-Kalman hidden-state stochastic process. This baby model is applied to growth rates that will subsequently be plugged into the Ramsey formula for determining discount rates.

The growth of consumption during period t is given by Eq. (1), which is the only part of the stochastic system that is directly observable (after its realization). For analytical convenience, I assume that in any period there is an infinite past record of previous observations.

When period t opens, with Y_{t-1} having just been observed at the very end of period $t - 1$, suppose that

$$X_{t-1} \sim \mathcal{N}(\mu_{t-1}, V_{xy}) \tag{13}$$

for some mean μ_{t-1} and variance V_{xy} yet to be determined. (Later it will become clear why V_{xy} carries the mnemonic subscript xy .) From the symmetry of the situation with an infinite number of observations of past growth rates, it is reasonable to postulate (and can be proved rigorously) that V_{xy} is a constant independent of t .

At the very beginning of period t (immediately after Y_{t-1} has been observed), the random walk described by (3) occurs and $X_{t-1} \rightarrow X_t$. The realization of the unobserved random variable w_t adds variance V_x to the variance V_{xy} in (13), so that the variance of the estimate of X_t increases to $V_{xy} + V_x$ throughout period t . At the very *end* of period t , the observation Y_t is made, which adds precision $1/V_y$ and thereby lowers the variance of X_t back down to V_{xy} according to the equation

$$\frac{1}{V_{xy} + V_x} + \frac{1}{V_y} = \frac{1}{V_{xy}}. \tag{14}$$

The corresponding Bayesian updating formula for μ_t is

$$\mu_t = \frac{\frac{1}{V_{xy}+V_x} \mu_{t-1} + \frac{1}{V_y} Y_t}{\frac{1}{V_{xy}+V_x} + \frac{1}{V_y}}. \tag{15}$$

Equations (15) and (14) reflect the fact that Bayesian updating of normal distributions is expressed in terms of additive weights of the *precision*, which equals one over the variance.

Solving the quadratic equation (14) for V_{xy} , one obtains the formula

$$V_{xy} = \frac{\sqrt{4V_x V_y + V_x^2} - V_x}{2}. \tag{16}$$

Notice in (16) that $V_{xy} = 0$ if *either* $V_y = 0$ or if $V_x = 0$. When V_y is much larger than V_x , then a good approximation for the formula (16) is

$$V_{xy} \approx \sigma_x \sigma_y. \tag{17}$$

Equation (15) can be solved recursively to yield (for time zero)

$$\mu_0 = (1 - \lambda) \sum_{s=0}^{\infty} \lambda^s Y_{-s}, \tag{18}$$

where

$$\lambda = \frac{V_y}{V_{xy} + V_x + V_y}. \tag{19}$$

To summarize, the mean of the normally distributed random variable ${}^t X_0$ is $t \mu_0$, while its variance is $t^2 V_{xy}$, so that

$${}^t X_0 \sim \mathcal{N}(t \mu_0, t^2 V_{xy}). \tag{20}$$

Decomposing Eq. (8) [or (12)] into the independent normal probability distributions (9), (11), (20) gives the basic result that

$$\ln C_t - \ln C_0 \sim \mathcal{N}\left(\mu_0 t, V_y t + V_{xy} t^2 + \frac{V_x}{3} t^3\right). \tag{21}$$

It is much easier to think in terms of average growth rates than changes in the log of consumption as given by formula (21). If we define the *average* future growth rate over the time horizon $[0, h]$ to be the random variable

$$G_h = \frac{\ln C_h - \ln C_0}{h}, \tag{22}$$

then from (21) it follows that

$$G_h \sim \mathcal{N}\left(\mu_0, \frac{V_y}{h} + V_{xy} + \frac{V_x}{3} h\right). \tag{23}$$

From formula (23) the variance of the *average* growth rate eventually *increases* with the length of the future horizon h (for $h > \sqrt{3V_y/V_x}$). In other words, the prediction interval for the average growth rate widens with the prediction horizon. I would view this as an intuitively desirable property because, once the law of large numbers has eliminated stationary background noise, I think that future average growth rates should become increasingly uncertain as the time horizon is more remote. However, the other side of this same coin is that the probability of large negative (or large positive) growth rates becomes increasingly likely over a longer time horizon. Both features are outcomes of the random walk specification (2). It is difficult to model the one feature without the other (while preserving a consistently symmetric treatment of time), so there may be a deep issue of model tradeoffs involved here. Perhaps the best interpretation is to view (23) as an acceptable approximation for large h (centuries?) but not for *too* large h (millennia?). Perhaps an asymmetric treatment of time is warranted for very large h . Perhaps some degree of mean reversion needs to be introduced, although this seems to make learning about hidden states an analytically intractable problem. Perhaps stochastic processes that rule out negative values (like expressing all variables in logarithms) would be more appropriate, but again analytical tractability would have to be sacrificed. In any event, this paper is devoted to analyzing the consequences for Ramsey discounting of the simplest standard model of a Muth-Kalman hidden-state stochastic growth process, to which application we now turn.

4 Ramsey Hidden-State Discount Rates

Viewed from the present time $t = 0$, consumption C_t at future time $t > 0$ is a random variable. The normative welfare criterion used throughout this paper is the standard formula for expected present discounted utility of consumption

$$E \left[\sum_{t=0}^{\infty} e^{-\rho t} U(C_t) \right], \tag{24}$$

where $U(C)$ is the utility function of consumption and $\rho \geq 0$ is the rate of pure time preference or utility discount rate.

Suppose that someone proposes a marginal “project” that costs some infinitesimally small amount δ of consumption now (at time $t = 0$) but will safely deliver some infinitesimally small amount ϵ of consumption at time $t > 0$ in the future, independent of the state of the world then. Should this project be accepted? This is the prototype question of investment criteria that we are seeking to answer from within the Ramsey hidden-state framework.

Applying the additively-separable expected present discounted form (24) to this problem, we should *accept* the marginal project if and only if

$$\epsilon e^{-\rho t} E[U'(C_t)] > \delta U'(C_0). \tag{25}$$

We can translate the investment criterion (25) into the more familiar language of a continuously-compounded discount rate as follows. Define the relevant discount rate for investment payoffs at time t to be the value r_t satisfying the equation $\exp(-r_t t) = e^{-\rho t} E[U'(C_t)]/U'(C_0)$, which can be rewritten as

$$r_t = -\frac{1}{t} \ln \left(\frac{e^{-\rho t} E[U'(C_t)]}{U'(C_0)} \right). \tag{26}$$

Then applying (26) to (25), the project should be approved if and only if

$$\epsilon \exp(-r_t t) > \delta. \tag{27}$$

Adding up linearly from (27), the basic principle of CBA here is that safe marginal investments should be undertaken if and only if their present discounted benefits exceed their present discounted costs, where the discount rate schedule to be applied to payoffs at future time t is given by formula (26).

To get a more interpretable formula than (26), we need to put more structure on the problem. Pursuing the spirit of utmost simplicity that this paper tries to embrace, without further ado I just postulate further that $U(C)$ is of the isoelastic or constant relative risk aversion (CRRA) form

$$U(C) = \frac{C^{1-\eta}}{1-\eta}, \tag{28}$$

where $\eta > 1$ is the coefficient of CRRA. Note from (28) that marginal utility here takes the simple form $U'(C) = C^{-\eta}$. The CRRA specification (28) is a familiar standby in applications, but its assumption is not innocuous. The same comment could apply to other familiar specifications and functional forms being assumed throughout this model. The saving grace is that one obtains a relatively simple discount-rate formula, which gives some clear basic insights that hopefully apply more generally.

Applying (28) to (26) enables the latter to be rewritten as

$$r_t = -\frac{1}{t} \left(\ln e^{-\rho t} + \ln E \left[(C_t/C_0)^{-\eta} \right] \right). \tag{29}$$

Plug the identity $(C_t/C_0)^{-\eta} = \exp(-\eta (\ln C_t - \ln C_0))$ into (29), thereby obtaining the expression

$$r_t = \rho - \frac{1}{t} \ln E \left[\exp(-\eta (\ln C_t - \ln C_0)) \right]. \tag{30}$$

Plug the distribution (21) into (30). Make use of the standard formula for the expectation of a lognormal distribution and rearrange terms to obtain, at last,

$$r_t = \rho + \eta \mu_0 - \frac{\eta^2}{2} \left(V_y + V_{xy}t + \frac{V_x}{3}t^2 \right). \tag{31}$$

Equation (31) is the Ramsey riskfree discount rate formula for the simplest model of a hidden-state Muth-Kalman stochastic growth process. The remainder of the paper is devoted to analyzing this time-dependent Ramsey discount rate schedule.

It is immediately obvious from Eq. (31) that discount rates eventually turn negative. As has been remarked upon already in connection with the stochastic behavior of growth rates, this disturbing feature can be removed by any one of several possible technical changes in the model, but they all destroy the simplicity of a formula like (31) and replace it by an analytically intractable equation. The unit root assumption implicit in the random walk specification is responsible for the extreme limiting situation that $r_t \rightarrow -\infty$ as $t \rightarrow \infty$. Some minimal amount of mean reversion (to a known positive mean) would bound r_t from below, while sufficiently strong mean reversion (to a known positive mean) would cause the probability of negative r_t to become vanishingly small. For what it is worth, two econometric studies using a time series of over 200 years of interest rate observations weakly favor the hypothesis of a unit root over the hypothesis of mean reversion.³ Alas, the presence of mean reversion seems to make learning about the hidden state analytically intractable.

The eventually negative discount rates in formula (31) can be viewed as an artifact of the possibility of highly negative growth outcomes from the normal distribution in the specification of the random walk Eq. (3), along with the assumption of a CRRA utility function (28). The representative agent in this setup is very averse to large negative growth rates, which become increasingly likely over time under a random walk. Eventually, the representative agent here would be willing to pay a positive price to guarantee the safe delivery of consumption over a long enough time period ahead. Again I point out that this undesirable feature can be avoided by a variety of ad hoc modifications of the model, but at a cost of sacrificing analytical tractability. Perhaps the simplest way to evade negative discount rates is to mechanically avoid considering the very long time periods over which they can manifest themselves. More basically and less literally, I think that the eventually negative discount rates in (31) are trying to tell us that, in a world where we are fundamentally unsure about how distant future growth prospects could evolve, there are potentially strong forces that might want to make us consider using very low discount rates for discounting distant-future events. All of this is testimony to the fear-power of being so very unsure about where the world is headed in the long run. This interpretation is not rigorously satisfying, but I think it contains a strong element of truth.

³ Newell and Pizer (2003) and Groom et al. (2007).

5 Analyzing Ramsey Hidden-State Discount Rates

The most intuitive way to understand formula (31) is to build it up piece by piece out of its constituent components.

When there is *no* uncertainty ($V_y = 0, V_x = 0$), then expression (31) reduces to the famous deterministic Ramsey equation for the discount rate

$$r_t = \rho + \eta \mu_0. \tag{32}$$

When there is *only* the uncertainty of transitory shocks ($V_y > 0, V_x = 0$), then (31) reduces to a familiar expression for the Ramsey equation modified for the presence of iid-normal growth uncertainty

$$r_t = \rho + \eta \mu_0 - \frac{\eta^2}{2} V_y. \tag{33}$$

Formula (33) indicates the well known result that the presence of iid growth shocks uniformly lowers the Ramsey discount rate, because the attractiveness of a risk-free investment increases when consumption is uncertain, but it does not cause the discount rate to change over time. For the Ramsey discount rate on a safe investment to exhibit time dependence of its term structure in an otherwise symmetric setting requires some persistence of growth uncertainty.

When there is *only* the uncertainty of the random walking permanent growth-rate component ($V_y = 0, V_x > 0$), then (31) reduces to the expression

$$r_t = \rho + \eta \mu_0 - \frac{\eta^2}{6} V_x t^2. \tag{34}$$

It is perhaps not a surprise that modeling the permanent component of future growth rates as a random walk causes a declining discount rate. What is perhaps more surprising is that the decline in the discount rate schedule is *quadratic* in time. If we think of V_x as being very small, then this quadratic time effect may take a long time to exert its influence. But once exerted, its influence can quickly expand quadratically to become very powerful.

Finally, when there is uncertainty about both the random-walking permanent component of growth rates *and* about transitory growth shocks ($V_y > 0, V_x > 0$), then the full equation (31) holds. If the quadratic time term $\eta^2 V_x t^2 / 6$ is expressing fear of where we might be *going* in the long term, the linear time term $\eta^2 V_{xy} t / 2$ is expressing fear of not knowing where we might be *now*. When V_y is much bigger than V_x , then the coefficient V_{xy} of the linear time component is much bigger than the coefficient $V_x / 3$ of the quadratic time component. So the linear time component causes the discount rate to decline more powerfully in the short run, while the quadratic component dominates in the long run.

In this hidden-state formulation, we are unsure what is the current underlying trend growth rate X_0 because we are unsure about how to read the noisy past record. This effect would remain even if future random walking suddenly ceased altogether. To isolate sharply the role of this linear time term $\eta^2 V_{xy} t / 2$ in (31), consider the following artificial thought experiment. Suppose that the full stochastic formulation of this paper describes perfectly the *past* behavior of the system, but that in the *future* all uncertainty will miraculously cease [(2) and (3) will no longer hold] so that effectively X_t (for $t > 0$) will henceforth remain frozen at X_0 . In this thought experiment, Eq. (31) would become

$$r_t = \rho + \eta \mu_0 - \frac{\eta^2}{2} V_{xy} t. \tag{35}$$

The “as if” story being told by (35) is the following. A take-it-or-leave-it decision must be made now, just before time zero (call it time 0^-), about whether or not to make a marginal investment that costs δ now but will yield a benefit payoff ϵ at future time t . Right now, at time 0^- , the constant permanent future value of the growth rate of consumption is a random variable whose realization is unknown. At time 0^+ , this constant permanent growth rate is realized as a random draw from $\sim \mathcal{N}(\mu_0, V_{xy})$. A decision must be made now (at time 0^- , just before the “true” state of the world is revealed at time $t = 0^+$) about whether or not to invest in the proposed project. Immediately after the investment decision is made, at time 0^+ , the true growth rate becomes known that will hold thenceforth, from time $t = 0^+$ to time $t = \infty$. What is the relevant discount rate?

In the above setup, the relevant discount rate for CBA is given by formula (35). The idea that uncertain future growth rates are permanent and materialize overnight is extreme, but it has been used in the literature as a point of departure for giving insights into declining discount rates.⁴ The point being made here is that such a story can be viewed as the hypothetical outcome of a situation that in the past was subject to hidden-state uncertainty, but in the future will not be subject to any further such uncertainty.

Equation (35) indicates that Ramsey riskfree discount rates will decline over time and eventually become negative, based just on uncertainty about what is the underlying permanent growth rate, even though this underlying permanent growth rate will henceforth stay put and not move away from its current (and currently unknown) value. As with the pure quadratic random-walking case (34), the pure linear fear-of-bad-growth case (35) comes from the long (even if thin) left tail of the normal distribution, along with the assumption of a CRRA utility function (28). Again here, a way to evade negative discount rates is to mechanically avoid considering the very long time periods in which they can manifest themselves. And again, more basically, I think that the eventually-negative discount rate in (34) is trying to tell us that, in a world where we are fundamentally unsure now about what will be distant future growth prospects, there are potentially strong forces that might want to make us consider using very low discount rates for discounting distant-future events.

In thinking about the declining discount rate schedules implied by formulas like (35), (34), or (31), the underlying normality assumption cuts two ways. On the one hand, the long negative tail of the normal distribution makes discount rates eventually become negative. If we artificially truncated the left tail, it would impair the ability to generate eventually-negative discount rates. On the other hand, the long negative tail of the normal distribution is very thin with probability measure, so that if it were made fatter with mean-preserving probability in the left side range where it is not truncated, discount rates could decline more rapidly to whatever low values they eventually attain.

The simplicity of formula (31), and the fact that it eventually yields negative values, depends on the specific functional forms of the model. However, I believe that the more general overarching message of the hidden-state model of this paper is the difficulty of avoiding eventually-low Ramsey discount rates in a world where there is fundamental uncertainty about interpreting past growth rates and projecting future changes in growth prospects.

6 A Numerical Example

Attempting to plug actual numbers into the model of this paper constitutes an unusually brazen act of hubris. The model itself is at a nearly cosmic level of abstraction. Whose

⁴ For example, [Gollier and Weitzman \(2010\)](#).

Table 1 Discount rates r_t (% per year) as function of time t (years)

t (years)	0	50	100	150	200	250
r_t (%)	3.8	3.5	2.9	2.1	1.2	.4

preferences and future prognostications is the representative agent in this model supposed to represent? So many different answers to this question are possible, with so many different possible numerical values to plug into formula (31), that it practically rules out decisive conclusions from the beginning.

For the record, I will attempt here just one such “illustrative” numerical exercise. I do not try hard to defend the numbers I use. The beauty of having a formula as simple as (31) is that readers can readily plug in their own favorite numbers and see almost immediately what emerges. Results are sensitive to the numbers being plugged in, but I let readers explore this topic for themselves.

For normative pure time preferences associated with intergenerational comparisons, I choose in this example to not favor any particular generation by selecting $\rho = 0$. Estimates of the relevant value of the coefficient of relative risk aversion vary according to the individual or the study. A not uncommon idea is that the CRRA coefficient should reasonably be somewhere between one and four. For this example I choose $\eta = 2$.

Worldwide per-capita growth rates have averaged about 2% per year in the time since 1950, but with enormous variation across regions and sub-periods. For the purposes of this example I choose $\mu_0 = 2\%$. The postwar variability of annual growth rates is relatively low in the advanced OECD economies, say $\sigma_y \approx 2\%$ on average, but it is considerably higher for the world as a whole or even for OECD countries prior to 1950. In this numerical example I choose $\sigma_y = 3\%$.

I think the trickiest, crudest, and most brazen calibration of all concerns the variability σ_x of the random walk for the permanent growth-rate component X_t . By (23), the term $\sigma_x^2 t/3$ is contributed to the variance of the *average* future growth rate by random walking of the underlying trend growth rate I choose σ_x by the following thought experiment. Suppose there is zero transient uncertainty ($\sigma_y = 0$) and $\mu_0 = X_0$ is *known* to be 2%. I then calibrate σ_x by requiring that the probability of a stagnant (no growth) century due to the random walk alone is one out of a million (10^{-6}). With a normal distribution, the one-sided probability 10^{-6} is a 4.265-sigma event. Therefore, the corresponding value of σ_x which makes such an extreme zero-average-growth random walk over a century have a probability of one out of a million satisfies

$$\frac{\sigma_x^2}{3} \times 100 = \left(\frac{.02}{4.265} \right)^2, \tag{36}$$

whose solution is $\sigma_x = .08\%$ per year. I readily admit the incredible crudeness of this calibration, but am unsure what else to do.

The corresponding Ramsey hidden-state discount schedule from formula (31) is enumerated in Table 1.

Considering the (unavoidable) casualness of the numbers plugged into formula (31), I do not want to over-interpret the results. Different people can plug in their own favorite numbers and may draw their own conclusions from what is but a single illustrative example. I think Table 1 indicates loosely that even a very small amount of random walking in the underlying trend growth rate can work its way through the hidden-state mechanism to have non-negligible long term impacts on lowering the discount rate.

7 Concluding Comments

I think that the paper suggests several themes.

First, there is a sense that how we model current fuzziness about future growth rates is likely to be an important ingredient in deriving and analyzing long-term discount rates. The extent to which predicting future growth is a highly unresolved issue is disconcerting. Modeling the growth process as a hidden-state stochastic process is one previously untried route to deriving discount rate schedules, which is long overdue for analysis. This route captures, however crudely, one aspect of what makes distant future growth so fuzzy. Perhaps we need to try other routes, some of which may be similarly unorthodox and are similarly overdue for analysis.

The hidden state model of this paper is suggesting a kind of crude taxonomy for discounting under uncertainty. In this approach, there are three stylized “source types” that push riskfree discount rates lower. One is independent noisy growth shocks, which lowers the Ramsey discount rate uniformly. A second “source type” of lower discount rates is the idea that we are ignorant of what is the current underlying growth rate trend amidst all the random noise of widely different past growth rates for various regions and periods. This current ignorance of the true state contributes to a discount rate that declines linearly over time. A third “source type” of lower discount rates is the idea that even if we knew the current underlying trend growth rate, it will likely change over time. The uncertainty of future random walking contributes to a discount rate that declines quadratically over time, although it is multiplied by a variance coefficient that is much smaller than the variance coefficient associated with the linear time term. A contribution of this paper is to spell out in detail how these three “source types” for declining discount rates interact with each other and depend differently upon the time horizon.

Over and above the crudeness of the basic model and the artificial precision of its formulas, I think the paper is hinting that random forces can act powerfully to lower discount rates over long time horizons. Even very small amounts of random walking can eventually work their way into having strong discounting effects for the distant future.

If, as I suspect, there is a fundamental problem with our long-term discounting models depending heavily on how we conceptualize and express fuzzy future growth rates, then we researchers have much work ahead of us.

References

- Gollier C (2012) Pricing the planet’s future: the economics of discounting in an uncertain world. Princeton University Press, Princeton (forthcoming October 2012)
- Gollier C, Weitzman ML (2010) How should the distant future be discounted when discount rates are uncertain? *Econ Lett* 107(3):350–353
- Groom B, Koundouri P, Panopoulou E, Pantelidis T (2007) Discounting the distant future: how much does model selection affect the certainty equivalent rate? *J Appl Econom* 22(3):641–656
- Maddison A (2001) The world economy: a millennial perspective. OECD Publishing
- Newell RG, Pizer WA (2003) Discounting the distant future: how much do uncertain rates increase valuations? *J Environ Econ Manag* 46:519–529