AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

Discounting the Distant Future: How Much Do Uncertain Rates Increase Valuations?

Richard Newell and William Pizer

Discussion Paper 00-45

Original version: July 19, 2000

Revised: October 24, 2000, May 14, 2001

Resources for the Future
1616 P Street, NW
Washington, D.C. 20036

Telephone: 202-328-5000

Fax: 202-939-3460

Internet: <http://www.rff.org>

© 2000 Resources for the Future. All rights reserved. No portion of this paper may be reproduced without permission of the authors.

Discussion papers are research materials circulated by their authors for purposes of information and discussion. They have not necessarily undergone formal peer review or editorial treatment.

Discounting the Distant Future: How Much Do Uncertain Rates Increase Valuations?

Richard Newell and William Pizer

Abstract

Costs and benefits in the distant future—such as those associated with global warming, long-lived infrastructure, hazardous and radioactive waste, and biodiversity—often have little value today when measured with conventional discount rates. We demonstrate that when the future path of this conventional rate is uncertain and persistent (i.e., highly correlated over time), the distant future should be discounted at lower rates than suggested by the current rate. We then use two centuries of data on U.S. interest rates to quantify this effect. Using both random walk and mean-reverting models, we compute the certainty-equivalent rate—that is, the single discount rate that summarizes the effect of uncertainty and measures the appropriate forward rate of discount in the future. Using the random walk model, which we consider more compelling, we find that the certainty-equivalent rate falls from 4%, to 2% after 100 years, 1% after 200 years, and 0.5% after 300 years. If we use these rates to value consequences at horizons of 400 years, the discounted value increases by a factor of over 40,000 relative to conventional discounting. Applying the random walk model to the consequences of climate change, we find that inclusion of discount rate uncertainty almost doubles the expected present value of mitigation benefits.

Key Words: Discounting, uncertainty, interest rate forecasting, climate policy, intergenerational equity

JEL Classification Numbers: D90, E47, C53, H43, Q28

Contents

1. Introduction.....	1
2. A Model of Discounting with Uncertain Rates.....	6
2.1 The Model.....	6
2.2 Implications of the Model.....	9
3. Estimation of Interest Rate Behavior.....	10
3.1 Data.....	13
3.2 Tests for Random Walks.....	14
3.3 Estimation Results	16
4. Forecasts and Application of Certainty-Equivalent Discount Rates	18
4.1 Certainty-Equivalent Discount Rates and Discount Factors.....	19
4.2 Relevance for Climate Change Policy Evaluation.....	22
5. Conclusions.....	25
Appendix.....	39
References.....	44

Discounting the Distant Future: How Much Do Uncertain Rates Increase Valuations?

Richard Newell and William Pizer*

1. Introduction

Implicit in any long-term cost-benefit analysis is the idea that costs and benefits can be compared across long periods of time using appropriate discount rates. Yet positive discount rates lead us to place very little weight on events in the distant future, such as potential calamities arising from global warming. For example, a dollar invested now yields \$51 after 100 years if the risk-free market return is 4%. Conversely, a promise to pay someone \$1 in 100 years—with complete certainty—is worth only two cents today at a 4% rate of discount. And a promise to pay someone (or rather, his or her descendants) \$1 in 200 years is worth only 4/100 of one cent today.

To bankers, financiers, or economists—people trained in the art of geometric discounting—this is a dull result. From an intuitive perspective, however, it is a bit more surprising. Furthermore, individuals consistently demonstrate the use of a declining discount rate in the future (Ainslie 1991). This effect can be particularly evident when valuations relate to an individual's own lifetime versus future generations (Cropper, Aydede, and Portney 1994). Use of a declining rate of discount is frequently referred to as hyperbolic discounting (in contrast to

* Newell and Pizer are Fellows at Resources for the Future. We thank Michael Batz for research assistance and Andrew Metrick and participants in seminars at the 2000 NBER Summer Institute and 2000 American Economic Association meetings for useful comments on previous versions of the paper. The research was supported in part by the Pew Center on Global Climate Change; such support does not imply agreement with the views expressed in the paper.

conventional, geometric discounting). However, the use of a deterministic declining rate, though consistent with individual preferences, produces time-inconsistent decisions.¹

An alternative to using continuously declining rates is to simply use a lower rate for long-term projects. Notable economists such as Ramsey (1928), for example, have argued that applying a positive rate of pure time preference to discount values across generations is “ethically indefensible.” More recently, Arrow et al. (1996) describe normative arguments for lower future discount rates under the rubric of a “prescriptive” approach to discounting, in contrast to a “descriptive” approach, which relies fully on historical market rates of return to measure discount rates. In practice, policymakers have, in some cases, begun applying lower discount rates to long-term, intergenerational projects (Bazerlon and Smetters 1999). Unfortunately, this approach comes close to causing the same time-consistency problems as long-term projects in the present become near-term projects in the future.

In contrast, the approach taken here fits squarely within the standard framework of geometric discounting based on market-revealed rates. The only distinction with most applications of geometric discounting is that we explicitly acknowledge that the discount rate itself is *uncertain*. As a direct consequence of this discount rate uncertainty, there is an increase in the expected value of future payoffs.² This underappreciated result is a straightforward implication of Jensen’s inequality. Because discounted values are a convex function of the discount rate, the expected discounted value will be greater than the discounted value computed

¹ For example, suppose we decide to use a 5% rate for 100 years and 0% afterward. From our perspective in the year 2000, a choice that trades a \$1 loss in 2150 for a \$2 gain in 2200 is desirable, because there is no discounting between these periods. After 2050, this choice begins to look worse and worse, as the interval between 2150 and 2200 begins to be discounted. Eventually, a utility-maximizing individual will want to reverse his or her initial decision, instead choosing the \$1 in 2150. Recognizing that this will assuredly happen represents a time inconsistency. That is, the mere passage of time will make the individual want to change his or her choice. For further discussion, see Cropper and Laibson (1999) and Heal (1998).

² If there is an opportunity to postpone a project, a second “option value” effect of interest rate uncertainty arises and creates a value of waiting to see whether interest rates rise or fall (Ingersoll and Ross 1992; Dixit and Pindyck 1994).

using an average rate. Put differently, the variable over which we should take expectations is not the discount *rate* (r), as is typically done, but rather the discount *factor* (e^{-rt}), which enters expectations linearly (Weitzman 1998).

Suppose, for example, that we are evaluating a project that yields \$1,000 in benefits 200 years in the future. The rate is thought to be 4% on average but is uncertain: it could be 1% or 7% with equal probability. If we simply use the average rate of 4% for discounting, then the present value benefit of this project ($\$1000e^{-0.04 \times 200}$) is 34 cents. On the other hand, if we take expectations properly, we find that the expected value of the project's benefits are now seen to be \$68, which is *200 times higher* than when we improperly used average rates ($0.5(\$1000e^{-0.01 \times 200}) + 0.5(\$1000e^{-0.07 \times 200}) = \68).

As shown by Weitzman (1998), a striking corollary of this result is that the effective “certainty-equivalent” rate of discount (corresponding to the expected discount factor) will decline over time. And this decline in the effective rate is especially dramatic as t becomes large. In fact, in the limit as t approaches infinity, Weitzman finds that the effective rate will decline to the *minimum possible* discount rate when a fixed but uncertain rate persists forever. Intuitively, the only relevant scenario in the limit is the one with the lowest possible interest rate, because all other possible higher interest rates have been rendered insignificant by comparison through the power of compounding over time. The same intuition holds in the intervening years—lower potential rates dominate as one moves further into the future because higher rates receive less and less weight as they are discounted away. As illustrated above, that result has potentially huge implications for the valuation of benefits in the distant future—such as those associated with mitigation of climate change, long-lived infrastructure, reduction of hazardous and

radioactive waste, and biodiversity—benefits that are discounted to a pittance when the discount rate is treated as if it is exactly known.³

It turns out that a crucial condition underlying the above results is that the discount rate is not only uncertain but also highly persistent. In the simple example above (and in Weitzman 2001, described below) the true but unknown rate is assumed to be fixed forever. Uncertainty without persistence will have little consequence for the effective rate if a high rate in one period is likely to be offset by low rates in subsequent periods. For this effect to have real punch, our expectation must be that periods of low rates will tend to be followed by more periods of low rates; the same goes for high rates. This is one of the key questions we will look to historical data to answer.

In order to quantify discount rate uncertainty, Weitzman (2001) conducts an email survey of more than 2,000 economists, asking them to *state* their “professionally considered gut feeling” about the appropriate real discount rate for valuing environmental projects. Here, uncertainty represents a *current* lack of consensus about the correct discount rate for all *future* time periods.⁴ We take a different approach: We assume there is a reasonable consensus about the correct discount rate today based on market rates, but that this rate is likely to change over long periods of time (as partially evidenced by the yield curve). We further assume that historical patterns of interest rate change *reveal* the likely patterns of change in the future. Two important features therefore distinguish our effort from previous work. First, we use market data to quantify uncertainty about future discount rates. Second, future beliefs about the appropriate discount

³ In a concrete application to climate policy, Pizer (1999) shows through simulations that uncertainty about future discount rates leads to the use of lower-than-average effective rates.

⁴ The difference of opinion expressed by the surveyed economists could be explained by different assumptions of each respondent about the incidence of environmental costs and benefits on consumption versus investment, taxes, and inflation—even though the survey twice specifies that this discount rate will be applied to “expected-consumption-equivalent real dollars.” In an unrelated survey of economists, Ballard and Fullerton (1992) find that conventional economic beliefs often supercede the specific features of a question presented for quick response.

rate evolve as the market evolves, meaning that any desire to revise a choice made in the past reflects the process of learning, rather than time-inconsistent behavior.⁵

We find significant empirical evidence that historical rates are indeed uncertain and persistent. Looking 400 years hence, the certainty-equivalent rate falls below 1%. However, the particular path of certainty-equivalent rates crucially depends on whether we believe that interest rates are a random walk or are mean reverting—a determination that is ambiguous in our data where the point estimate is 0.98 for the largest autoregressive root. Under the random walk assumption, which we find more compelling, the certainty equivalent rate falls from 4%, to 2% after 100 years, 1% after 200 years, and 0.5% after 300 years. In contrast, a mean-reverting model indicates certainty-equivalent rates that remain above 3% for the next 200 years, falling to 1% after 400 years. After 400 years, the cumulative effect of these rates is to raise valuations by a factor of more than 40,000 under the random walk model (and a factor of 130 under the mean-reverting model).

When we apply these rates to a real problem—the potential path of future damages due to current carbon dioxide emissions—we find that uncertainty almost doubles the present value of those damages based on the random walk model, while modestly raising the value by 14% based on the mean-reverting model. We find the random walk model more credible because the mean rate over the distant past is far less informative than the recent past when we forecast at any horizon in the future. We explore this distinction in detail and discuss ways that both approaches can be combined.

⁵ In an uncertain world there is always the possibility that *ex ante* good decisions turn out to be regrettable *ex post*, once nature has revealed herself—much like the purchase of insurance seems wasteful once the risk has uneventfully passed. This stands in contrast to time-inconsistent behavior where we know *with certainty* that our choice now opposes our choice in the future.

2. A Model of Discounting with Uncertain Rates

There are two natural approaches to modeling interest rate behavior. The first is a structural approach, which views interest rates as an endogenous outcome of a dynamic general equilibrium model. For example, the interest rate equilibrates demand for new capital and the supply of savings in a Ramsey growth model. Stochastic versions of this model, with either single or multiple technology trends, have been well studied (Brock and Mirman 1972; Cox, Ingersoll, and Ross 1985). In these types of models, exogenous and stochastic variables describing technology and growth determine the long-run behavior of interest rates.

The second approach models the time-series behavior of interest rates directly in a reduced-form manner. For example, we could specify that interest rates follow a first-order autoregressive process. This approach does not directly model the underlying determinants of the interest rate or impose structural restrictions. However, the reduced-form approach can be more transparent than a structural stochastic growth model. More importantly, the apparent strength of a structural model—the representation of underlying economic relations—can turn out to be a weakness because such models are typically based on restrictive, simplifying assumptions. In particular, over long horizons it is questionable whether we should assume that simplistic structural relations will continue to hold in the same way. A flexible reduced-form model, on the other hand, can have better predictive power for forecasting interest rates and modeling their inherent variability.

2.1 The Model

We begin by specifying the following stochastic model of interest rate behavior. The interest rate in period t , r_t , is uncertain, and its uncertainty has both a permanent component η , and a more fleeting yet persistent component ε_t :

$$r_t = \eta + \varepsilon_t. \quad (1)$$

The permanent component η has a normal distribution with mean $\bar{\eta}$ and variance σ_η^2 . The other random component, ε_t , is an autocorrelated mean-zero shock to the discount rate governed by:

$$\varepsilon_t = \rho \cdot \varepsilon_{t-1} + \xi_t, \quad (2)$$

where the correlation parameter ρ describes the persistence of deviations from the mean rate, and ξ_t is an independent identically distributed (iid), mean-zero, normally distributed random variable with variance σ_ξ^2 . The two uncertain components η and ε_t are assumed to be independent. Thus, the mean of r_t is $\bar{\eta}$.

To summarize, the interest rate has some mean value $\bar{\eta}$, which is itself uncertain, as well as persistent deviations from this mean rate. A value of ρ near one means that the interest rate can persistently deviate from the mean rate, staying consistently above or below it for many periods. In this case, the best guess of next year's rate is about equal to last year's rate. A value of ρ near zero, in contrast, means that a period of abnormally high interest rates may be followed by rates above or below the mean with equal probability, so that the next year's expected rate will be about equal to the mean rate. As we noted earlier, allowing for this persistence is essential. Uncertainty without persistence will have little consequence for discounting the distant future.

Discounting future consequences in period t back to the present is typically computed as the discount factor P_t , where

$$P_t = \exp\left(-\sum_{s=1}^t r_s\right).$$

Because r is stochastic, the *expected* discounted value of a dollar delivered after t years (i.e., the certainty-equivalent discount factor) will be

$$E[P_t] = E\left[\exp\left(-\sum_{s=1}^t r_s\right)\right]. \quad (3)$$

Following Weitzman (1998), we define the corresponding certainty-equivalent rate for discounting between adjacent periods at time t as equal to the rate of change of the expected discount factor:

$$\tilde{r}_t = -\frac{dE[P_t]/dt}{E[P_t]}. \quad (4)$$

Note that $\tilde{r}_t$ is the instantaneous period-to-period rate at time t in the future—not an average rate for discounting between period t and the present.

We evaluate Equation (3) by first using Equation (1) to separate it into two parts depending on the two components of discount rate uncertainty:

$$E[P_t] = E[\exp(-\eta t)] \cdot E\left[\exp\left(-\sum_{s=1}^t \varepsilon_s\right)\right]. \quad (5)$$

One can show that the first part of Equation (5) is equal to

$$E[\exp(-\eta t)] = \exp\left(-\bar{\eta}t + \frac{t^2 \sigma_\eta^2}{2}\right), \quad (6)$$

and that the second part is given by

$$E\left[\exp\left(-\sum_{s=1}^t \varepsilon_s\right)\right] = \exp\left(\frac{\sigma_\xi^2}{2(1-\rho)^2} \left(t - \frac{2(\rho - \rho^{t+1})}{1-\rho} + \frac{\rho^2 - \rho^{2t+2}}{1-\rho^2}\right)\right), \quad (7)$$

assuming $|\rho| < 1$. The limit expression for $\rho = 1$ is also finite. Finally, substituting Equations (6) and (7) into (5) we find that the expected discount factor for period t is

$$E[P_t] = \exp\left(-\bar{\eta}t + \frac{t^2 \sigma_\eta^2}{2}\right) \exp\left(\frac{\sigma_\xi^2}{2(1-\rho)^2} \left(t - \frac{2(\rho - \rho^{t+1})}{1-\rho} + \frac{\rho^2 - \rho^{2t+2}}{1-\rho^2}\right)\right). \quad (8)$$

The corresponding certainty-equivalent discount *rate* at time t is given by the rate of change of $E[P_t]$, as shown in Equation (4); it is equal to

$$\tilde{r}_t = \bar{\eta} - t\sigma_\eta^2 - \sigma_\xi^2 \Omega(\rho, t), \quad (9)$$

where $\Omega(\rho, t)$ is the effect of autocorrelation in the interest rate shocks and is given by

$$\Omega(\rho, t) = \frac{1 - \rho^2 + 2\log(\rho)\rho^{t+1}(1 + \rho - \rho^{t+1})}{2(1-\rho)^3(1+\rho)},$$

assuming $|\rho| < 1$.⁶

⁶ For the case of $\rho = 1$, $\Omega(\rho, t) = \frac{1}{12}(1 + 6t + 6t^2)$, while for $\rho = 0$, $\Omega(\rho, t) = \frac{1}{2}$.

2.2 Implications of the Model

Equation (9) gives the basic result of the model of discount rate uncertainty. The certainty-equivalent rate declines from the mean rate with increases in the forecast period (t), uncertainty in the mean rate (σ_η^2), uncertainty in deviations from the mean rate (σ_ξ^2), and the degree of persistence in those deviations (ρ). The second term of Equation (9) captures the influence of uncertainty in the mean interest rate, and the third term captures the effect of persistent deviations away from the mean. Positive correlation reduces the certainty-equivalent discount rate because $\Omega(\rho, t)$ will always be positive; this effect is increasing in ρ and t .

Thus, the certainty-equivalent discount rate will become lower the further one forecasts into the future, the greater the uncertainty in the mean interest rate, and the greater the variance and persistence (correlation) of shocks in the interest rate. As mentioned earlier, the degree of persistence in discount rate fluctuations turns out to be a critical component of what drives the certainty-equivalent rate down over time. This is illustrated in Figure 1, which shows the certainty-equivalent discount rate path based on a mean rate of $\bar{\eta} = 4\%$ —the average rate of return to government bonds over the past 200 years—and using parameter estimates from the simple autoregressive model given by Equations (1) and (2) using two centuries of market interest rate data (see Section 3 and Table 2 for estimation details). Specifically, we set $\sigma_\xi = 0.23\%$, $\sigma_\eta = 0.52\%$ and, $\rho = 0.96$. We also indicate the path when $\rho = 1$, as well as simulation results that treat ρ as uncertain.

Two important features are apparent in Figure 1. The first feature is the dramatic effect of different values of ρ near one. Although a value of 0.96 (or lower) leads to virtually no consequence over 200 years, a value of $\rho = 1$ leads to negative rates after only 130 years. In practice, the estimated mean and standard error of ρ do not rule out the possibility of ρ equal to or very close to one, as we discuss further below. The substantial decline of the simulated path (which treats ρ as uncertain) within 200 years demonstrates this real possibility.

The second feature is that the certainty-equivalent rate becomes negative when ρ is high. This is a straightforward feature of our model that fails to rule out negative rates based on the

specification given in Equation (1) with normally distributed errors (which we assume for tractability). Although one could imagine circumstances supporting negative expected rates, none are observed in the millennia of data covered by our primary source on interest rates, Homer and Sylla (1998). Further, common sense about the pure rate of time preference suggests that persistently negative rates are unlikely.⁷ This is a second issue we need to address if we want to model interest rate behavior realistically.

3. Estimation of Interest Rate Behavior

Our simple analytical model of interest rate behavior provides a useful guide to the various parameter combinations that make discounting with uncertain rates substantially different from discounting with certain rates. The model parameters are easily estimated from an autoregression of appropriate historical interest rate data. As we previously noted, however, there are several problems with the simple model as a realistic model of historical interest rate behavior.

First, the analytical model does not rule out the possibility of persistently negative discount rates even though rates below 1% have rarely been observed. Second, when the true value of ρ is near one, standard estimates of ρ will be biased downward in finite samples (and biased asymptotically when the true value equals one; see Chapter 17 of Hamilton 1994). Finally, the value of ρ , which is fixed in our analytical model, is estimated with considerable error using historical data. The range of this error is sufficient to lead to dramatically different estimates of the certainty-equivalent rate, as shown in Figure 1.

These problems can be overcome by modifying our earlier model. First, we transform the model in a way that prevents negative rates. In particular, we assume

$$r_t = \eta \exp(\varepsilon_t), \quad (10)$$

⁷ See chapter 7 of Mandler (1999).

or after taking logs,

$$\ln r_t = \ln \eta + \varepsilon_t, \quad (11)$$

with η now modeled as a log-normally distributed random variable⁸ with mean $\bar{\eta}$ and variance σ_η^2 , and where we generalize the autoregressive form for ε_t given by Equation (2) to allow for ε_t to depend on more than one past value:

$$\varepsilon_t = \rho_1 \varepsilon_{t-1} + \rho_2 \varepsilon_{t-2} \cdots + \rho_L \varepsilon_{t-L} + \xi_t, \quad (12)$$

where L is the number of lagged values included in the model and the ρ s are autoregressive coefficients.⁹ This is called the “mean-reverting” model because, with $\sum_s \rho_s < 1$, the series would eventually tend toward its long-run mean.

Second, we test the null hypothesis that historical interest rates are random walks (with $\sum_s \rho_s = 1$) and, if the test fails to reject, we estimate the random walk version of Equation (11). This is given by

$$\ln r_t = \ln r_0 + \varepsilon_t, \quad (13)$$

where now the constraint is imposed that $\sum_s \rho_s = 1$ for the autoregressive shocks ε_t in Equation (12). Note that r_0 in the random walk model replaces η in the mean-reverting model because interest rates are now modeled as an accumulation of permanent innovations from some initial rate (r_0), rather than deviations from a long-run mean (η).¹⁰

Finally, we use the estimated covariance matrix of the model parameters, including ρ , to randomly draw combinations of parameters as well as stochastic shocks ξ_t . We use these draws

⁸ A normality test on the data rejects a normal distribution for the data in levels but does not reject normality once the data have been logged.

⁹ Here and throughout, ρ refers to the single autoregressive parameter in a simple AR(1) process or the largest root in a more general AR(L) process.

¹⁰ Note that estimating the model conditional on initial observations (described in the next two paragraphs), $\ln r_t - \sum_s \rho_s \ln r_{t-s} = \xi_t$ under the random walk assumption and the likelihood is independent of r_0 . Setting $\varepsilon_0 = 0$, r_0 becomes the interest rate in period 0.

to simulate one hundred thousand alternative future paths for the discount rate. We then use these simulations to compute the certainty-equivalent rate numerically, rather than analytically.¹¹

We estimate the model parameters conditional on the initial observations, dropping those for which lagged values are not directly observed. With a single lagged value in the autoregression, this is equivalent to the Cochrane-Orcutt (1949) method; with more lagged values, this approach is referred to as conditional maximum likelihood (Hamilton 1994). We pick the number of lagged values in the autoregression based on the Schwarz-Bayes information criterion (Schwarz 1978).¹² For comparison, we also estimate the “simple autoregressive model” given by Equations (1) and (2), which includes only a first-order autoregressive lag and is not estimated on logged data.

In the remainder of this section we discuss the data used to estimate the models given by Equations (11) and (13), present the results of tests for a random walk, and then give estimation results for the different models. In Section 4 we use the estimation results to simulate the certainty-equivalent discount factors and rates, and then apply these to the case of climate change damages to gauge its ultimate significance for this important problem.

¹¹ With normally distributed errors, $E[\exp(r_t)] = r_0 \cdot \exp(\sigma_{\varepsilon_t}^2/2)$ for the random walk and $E[\exp(r_t)] = \bar{\eta} \cdot \exp(\sigma_{\varepsilon_t}^2/2) \cdot \exp(\sigma_{\eta}^2/2)$ for the mean-reverting models. This change in the mean rate confuses any attempt to show the effect of uncertainty relative to a constant discount rate and requires a small adjustment. In the random walk model, the factor $\exp(\sigma_{\varepsilon_t}^2/2)$ has a limiting value of $\exp(t \cdot \sigma_{\varepsilon}^2 / (2(1 + \sum_{l=2}^L (l-1)\rho_l)^2))$ where L is the number of lags. Therefore, a simple correction is to impose a deterministic trend of $-t \cdot \sigma_{\varepsilon}^2 / (2(1 + \sum_{l=2}^L (l-1)\rho_l)^2)$ in Equation (13) that exactly offsets this drift due to increasing variance. The correction for the mean-reverting model is more complex because the limiting value of $\exp(\sigma_{\varepsilon_t}^2/2)$ may not be achieved for many years (due to autoregressive roots arbitrarily close to unity). We impose a trend $-\sigma_{\eta}^2/2 \cdot t$ in Equation (11) to offset the first part of the drift. We then compute the expected variance $\sigma_{\varepsilon_t}^2$ each period based on the particular values of the ρ_s 's and σ_{ε} for each state of nature, and subtract that value (divided by two) from the simulated value of ε_t .

¹² The Schwarz-Bayes information criterion equals the log-likelihood minus a penalty, $k/2 \ln(n)$, where k is the number of parameters and n is the number of observations. This criterion reveals the (asymptotic) logged odds ratios for any collection of models, regardless of dimension.

3.1 Data

To estimate the model of interest rate behavior, we have compiled a series of market interest rates over the two-century period 1798 through 1999. The question of which rate to use is naturally contentious and has been discussed at length by numerous authors (see, for example, Portney and Weyant 1999; Arrow et al. 1996; Lind 1982). Plausible candidates include rates of return from bonds and other debt instruments, equities, or direct investment. Even within these broad categories there are a variety of possibilities with various risk levels, time horizons, and other characteristics. For the present analysis, we have focused on U.S. market interest rates for long-term, high-quality, government bonds (primarily U.S. Treasury bonds).¹³ This decision was based mainly on our desire to construct a very long time series of relatively low-risk rates of return and is described more thoroughly in an appendix.

We compiled a series of bond yields based on Homer and Sylla's (1998) monumental *History of Interest Rates* and used their assessments to determine the best instrument among high-quality, long-term government bonds available each year. Based on these nominal rates, we create a series of real interest rates by subtracting a measure of expected inflation, also described in the appendix. We then convert these rates to their continuously compounded equivalents. The final data set has 202 observations. We estimate the models using a three-year moving average of the real interest rate series to smooth very short-term fluctuations, which we are not interested in modeling here and which would otherwise mask the longer-term behavior of the data in which we are interested. The interest rate series are shown in Figure 2. We also note that our basic results are robust to the market interest rate used, the approach selected for inflation adjustment, and whether we smoothed the data before estimation; see the appendix for further detail.

¹³ Long-term bond yields in other industrialized countries have historically followed similar trends and general fluctuations, although the magnitude of fluctuations has at times differed substantially (Homer and Sylla 1998).

Although a full discussion of the causes of various trends and patterns of these rates is beyond the scope of this paper, a few points are worth noting. There appears to be a fairly steady downward trend from 1800 through at least 1940. Viewed on a grander historical scale, Figure 2 shows that this general trend has been apparent for at least the last millennium.¹⁴ Deviations from any mean or trend are persistent: for example, interest rates over 1860–1870 and 1910–1920 are noticeably higher than in adjacent periods. Finally, the rates observed during the early 1980s—even adjusted for inflation—are the highest since the early 1800s.

3.2 Tests for Random Walks

An important observation in our analytical results was that small differences in the estimate of the autoregressive parameter ρ have significant consequences for the certainty-equivalent discount rate. At the same time, an enormous literature in time-series econometrics has emphasized that ordinary inference is inappropriate when the data possess a “unit root” ($\rho = 1$)—that is, when the data follow a random walk (Dickey and Fuller 1979; Phillips 1987; Sims, Stock, and Watson 1990). Specifically, the estimates of ρ are asymptotically biased downward, with a standard 5% t -test of the $\rho = 1$ null hypothesis falsely rejecting 65% of the time (Nelson and Plosser 1982). When ρ is near but not equal to one, there will also be a finite-sample bias.

The standard approach to univariate modeling of time series is to specifically test the random walk hypothesis using one of several approaches.¹⁵ If the model rejects $\rho = 1$, the series is presumed to be stationary and standard methods can be applied directly to the data. If the

¹⁴ Using data from Homer and Sylla (1998), the figure shows the 50-year minimum and maximum yields on long-term debt for the United States (1800s and 1900s), England (1700s), and several European countries for the 11th to 17th centuries.

¹⁵ It is also possible to test the null hypothesis that the data are stationary (Kwiatkowski and et al. 1992). This approach is less common because the size of the test is unreliable when the data are autocorrelated, making it difficult to control Type I error.

model fails to reject, there is an unfortunate ambiguity because these tests have notoriously low power—that is, they are unable to distinguish among true values of ρ near unity. Thus, either of two approaches may be reasonable. One can impose a unit root and estimate the remaining parameters using standard methods. Or one can continue to estimate the unconstrained model but recognize the likely bias.

Table 1 presents our tests of the random walk hypothesis for the historical interest rate data in both levels and logs. The test consists of applying standard ordinary least squares to the model

$$x_t = \rho x_{t-1} + \alpha_0 + \sum_{l=1}^L \alpha_l \Delta x_{t-l} + \vartheta_t \quad (14)$$

where x_t is the time series being tested, Δx_t is the differenced time series, L is some number of lags included in the regression, and ϑ_t is an iid random disturbance. We choose the number of lags L of lagged first-differences by maximizing the Schwarz-Bayes information criterion over the choice of L (see footnote 12). The test statistic is the standard t -statistic computed as a test of $\rho = 1$. The 5% critical value is -3.43 if a trend is included in the regression and -2.88 otherwise.¹⁶

With a point estimate of $r = 0.976$ and standard error of 0.011 , we fail to reject the hypothesis of a random walk in the model with logged interest rates and no trend, as well as in other models based on unlogged interest rates and including trends. Because our results are therefore likely to be extremely sensitive to either a downward bias in ρ (if we estimate without imposing a unit root) or a specific assumption that $\rho = 1$, we estimate and present results for both the random walk and the mean-reverting models (Equations (13) and (11), respectively). We discuss the implications of these models below.¹⁷

¹⁶ This is the augmented Dickey-Fuller test presented in Dickey and Fuller (1979) with critical values in Fuller (1976) and Hamilton (1994).

¹⁷ It is possible using a Bayesian approach to compute the posterior probabilities for the mean-reverting and random walk models. Sims (1988) describes such an approach. Placing equal prior weight on each model (and a flat prior

It is useful to note that the ambiguity between mean-reverting and random walk models of interest rate behavior is not a recent phenomenon. Figure 3 shows the longer (but less reliable) history of interest rates over the last millennium. Periods of high rates are often followed by periods of lower rates—suggesting mean reversion. Yet over the entire millennium there have been highly persistent changes, from rates that averaged near 10% around 1000 A.D. to rates averaging less than 5% around 2000 A.D. In the end, we believe the only convincing way to decide between random walk and mean-reverting models is to ask whether—having observed unusually low rates for an extended period (say 30 to 40 years)—one anticipates a return to the longer-run average or a continuation of low rates.

3.3 Estimation Results

Table 2 presents estimates of the three models. We estimate each model with and without a time trend. The time trend is significant only in the mean-reverting model. However, this apparent significance is suspect because the coefficient will have a nonstandard distribution if logged interest rates follow a random walk, which we cannot reject. Nelson and Plosser (1982) show that a standard 5% *t*-test will falsely reject the null hypothesis of no trend 30% of the time even when the underlying series is truly a random walk without a trend. Further, it is unclear whether and how to extrapolate an observed historical trend into the future. Any extrapolation will be extremely sensitive to functional form; for example, a trend in the log of interest rates has very different implications than a trend in the level of interest rates. For these reasons, the remainder of our discussion focuses on the models without time trends.¹⁸

over the mean-reverting interval $\rho \in (0.5, 1.0)$) yields a posterior likelihood of 59% for the random walk; using Sims' preferred prior of a 20% weight on the random walk model yields a posterior likelihood of 26% for the random walk model. We discuss how one might combine the results of the two models using such probabilities at the end of the paper.

¹⁸ Note that we continue to include a trend correction in the random walk model so that the expected rate remains constant as described in footnote 11. This correcting trend is constrained to be equal to $-\sigma_{\xi}^2 / (2(1 + \rho_2 + 2\rho_3)^2)$,

All three models yield remarkably consistent results. The mean-reverting and simple unlogged models provide similar estimates of both the mean interest rate and its associated standard error. The random walk and mean-reverting models provide similar estimates of the autocorrelation parameters—not surprising, since we are unable to resolve whether a unit root exists. The largest “root” of the mean-reverting parameter estimates also equals 0.95, almost exactly the estimate in the simple, unlogged, first-order autoregressive model.¹⁹ Finally, the estimates of σ_{ξ}^2 are also consistent across these models.²⁰

Despite similar parameter estimates, the mean-reverting and random walk models paint starkly different pictures of the future. This is evident when we consider simple out-of-sample prediction intervals. For example, Figure 4 shows the 95% prediction intervals on forecasts for 1900–1950 using the two models estimated over the first half of the sample period (1799–1899). The mean-reverting model indicates that our prediction interval will remain virtually unchanged after 30 or 40 years. It also remains centered on the long-run mean (4% over 1799–1899). In contrast, the random walk model indicates that the prediction interval will continue widening and is centered on the last observation.²¹ This difference arises because deviations in the mean-reverting model are attenuated over time. New deviations occur, but because they do not build on

which exactly offsets the positive trend that would otherwise exist in the rate due to an increasing variance in the logged rate over time. Inclusion of this trend correction has no significant effect on the parameter estimates.

¹⁹ The polynomial expression $1 - 1.88L + 1.31L^2 - 0.40L^3$, based on the mean-reverting parameter estimates, can be factored into $(1 - 0.95L)(1 - 0.93L + 0.42L^2)$, revealing that the largest autoregressive root is 0.95.

²⁰ In the mean-reverting and random walk models (both estimated in logs), the stochastic innovations have an estimated variance of 0.0015, or a standard error of 0.04. Because differences in logs correspond to percent changes in levels, one would think that multiplying the mean interest rate of 3.52 by the standard error of the logged models (0.04) ought to roughly yield the standard error in the simple unlogged model (0.23). However, 0.04×3.52 equals 0.14. This difference is accounted for by the fact that the simple unlogged model includes only a single lag (to make it comparable to the analytical model). Specifying the unlogged model with three lags yields a direct estimate of the standard error of 0.16, close to the above calculation of 0.14.

²¹ Both models are “centered” in logarithms. The distribution becomes skewed when the predictions are exponentiated.

previous ones, the series tends toward its long-run mean. In contrast, deviations in the random walk model persist forever, and there is no tendency to revert to previous levels.

Comparing those prediction intervals to realized data, we find some impetus to favor the random walk model in our application. The realized interest rate over 1900–1950, for example, first appears to revert up to the estimated long-run mean in the mean-reverting model. Toward the end of the period, however, the interest rate lies below the 95% prediction interval for the mean-reverting model—and remains there for ten years without any sign of mean reversion. Similar patterns appear in virtually any out-of-sample forecast. This inconsistency between the mean-reverting forecasts and the realized interest rate is particularly troubling because we know that the lower range of possible interest rates ultimately determines the future certainty-equivalent rate. Because the random walk model does a better job of predicting this possibility, we find it more compelling for our application, even though evidence based on standard statistical tests is ambiguous.

4. Forecasts and Application of Certainty-Equivalent Discount Rates

Rather than using our analytical model to compute an exact certainty-equivalent discount rate for a particular value of ρ , we now construct a numerical approximation to the certainty-equivalent discount rate based on simulations that incorporate estimated uncertainty about ρ . This also allows us to handle the more complex models for the discount rate estimated in Section 3, which include higher-order lag structures and data in logs.

To construct the numerical approximation, we simulate 100,000 possible future discount rate paths for each model starting in 2000 and extending 400 years into the future. We begin with point estimates and a joint covariance matrix for the parameters given in Table 2. For each simulated discount rate path, we assume the parameters in Table 2 are jointly normal and draw

values for each parameter.²² We then draw values for the stochastic shocks ξ . We create the ε shocks by recursively defining ε based on Equation (12). Finally, we use the simulated values of ε to construct simulated discount rates for each of the three models based on Equations (1), (11), and (13).²³

After simulating 100,000 future paths for each model, we compute the expected discount factor $E[P_t]$ based on Equation (3). The certainty-equivalent discount rate is computed as the discrete approximation to Equation (4) given by $\tilde{r}_t = E[P_t]/E[P_{t+1}] - 1$. Table 3 presents our estimates of discount factors over the next four hundred years based on a 4% rate of return in 2000 and using our historical data on long-term government bonds to quantify interest rate uncertainty.²⁴ We report results for both the random walk model as well as the mean-reverting model. For comparison we present discount factors associated with a constant rate of 4%. Table 4 presents a sensitivity analysis for the alternative rates of 2% and 7%. Figure 5 shows the certainty-equivalent rates corresponding to the discount factors given in Table 3, based on an initial 4% rate.

4.1 Certainty-Equivalent Discount Rates and Discount Factors

The two important features to notice in Figure 5 are the effect of using a model in logs (Equation (11) versus (1)) and the effect of a random walk versus mean-reverting model

²² In simulations of the mean-reverting and simple unlogged models, we also replace any random draws that result in explosive autoregressive models; i.e., when $\sum \rho_t > 1$, which would unrealistically imply that current innovations ξ have an increasing rather than diminishing effect on future discount rates.

²³ As noted in footnote 11, we make a small adjustment to the simulated mean-reverting discount rates to correct for changes in the expected value. These corrections are not necessary in the random walk model, where a small trend fixes the problem.

²⁴ We begin with an initial rate of 4%, rather than the actual bond rate of about 3% in 2000, because short-term forecasts of the interest rate suggest it is likely to rise over the next few years—making a fair comparison with a constant interest rate impossible. A rate of 4% reflects the approximate 200-year average as well as the average over the past 20 years. It also falls close to the middle of the range of defensible consumption rates of interest (2–7%).

(Equations (13) versus (11)). The model in logs successfully avoids the possibility of negative interest rates in the future, eliminating negative certainty-equivalent rates. Eliminating the possibility of negative rates also slows the decline in future certainty-equivalent rates. Although the unlogged model generates a certainty-equivalent rate declining from 4% to 1% after 210 years, the comparable (mean-reverting) logged model requires almost 400 years to decline to 1%.

Because the exclusion of negative interest rates is relatively uncontroversial, the more interesting comparison is between the random walk and mean-reverting models. Both will decline toward zero because over the years, simulated discount rates have more time to wander closer to zero. In the mean-reverting model, however, the chance that discount rates will be persistent enough to wander very far is relatively small. In contrast, the random walk model assumes such persistence with certainty. This distinction can be seen in part in Figure 4, which shows that the estimated range of uncertainty (i.e., prediction interval) for the mean-reverting forecast is narrower than for the random walk model, implying a higher likelihood that the random walk model can more quickly drift close to zero.

The impact on certainty-equivalent discount rates is enormous. The random walk model implies rates that decline from 4%, to 2% after 100 years, down to 1% after 200 years, and further declining to 0.5% after about 300 years. Meanwhile, the mean-reverting model indicates that certainty-equivalent rates stay above 3% for 200 years. Over the next 200 years, the decline is similar to the random walk model, with rates hitting 2% after 300 years and 1% after 400 years.²⁵

These certainty-equivalent discount rates translate into dramatic differences in the valuation of future consequences. Table 3 shows the expected value today of \$100 delivered at

²⁵ Recall that the certainty-equivalent rate is the period-to-period rate at some time t in the future—not an average rate for discounting between period t and the present. To get a better sense of the effect of discount rate uncertainty on present values, we must look at expected discount factors, as we do below.

various points in the future (e.g., $100 \cdot E[P_t]$ in Equation (3)). The expected value is first evaluated using a constant discount rate of 4%—the average rate of the 200-year sample. We then evaluate the expected value based on simulated discount rates from both the mean-reverting and the random walk models.

After only 100 years, conventional discounting at 4% undervalues the future by a factor of 3 based on the random walk model of interest rate behavior. After 200 years, that factor rises to about 40. That is, conventional discounting values \$100 in the year 2200 at 4 cents. The random walk model values the same \$100 at \$1.54—about 40 times higher. Going further into the future, conventional discounting is off by a factor of over 40,000 after 400 years. The same dramatic effects occur with the mean-reverting model, but lagged by 100 to 200 years (a factor of 3 after 200 years, and a factor of over 40 after 360 years).

Table 4 presents an alternative comparison using initial interest rates of 2% and 7% for the random walk model—what you might think of as upper and lower bounds on the consumer rate of interest. We use the same assumptions about random disturbances estimated from data on government bond rates, but initialize the random walk at a different rate.²⁶ Again, we compute discount factors based on the corresponding constant rate as a benchmark. When you compare the ratio of random-walk to constant-rate discount factors, these valuations show that the *relative* effect of interest rate uncertainty (measured by this ratio) rises as the initial rate rises. The effect at a horizon of 400 years raises the valuation by a factor of 530 million based on a 7% rate. Meanwhile, the effect is a factor of a little over 100 based on a 2% rate (from the bottom line of Table 4). Intuitively, the effect must be smaller for low discount rates (e.g., 2%) because the range of possible lower rates (0-2%) is narrower.

²⁶ Note that because the logged interest rate follows a random walk, it implies that the disturbances to the (unlogged) interest rate are scaled by the magnitude of the interest rate—larger rates experience larger disturbances, and vice-versa for smaller initial rates. Casual observation of historic fluctuations supports such an assumption (see figures on pages 369, 394 and 424 of Homer and Sylla 1998).

This means that the difference between valuations using different initial rates is generally *smaller* when uncertainty about future rates is incorporated. Note that the ratio of discount factors based on the random walk model, but starting with initial rates of 2% versus 7%, is a factor of about 40 after four hundred years (see bottom line of Table 4: $3.83 \div 0.09 \approx 40$). Compare that to a factor of 200 million based on constant discount rates.²⁷ In other words, the choice of discount rate is less important when you consider the effect of uncertainty—though obviously a factor of forty is still substantial.

Thus, Table 3 and Table 4 provide a precise answer to the question posed by the title of this paper: “How much do uncertain rates increase valuations?” At horizons of several hundred years, the answer is, “Quite a lot.” After 400 years the random walk model, which we consider more convincing, increases the value by a factor of over 40,000 relative to conventional discounting at 4%. The mean-reverting model increases the value by a factor of 130.²⁸ Because the numbers in these tables represent expected values, one can alternatively combine the results based on the probability given to each model: simply weight the values by the appropriate probability (see footnote 17). For example, placing an equal probability on both models yields an increase in value of roughly 21,000 after 400 years. A weight of 25% on the random walk model yields a factor of roughly 11,000. As we will see below, these large increases in intertemporal prices can substantially alter the evaluation of policy consequences over long horizons.

4.2 Relevance for Climate Change Policy Evaluation

An important application of discounting the distant future is valuation of the consequences of climate change due to human activities, namely the burning of fossil fuels and

²⁷ Discounting at a constant rate of 7%, \$100 delivered 400 years in the future would be valued at 2×10^{-10} today, a fact obscured due to rounding in the bottom line of Table 4, column 6. Note that $0.04 \div 2 \times 10^{-10} = 2 \times 10^8$ or 200 million.

²⁸ After 500 years—which we believe stretches the credibility of the model—the factors are 160,000 and 600.

emission of carbon dioxide. Conventional analyses, using constant rates of 4% to 5%, tend to produce extremely low estimates of climate change damages (see, for example, Nordhaus 1994). These analyses recommend moderate if not marginal mitigation action. Other analyses, based on much lower discount rates, produce significantly higher climate damage estimates and recommend aggressive action (Cline 1992).

How does our approach weigh in? We can simulate Nordhaus's model to determine the marginal damages in future years from an extra ton of emissions in the year 2000. This profile of climate damages is shown in Figure 6; it depicts a sensible pattern of damages having a delayed, then an increasing, and finally a declining effect as emissions decay.²⁹ We can then apply alternative discount rate paths to estimate the present value of damages, PV , from a marginal ton of emissions today—and hence the marginal benefit of avoided emissions:

$$PV = \sum_{t=2000}^{2400} MD_t \exp\left(-\sum_{s=2000}^t \tilde{r}_s\right) = \sum_{t=2000}^{2400} MD_t E[P_t] \quad (15)$$

where MD_t is the marginal damage illustrated in Figure 6, $\tilde{r}_s$ are the various paths for the certainty-equivalent discount rate (e.g., shown in Figure 5), and $E[P_t]$ are the various discount factors summarized in Table 3 and Table 4.

The results of these calculations are shown in Table 5 based on the 4% rate that reflects the historic return to government bonds, as well as our high and low sensitivity calculations. For each rate, we report the valuation based on both the random walk and mean-reverting models, as well as for a constant discount rate, and indicate the relative effect of uncertainty—the ratio of the valuation under the two uncertainty models to the corresponding valuation with constant

²⁹ We construct this profile by simulating the model described in the Appendix to Nordhaus (1994) in the absence of any policy intervention. We then modify the model to emit an extra ton of emissions in 1995, the base year of the model, resimulate the model, and measure the impacts on consumption. In practice, because the effect of a single ton is beyond measurement, it is necessary to model the effect of an extra 100 million tons and divide by 100 million.

rates.³⁰ We first observe that the value associated with the constant discount rate of 4% is reasonably close to the value of \$5.29 reported in Table 5.7 of Nordhaus (1994).³¹

In contrast, the effect of discount rate uncertainty based on the random walk model is quite large—increasing the estimated benefits of mitigation by over 80% to more than \$10.³² The mean-reverting model yields a more modest increase (a 14% increase in present value, to about \$7). As noted above, one could combine these estimates by assigning probabilities to each of the two models and taking expectations. Equal weighting, for example, would yield an estimate of \$8.50.

As one would expect based on the discount factors from Table 4, the relative effect of uncertainty on present value of expected mitigation benefits is larger when the comparison involves higher discount rates. This reflects the greater opportunity for uncertainty to lower rates when the initial rate is higher (versus a low initial rate where the rate simply cannot go much lower). The effect of uncertainty is a 95% increase in discounted mitigation benefits with a 7% rate and a 56% increase with a 2% rate, based on the random walk model. The mean reverting model again yields a more modest 21% increase using the 7% rate, and a 7% increase using the 2% initial rate.

Note that while the dollar value of discounted climate benefits is sensitive to the magnitude of the benefits profile we have chosen for illustration (Figure 6), the proportional increase due to incorporation of the effect of discount rate uncertainty depends only on the general shape of the benefits profile. In addition, because we focus on a 400-year horizon, our

³⁰ For brevity, we did not report the mean-reverting discount factors based on initial rates of 2% and 7% in Table 2.

³¹ A constant discount rate of 4.15% delivers a net present value of exactly \$5.29. Note that Nordhaus does not use a constant rate based on his assumption of a growth slowdown.

³² At first glance, one might think based on Table 3 that the effect should have been even greater. However, even though benefit flows in the distant future may be severely undercounted by discounting that ignores uncertainty, the error is not as great for near-term benefit flows, and near-term flows still carry greater weight than distant flows in the overall calculation.

results are in some sense conservative; extending the horizon further introduces damages that are counted more heavily in the presence of uncertainty.³³ Applying the uncertainty-adjusted discount factors to other greenhouse gases with longer atmospheric residence (e.g., methane or sulphur hexafluoride), or to climate damage profiles that include catastrophic events or other permanent consequences (e.g., species loss), would also generate larger increases in discounted climate damages because the consequences would be more heavily concentrated in the future. In general, the greater the proportion of benefit flows occurring in the distant future, the greater the error introduced through discounting that ignores uncertainty in the discount rate itself. Other applications can be explored by applying the random walk and mean-reverting discount factors (or a weighted combination thereof) from Table 3 and Table 4.

5. Conclusions

Properly discounting over long horizons requires that one consider the uncertainty surrounding future discount rates. Uncertainty coupled with persistence implies that appropriate future certainty-equivalent rates will decline toward the lower bound of possible future discount rates. Quantifying this effect requires not only a forecast of future rates but also estimates of their variance and autocorrelation.

Using 200 years of historical data on high-grade, long-term government bonds in the United States (primarily U.S. Treasury bonds), we estimated several models of interest rate behavior. We found that the sampling error associated with our estimates of interest rate autocorrelation makes it difficult to draw conclusions based on our simple analytical model, which treats the autocorrelation as known. In fact, this sampling error is large enough that we are unable to reject the hypothesis that interest rates in either levels or logs follow a random walk.

³³ We have estimated that this understatement might be about 4% based on an extrapolation of the rate of decline in discounted carbon benefits at the end of our 400-year horizon.

Turning to simulations that both treat the autocorrelation as uncertain and allow a log-normal model for interest rates (which avoids negative values), we find significant effects associated with uncertainty in the long run, regardless of the underlying model. Although the average long-term real rate of return on government bonds is around 4%, the appropriate rate to discount the distant future (more than 400 years) is around 0.5% based on a random walk model, and 1% based on a mean-reverting model. Over horizons of less than 400 years, the random walk model suggests declines to much lower rates: 2% after 100 years, 1% after 200 years, and 0.5% after 300 years. Certainty-equivalent rates for the mean-reverting model, on the other hand, remain above 3% for next 200 years, declining to 2% after 300 years and 1% after 400 years.

At horizons of several hundred years, we find that uncertainty considerably raises the valuation of future consequences. After 400 years, the random walk model (which we consider more compelling) increases discounted values by a factor of more than 40,000; the mean-reverting model increases the value by a factor of 130. When applied to the path of future damages from current carbon dioxide emissions, uncertainty nearly doubles the present value of those damages based on the random walk model while modestly raising the value by 14% based on the mean-reverting model.

Although the 200 years of data used in our analysis does not give us statistical evidence to unambiguously favor either the mean-reverting or the random walk models, our opinion is that the latter makes more sense. Over long horizons, such as millennia, we see persistent changes in observed interest rates. Moreover, after a long period of relatively low rates, we think that such rates are more likely to persist than return to a long-run average. These propositions—that changes in rates are persistent over hundreds if not thousands of years, and that low rates for many recent years suggest low rates in the future—are consistent with the random walk model. At a practical level, the random walk model generates more sensible prediction intervals than the mean-reverting model.

This in turn suggests that for climate change mitigation the expected marginal benefits could be understated by a factor of 2 in analyses that ignore uncertainty in the discount rate itself. We would expect similar understatements in the analysis of other problems with very long time

horizons, including long-lived infrastructure projects, hazardous and radioactive waste disposal, and biodiversity preservation.

Table 1: Tests of Random Walk Hypothesis

	Log of interest rate		Level of interest rate	
	No trend	Trend included	No trend	Trend included
$\hat{\rho} - 1$	-0.024 (0.011)	-0.048 (0.016)	-0.024 (0.010)	-0.045 (0.015)
t -statistic	-2.255	-3.098	-2.300	-3.339
Lags included**	2	2	5	2

*Standard errors are shown in parentheses. The tests reject the null hypothesis of a random walk when the t -statistic falls below -2.88 for models without a trend and below -3.43 for models with a trend.

**Two lags in this specification correspond to three lags in the estimated model.

Table 2: Estimation Results

$$\ln r_t = \ln \eta + \varepsilon_t \text{ and } \varepsilon_t = \rho_1 \varepsilon_{t-1} + \rho_2 \varepsilon_{t-2} \cdots + \rho_L \varepsilon_{t-L} + \xi_t$$

	Random walk model ($\sum \rho_s = 1$)		Mean-reverting model		Simple unlogged model ^e	
mean rate ($\bar{\eta}$)			3.69 ^{*d}	3.95 ^{*d}	3.52 [*]	3.92 [*]
std error (σ_{η})			0.45	0.23	0.52	0.31
autoregressive coefficients ^a						
ρ_1	1.92* (0.06)	1.92* (0.06)	1.88* (0.07)	1.85* (0.07)	0.96* (0.01)	0.94* (0.02)
ρ_2	-1.34* (0.12)	-1.34* (0.12)	-1.31* (0.12)	-1.26* (0.12)		
ρ_3	0.43* (0.07)	0.43* (0.07)	0.40* (0.07)	0.36* (0.07)		
trend	^b	-0.0037 ^c (0.0055)		-0.0033 ^{*c} (0.0010)		-0.010 ^c (0.005)
σ_{ξ}^2	0.0015* (0.0002)	0.0015* (0.0002)	0.0015* (0.0002)	0.0015* (0.0002)	0.0522* (0.0052)	0.0517* (0.0052)

* Significant at the 5% level.

^a Number of autoregressive terms chosen using the Schwarz-Bayes information criterion.

^b As noted in footnote 18, the parameter estimates are unaffected by either constraining the trend to be zero or constraining it to equal $-\sigma_{\xi}^2 / (2(1 + \rho_2 + 2\rho_3)^2) = -0.003$, which exactly offsets the positive trend that would otherwise exist in the actual (unlogged) rate due to an increasing variance in the logged rate over time.

^c Indicates a linear time trend estimated alongside ε_t in each model.

^d The mean rates in this table were constructed, for simulation purposes, to reflect a continuously compounded rate. To convert to simple annual rates, simply compute $100 \times (\exp(\eta/100) - 1)$, e.g., $100 \times (\exp(3.69/100) - 1) = 3.76$.

^e Simple unlogged model is $r_t = \eta + \varepsilon_t$, where $\varepsilon_t = \rho \varepsilon_{t-1} + \xi_t$.

Table 3: Value Today of \$100 in the Future

Years in future	Discount rate model			Value relative to constant discounting	
	Constant 4% rate	Mean reverting	Random walk	Mean reverting	Random walk
0	\$100.00	\$100.00	\$100.00	1	1
20	45.64	46.17	46.24	1	1
40	20.83	21.90	22.88	1	1
60	9.51	10.61	12.54	1	1
80	4.34	5.23	7.63	1	2
100	1.98	2.61	5.09	1	3
120	0.90	1.33	3.64	1	4
140	0.41	0.68	2.77	2	7
160	0.19	0.36	2.20	2	12
180	0.09	0.19	1.81	2	21
200	0.04	0.10	1.54	3	39
220	0.02	0.06	1.33	3	75
240	0.01	0.03	1.18	4	145
260	0.00	0.02	1.06	5	285
280	0.00	0.01	0.97	7	568
300	0.00	0.01	0.89	11	1,147
320	0.00	0.01	0.83	16	2,336
340	0.00	0.00	0.78	26	4,796
360	0.00	0.00	0.73	43	9,915
380	0.00	0.00	0.69	74	20,618
400	0.00	0.00	0.66	131	43,102

Table 4. Sensitivity of Valuation to Initial Interest Rate

Years in future	2% initial rate			7% initial rate		
	random walk	constant rate	ratio	random walk	constant rate	ratio
0	100.00	100.00	1	100.00	100.00	1
20	67.54	67.30	1	26.89	25.84	1
40	46.48	45.29	1	8.67	6.68	1
60	33.05	30.48	1	3.52	1.73	2
80	24.42	20.51	1	1.75	0.45	4
100	18.76	13.80	1	1.02	0.12	9
120	14.93	9.29	2	0.67	0.03	22
140	12.25	6.25	2	0.47	0.01	62
160	10.32	4.21	2	0.36	0.00	181
180	8.89	2.83	3	0.29	0.00	557
200	7.81	1.91	4	0.24	0.00	1,778
220	6.97	1.28	5	0.20	0.00	5,851
240	6.30	0.86	7	0.17	0.00	19,726
260	5.77	0.58	10	0.16	0.00	67,829
280	5.33	0.39	14	0.14	0.00	236,788
300	4.97	0.26	19	0.13	0.00	837,153
320	4.66	0.18	26	0.12	0.00	2,992,921
340	4.40	0.12	37	0.11	0.00	10,804,932
360	4.18	0.08	52	0.10	0.00	39,298,213
380	3.99	0.05	74	0.10	0.00	143,866,569
400	3.83	0.04	105	0.09	0.00	529,656,724

Table 5: Expected Discounted Value of Climate Mitigation Benefits (per ton Carbon)

		Benefits from 1 ton of carbon mitigation	Relative to constant rate
Government bond rate (4%)	Constant 4% rate	\$5.74	—
	Random walk model	\$10.44	+82%
	Mean-reverting model	\$6.52	+14%
2% rate	Constant 2% rate	\$21.73	—
	Random walk model	\$33.84	+56%
	Mean-reverting model	\$23.32	+7%
7% rate	Constant 7% rate	\$1.48	—
	Random walk model	\$2.88	+95%
	Mean-reverting model	\$1.79	+21%

Figure 1. Importance of Persistence for Certainty-Equivalent Discount Rates

Figure 2. Market Interest Rate on U.S. Long-Term Government Bonds (1798–1999)

Figure 3. A Millennium of Market Interest Rates on Long-term Debt (50-year Intervals)

Figure 4. Forecasts of Interest Rate Uncertainty (1900–1950)

Figure 5: Forecasts of Certainty-Equivalent Discount Rates

Figure 6: Profile of Marginal Climate Damages from 1 Ton of Carbon Emissions in 2000

Appendix

Several concerns guided our choice of data. First, we sought a single instrument for which hundreds of years' worth of data were available. The empirical basis and plausibility of forecasts over long horizons depend on the use of historical data of a comparable duration. Because we hope to provide guidance for discounting over hundreds of years into the future, we need data over hundreds of years in the past. Although it is possible to splice together interest data from many different instruments, this can be confusing and introduces subjective factors related to the choice of instrument at various times.

Second, we wanted the level of investment risk associated with the selected financial instrument to be both low and relatively constant over time. Rates of return typically include a risk premium depending on the risk associated with the investment. Investments that may have an unexpectedly low or even negative return must, on average, provide a higher return to attract investors. This premium depends on the riskiness of the investment, the risk aversion of individuals, and the degree to which the particular investment risk is correlated with other risks. When we evaluate social projects, the riskiness of the investment may not be comparable to commonly observed market risks, and we may want to value that risk differently. Therefore, it is useful to separate the issue of risk from the issue of discounting by focusing on rates of return on high-quality, low-risk investments.³⁴

The only investment instrument that nearly satisfies those criteria is long-term U.S. Treasury bonds, backed by the faith and credit of the U.S. government. Unfortunately, the market for long-term federal debt was very thin or nonexistent between 1829 and 1843 as the

³⁴ Thus, a standard approach to policy analysis under uncertainty is to convert uncertain cost and benefit flows to certainty equivalents and then discount these flows using a risk-free (or at least low-risk) rate of return (see, for example, Arrow et al. 1996; Lind 1982). Our approach takes this framework one step further by recognizing that uncertain discount rates in the future must also be converted to a certainty-equivalent path.

government paid down the debt, and the market for U.S. Treasury bonds was distorted by banking reform between 1865 and 1920.³⁵ During these periods we make use of high-quality, long-term municipal bonds. The data are from Tables 38, 45, 48, 51, and 84 of Homer and Sylla (1998), with recent data for 1996–1999 from the Federal Reserve (2000).

Our focus on long-term bonds raises some issues. The reported yield on a long-term bond measures the annual rate of return if the bond is held until it matures—essentially the average interest rate over the life of the bond.³⁶ Although treating this reported yield like an annual return is not quite right, it turns out not to matter much for our analysis. For example, suppose we use the reported yield on ten-year bonds over 1981–2000 to compute the return over that period. The reported yield in 1981 is actually an average of expected interest rates over 1981–1990; the yield in 1982, an average of 1982–1991; and so on. By treating the reported yields like annual interest rates, we end up underweighting the interest rate in early years and, by including the bond yields in the final years, incorporating interest rates beyond the period we care about. The important observation, however, is that this error will be small over horizons of hundreds of years because it is limited to a misweighting of the first few and last few rates.³⁷

A second issue arises when we recognize that the reported yield on a ten-year bond is not exactly the expected rate on a sequence of ten one-year bonds. Long-term bond rates have typically traded at a premium (of about 1%) relative to short-term rates since the 1950s.³⁸ To the

³⁵ From about 1865 until 1920, U.S. government bond yields were distorted by major changes in banking policies. To establish a single national currency, banks were required to hold government bonds in exchange for the right to circulate government notes. This was followed by a period when government surpluses led to repurchase of outstanding debt and a shortage of government bonds.

³⁶ Specifically, it is the constant discount rate that equates the purchase price today with the stream of payments promised by the bond in the future.

³⁷ A second consequence for our econometrics is that long-term bonds effectively smooth the series, removing short-term fluctuations. This is, in fact, desirable, because we are interested in long-term, persistent movements in interest rates.

³⁸ The premium has been justified in terms of borrowers' preference for longer-term loans and lenders' preference for shorter-term investments. Campbell (1995) provides an excellent discussion of these issues.

extent that this premium remains positive over time, it suggests that our results may be conservative and undervalue the future. Before 1930, however, and especially during the 1800s, reported short-term rates were much higher than long-term rates.³⁹ In this case, our use of long-term rates is appropriate because it avoids any premium placed on short-term loans due to either risk preferences or market failures.

To check the robustness of our results to the construction of our choice of data series, we also estimate our models using series that (i) employ the rate on one-year Treasury bills starting in 1956;⁴⁰ (ii) use corporate rather than government bonds starting in 1888;⁴¹ (iii) adjust for inflation using a number of alternative measures of inflation expectations; and (iv) are unsmoothed. We find that our basic results are robust to the market interest rate used, the approach used for inflation adjustment, and whether the data were smoothed before estimation.

Returning now to the issue of inflation adjustment, we begin by noting that the creation of an index of expected inflation for the purpose of adjusting nominal interest rates is a necessarily subjective process. We base the adjustment on the surveyed expectations of professional economists and several historical facts that are described in more detail in Homer and Sylla (1998) and Spiro (1989). Before the movement from the gold standard to a paper money standard at the end of World War II, the average annual rate of inflation was very close to zero. There were some years of high inflation, usually associated with wars or other temporary disruptions, but they were always offset by subsequent years of deflation. In fact, price levels as measured by the Consumer Price Index were the same in 1940 as they were in 1800. As a result, financial market participants anticipated that the future rate of inflation would be about zero. In

³⁹ This reflected the lack of well-formed markets for high-quality, short-term debt; see Homer and Sylla (1998).

⁴⁰ We use the one-year Treasury bill rate over the period 1956–1999 using data from the Federal Reserve (2000) and Table 52 of Homer and Sylla (1998).

⁴¹ We use the long-term corporate bond rate over the period 1888–1999 using data from Tables 38, 45, 47, 50, and 84 of Homer and Sylla (1998).

individual years when inflation rose sharply, there was very little effect on long-term bond yields because any such inflation was viewed as temporary and likely to be offset by deflation in subsequent years.

In contrast, prices have risen every year since 1955, eventually giving rise to an expectation that inflation would be a persistent phenomenon. But expectations of persistent inflation were slow to develop and slow to evolve. Even after a period of substantial inflation in the 1940s, expectations were that the deflation of the 1930s would resume. Indeed, it was not until the Eisenhower and Kennedy years that inflationary concerns even began to enter into the public consciousness. There is therefore no consistent effect of inflation on long-term bond yields evident until the mid-1950s.

Thus, starting in 1955, we adjust the nominal market interest rates for inflation by subtracting a moving average of the expected inflation rate of the CPI over the previous ten years, as measured by the Livingston Survey of professional economists.⁴² For the purposes of adjusting long-term bond yields, a ten-year lagged moving average provides an appropriate measure of expected inflation because it captures the basic trends yet is not subject to extreme short-term volatility. The resultant measure of expected inflation follows the basic trends in nominal bond yields over the relevant period reasonably well and results in a real interest rate series that is consistent with historical information on the timing of recessions and expansions. Before 1955 we assume expectations of 0% inflation, so the nominal and the real rates are equal.

⁴² Thomas (1999) provides an excellent overview of the merits of different available measures of expected inflation. We chose to use the Livingston Survey because it measures the expectations of individuals directly engaged in bond and other financial markets (e.g., economists at commercial banks, investment banking firms, corporations, academia, government, and insurance companies), and it has been conducted over a long enough timeframe to be useful for our study. Other available measures either surveyed households rather than investment professionals (e.g., University of Michigan Institute of Social Research Survey) or were not conducted over the timeframe necessary for our analysis (e.g., NBER Survey of Professional Forecasters). We phase in the moving average beginning with a one-year lag in 1955, followed by a two-year lagged average in 1956, and so on.

Even though the inflation adjustment we chose was the most compelling to us for the task at hand, there are obviously other ways one could adjust the data for inflation. To check the robustness of our results to the inflation adjustment approach, we estimated our models using series that were unadjusted (i.e., the nominal rate), that used the one-year lagged Livingston inflation measure (rather than the ten-year lagged moving average), and that used a ten-year lagged average of actual inflation in the CPI. Again, the basic results were robust to these alternative treatments of inflation.

References

- Ainslie, George. 1991. Derivation of 'Rational' Economic Behavior from Hyperbolic Discount Curves. *American Economic Review* 81:334-340.
- Arrow, K.J., W.R. Cline, K.-G. Maler, M. Munasinghe, R. Squitieri, and J.E. Stiglitz. 1996. Intertemporal Equity, Discounting and Economic Efficiency. In *Climate Change 1995 - Economic and Social Dimensions of Climate Change*, edited by J. P. Bruce, H. Lee and E. F. Haites. Cambridge, UK: Cambridge University Press.
- Ballard, Charles L., and Don Fullerton. 1992. Distortionary Taxes and the Provision of Public Goods. *Journal of Economic Perspectives* 6 (3):117-31.
- Bazerlon, Coleman, and Kent Smetters. 1999. Discounting Inside the Beltway. *Journal of Economic Perspectives* 13 (4):213-228.
- Brock, William A., and Leonard J. Mirman. 1972. Optimal Economic Growth and Uncertainty: The Discounted Case. *Journal of Economic Theory* 4:479-513.
- Campbell, John Y. 1995. Some Lessons from the Yield Curve. *Journal of Economic Perspectives* 9 (3):129-52.
- Cline, William. 1992. *The Economics of Global Warming*. Washington: Institute of International Economics.
- Cochrane, D. and G. H. Orcutt. 1949. Application of Least-squares Regression to Relationships Containing Autocorrelated Error Terms. *Journal of the American Statistical Association* 44:32-61.
- Cox, John C., Jonathan E. Ingersoll, Jr., and Stephen A. Ross. 1985. An Intertemporal General Equilibrium Model of Asset Prices. *Econometrica* 53 (2):363-384.
- Cropper, Maureen L., Sema K. Aydede, and Paul R. Portney. 1994. Preferences for Life Saving Programs: How the Public Discounts Time and Age. *Journal of Risk and Uncertainty* 8:243-265.

- Cropper, Maureen L., and David Laibson. 1999. The Implications of Hyperbolic Discounting for Project Evaluation. In *Discounting and Intergenerational Equity*, edited by P. R. Portney and J. P. Weyant. Washington: Resources for the Future.
- Dickey, D.A., and W.A. Fuller. 1979. Distribution of the Estimators for Autoregressive Time Series with a Unit Root. *Journal of the American Statistical Association* 74 (366):427-431.
- Dixit, Avinash K., and Robert S. Pindyck. 1994. *Investment Under Uncertainty*. Princeton: Princeton University Press.
- Federal Reserve. 2000. Federal Reserve Statistical Release H-15. Washington: Federal Reserve.
- Fuller, W.A. 1976. *Introduction to Statistical Time Series*. New York: Wiley.
- Hamilton, James D. 1994. *Time Series Analysis*. Princeton, NJ: Princeton University Press.
- Heal, G.M. 1998. *Valuing the Future: Economic Theory and Sustainability*. Edited by G. Chichilnisky and G. Heal, *Economics for a Sustainable Earth*. New York: Columbia University Press.
- Homer, Sidney, and Richard Sylla. 1998. *A History of Interest Rates*. 3rd ed. New Brunswick, NJ: Rutgers University Press.
- Ingersoll, Jonathan E., and Stephen A. Ross. 1992. Waiting to Invest: Investment and Uncertainty. *Journal of Business* 65 (1):1-29.
- Kwiatkowski, Denis, and et al. 1992. Testing the Null Hypothesis of Stationarity against the Alternative of a Unit Root: How Sure Are We That Economic Time Series Have a Unit Root? *Journal of Econometrics* 54 (1-3):159-78.
- Lind, Robert. 1982. A Primer on the Major Issues Relating to the Discount Rate for Evaluating National Energy Options. In *Discounting for Time and Risk in Energy Policy*, edited by R. Lind. Washington: Resources for the Future.
- Mandler, Michael. 1999. *Dilemmas in economic theory : persisting foundational problems of microeconomics*. New York, NY: Oxford University Press.

- Nelson, C. R., and C. I. Plosser. 1982. Trends and Random Walks in Macroeconomic Time Series. *Journal of Monetary Economics* 10 (2):139-162.
- Nordhaus, William D. 1994. *Managing the Global Commons*. Cambridge: MIT Press.
- Phillips, P.C.B. 1987. Time Series Regression with a Unit Root. *Econometrica* 55:277-301.
- Pizer, William A. 1999. Optimal Choice of Climate Change Policy in the Presence of Uncertainty. *Resource and Energy Economics* 21:25-287.
- Portney, Paul R., and John P. Weyant, eds. 1999. *Discounting and Intergenerational Equity*. Washington: Resources for the Future.
- Ramsey, F.P. 1928. A Mathematical Theory of Saving. *The Economic Journal* 38 (152):543-559.
- Schwarz, Gideon. 1978. Estimating the Dimension of a Model. *Annals of Statistics* 6 (2):461-464.
- Sims, Christopher A. 1988. Bayesian Skepticism on Unit Root Econometrics. *Journal of Economic Dynamics and Control* 12 (2/3):463-74.
- Sims, Christopher A., James H. Stock, and Mark W. Watson. 1990. Inference in Linear Time Series Models with Some Unit Root. *Econometrica* 58 (1):113-144.
- Spiro, Peter S. 1989. *Real Interest Rates and Investment and Borrowing Strategy*. New York: Quorum Books.
- Thomas, Lloyd B. 1999. Survey Measures of Expected U.S. Inflation. *Journal of Economic Perspectives* 13 (4):125-144.
- 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.
- . 2001. Gamma Discounting. *American Economic Review* 91 (1):260-271.