EE
7/29/03
9:30 AM

July 16, 2003
\UncertainDiscounting

Discounting a Future Whose Technology is Unknown

Abstract

This paper derives, from a stochastic neoclassical optimal growth model, a complete analytical expression for the term structure of interest rates in a dynamic competitive rational-expectations equilibrium. The primary uncertainty driving the model is the unknown trend rate of future technological progress. The growth rate of the technological progress residual is modeled here as a noisy-signal diffusion process intrinsically synchronized with a classic problem of information extraction and prediction. The paper shows that discount rates decrease over time and analyzes the extent of the decline in terms of fundamental parameters of the underlying real economy.

Martin L. Weitzman
Harvard University

Preliminary and Incomplete
Comments Appreciated

1. Introduction

Perhaps no other topic in economics combines so explosive a "triple threat" to good analysis (because it is simultaneously contentious, important, and unresolved) as the issue of how best to discount the distant future. Taking global warming as a prototype motivating example here, the economic substance of the international debate concerning what to do largely hinges on differing perceptions about this one great overriding issue of how to discount future events.

There is, of course, a great deal of uncertainty about what will be the possible results of global climate change. There is scientific uncertainty concerning geophysical and biological effects. There is economic uncertainty about possible future flows of costs and benefits under different scenarios. However, it seems a fair empirical generalization from the numbers being tossed about in the literature on the economics of global warming that the most significant uncertainty of all concerns the interest rates to be used in discounting distant future events.[1] The logic behind this empirical generalization is easily stated, as it does not involve subtle or complicated reasoning.

Suppose we take as a point of departure the ballpark estimate that the overall netted-out aggregate economic impact of global climate change might be equivalent to a loss of world income of approximately one percent per year coming on line about a century hence.[2] (Such "estimates" are more of an educated guess than anything else, but an analysis must begin somewhere.) Discounted at a "medium-low" real interest rate of four percent a year, a one-percent loss of income a hundred years from now is equivalent to a 0.02% loss of current income. To place such a number in comparative perspective, the loss of stationary-equivalent world income from the ultimate exhaustability of oil reserves is maybe 1%, while to reflect accurately the historical impact of technological progress might require an upward adjustment to

---

[1] To some extent this has been acknowledged in the literature (see, e.g., Pizer, Nordhaus, etc.).

[2] This is consistent with Nordhaus, Schneider, etc.

conventionally measured income by some 50%.[3] Viewed in this light, the issue of what to do about global warming might seem far less alarming than the issue of what to do about petroleum exhaustion – and both could appear to be vastly less significant than the issue of how to ensure that the historical pace of technological progress is maintained over future time.

To summarize thus far, if historical real interest rates from recent past centuries (and, implicitly, the underlying causative rates of growth of technological progress) are allowed to be projected forward to future centuries, then the force of compound interest effectively nullifies any plausible economic effects (good or bad) of global climate change. However, the flip side of the same coin is that, if future technological progress stalls (and real interest rates correspondingly decline), then the consequences of not now beginning measures to slow greenhouse gas emissions may be seen in a considerably more ominous hue. Therefore, as an empirical matter, with the numbers that are currently in vogue as being "reasonable" estimates, the most decisive force in conceptualizing the economics of global climate change (and in deciding what, if anything, should be done about global warming) seems to be the force of compound interest exerted over time horizons of a century or more.

Behind all projections of distant-future interest rates is the long run productivity of capital, which depends upon a host of factors unknowable to us at the present time. Foremost among these currently-unknown factors is the ability of future technological change to progress at a sufficiently rapid pace to offset diminishing returns from the accumulation of capital (for a given, fixed technology). The two centuries since the industrial revolution have been kind to us in this regard, but how can we presently be sure that the unknowable forces propelling future technological progress will continue at the same pace over the coming centuries? As has been pointed out many times, the residual of technological progress remains more a measure of our ignorance than a variable whose behavior we genuinely understand. It is essential to recognize that the typical use of interest rates (at values seeming like reasonable projections of past realizations) to discount distant-future events essentially comes down to an implicit extrapolation

---

[3] For exhaustible resources, see the estimates in Weitzman [ ]. For technological progress, see the estimates in Nordhaus [ ], Weitzman [ ], or Weitzman and Lofgren [ ].

very far forward of the growth rate of this mysterious and poorly understood residual measure of our ignorance.

If we try to imagine how the future world might look a century or so from now, we might begin by trying to conceptualize how people about a century ago might have attempted to envision our world today. We now have some very important new technologies, like computers or genetic engineering, that were essentially unimaginable a hundred years ago. Maybe a now unimaginable "photon-silicon technology" will replace today's "electron-carbon technology" and deliver such prodigious rates of technological progress with a clean environment that historians then will look back on the previous hundred years and smile at the modest projections of even the growth optimists at the beginning of the twenty-first century. Or, who knows, maybe a century from now people will feel constrained and polluted and very disappointed in a pace of technological change that failed to maintain the productivity growth of the "golden age" of the industrial revolution during the earlier two centuries from 1800 to 2000. It should be apparent that these two alternative scenarios imply very different rates of return to capital.

It follows from the above discussion that, for analyzing the process of discounting the far-distant future, it is essential to understand the nexus connecting uncertainty about future technological progress with uncertainty about future interest rates. To attain such an understanding requires a carefully constructed model that can show clearly the interrelationship between uncertain future rates of technological progress, uncertain future interest rates, and the schedule of effective forward-yield discount rates to be used now for discounting the flows of costs and benefits from projects whose expected payoffs occur in the distant future. The standard approach to modeling the term structure of interest rates in the finance literature is not fully satisfactory for such purposes, because the models are primarily a reduced-form partial-equilibrium description of interest rate formation, whose exact connection with the dynamic production process of the underlying real macro-economy is not specified.

Interest rates are essentially relative prices of consumption at different times. Like any other competitive price, an interest rate is determined by considerations of demand and supply, based, ultimately, on the outcome of a general equilibrium process operating upon the underlying specifications of tastes and technology. An interest rate is not an exogenously given primitive (not even when it is given by the solution of some particular stochastic diffusion equation), but is rather an endogenously-derived variable. For the purposes of understanding which interest rates should be used for long-term discounting, we really want to go behind the somewhat mechanical partial equilibrium specifications of interest rate formation in the term-structure literature, in order to see clearly what assumptions are being made about technological progress and other important primitives of the underlying real economy.

The aim of this paper is to make interest rate formation be a true general equilibrium phenomenon by modeling it as an endogenous process derived from a stochastic version of the standard neoclassical optimal growth problem. A key feature of the model developed here is that information about future trend rates of technological progress comes in the form of a noisy observation of past rates, whose information content must be continually extracted, filtered, and updated simultaneous with the real-time unfolding of the stochastic process. The paper focuses on highlighting the connection between the superstructure of forward-yield interest rates that show up in current financial markets and the uncertain growth of, and information about, the hidden substructure of the future real economy. More particularly, the paper focuses on explaining the evolution of interest rates in terms of three basic interconnected forms of uncertainty about technological progress: (1) *present uncertainty* about what is now the current underlying trend rate of technological progress; (2) *future uncertainty* about what will be the future underlying trend rate of technological progress; (3) *past, present and future uncertainty* about the exact relationship between the observed historical rates of technological progress (measured, from growth accounting, as a Solow residual) and the unobserved underlying trend rates of technological progress.

The deterministic case of the neoclassical optimal growth model has a flat yield curve in steady state, but this standard image of a trendless term structure is misleading because it does not survive the introduction of uncertainty. The paper shows that the stochastic generalization of such a model generates a downward sloping yield curve. Because the model yields an explicit closed-form expression for the term structure of interest rates, it is easy to analyze how the time profile depends upon the fundamental parameters of the underlying real economy, as well as upon the passage of time itself.

The interconnected cluster of issues this paper attempts to untangle is enormously complicated, involving as it does a general equilibrium interaction between discounting, uncertainty, learning, prediction, consumption, capital accumulation, and growth. The construction of a good model here involves a balancing act between sufficient richness of general structure to be interesting and sufficient specificity of imposed functional forms to allow for analytically tractable solutions and useful insights.

Although the model uses particular functional forms for the sake of analytical tractability (in order to be able to derive an especially transparent closed-form expression for the general equilibrium time profile of discount rates), it will become apparent that the basic conclusions generalize. Even with its simplifying assumptions, therefore, the paper may have implications for real business cycle theory and for the term structure of interest rates, as well as for long term discounting. So far as I am aware, the model of this paper is the first to combine together a model of noisy learning and prediction with a neoclassical model of optimal economic growth under uncertainty. The dynamic system being analyzed is simultaneously in an economic equilibrium and in an informational equilibrium, which creates a novel interaction between capital accumulation and signal extraction that has previously gone unnoted in the optimal growth literature.

The plan of the remainder of the paper is as follows. The next section reviews the deterministic neoclassical optimal growth problem (with a fixed rate of labor-augmenting technological progress) – as a point of departure for deriving the process of interest rate formation in the stochastic case. The section after that introduces noisy-signal uncertainty in technological progress. Then the stochastic optimal growth model is solved (for a Bernoulli utility function and a Gompertz production function), and the corresponding general-equilibrium time profile of interest rates is given as a closed-form expression and analyzed. Finally, the paper concludes by discussing possible implications and extensions of the basic model.

2. The Deterministic Case as a Point of Departure

There are two standard representative agents in the economy: a consumer representing all households and a producer representing all firms. In what follows, it will sometimes be convenient to see things through the eyes of a third, non-standard, agent: a banker representing financial institutions that intermediate between consumers and producers by offering bonds of all different maturities. Representative agents act perfectly competitively by optimizing as if all relevant prices are given.

Each dynastic household has constant population and lasts forever. The representative consumer has instantaneous utility function $U(C)$ and rate of pure time preference ρ. In a deterministic setting, households choose a consumption trajectory $\{C(t)\}$ to maximize, over all feasible possibilities available to them, the expression

$$\int_0^\infty U(C(t))\ e^{-\rho t}\ dt\ . \qquad \textbf{(1)}$$

To ensure that a steady-state limiting path is feasible in the presence of exponentially growing trend productivity expansion, the instantaneous utility function is chosen to be of the standard iso-elastic form

$$U(C) \ = \ \frac{C^{1-\eta}}{1-\eta}\ , \qquad \textbf{(2)}$$

where the coefficient of relative risk aversion is the positive constant $\eta = -CU''(C)/U'(C)$, which is independent of C. (A Bernoulli logarithmic utility function is a special limiting case of (2), corresponding to $\eta = 1$.)

The economy produces a single homogeneous output, which can be divided in any proportions between consumption and investment. Each firm has a (net) production function of the form

$$Y \ = \ F(K,A)\ , \qquad \textbf{(3)}$$

where Y is output net of depreciation, K is capital stock, and A is the level of labor-augmenting (or, more generally, fixed-factor-augmenting) technological progress throughout the economy. The neoclassical production function $F(K,A)$ has constant returns to scale. It is assumed that technological progress grows exogenously at fixed rate g, so that

$$d\ln A(t) = g\ dt, \tag{4}$$

or, equivalently,

$$A(t) = A_0 e^{gt}\ , \tag{5}$$

where $A_0 = A(0)$ is given.

The neoclassical optimal growth problem here takes the form of maximizing the objective function (1) subject to the constraint

$$\dot{K}(t) = F(K(t), A_0 e^{gt}) - C(t) \tag{6}$$

and with the initial condition

$$K(0) = K_0\ , \tag{7}$$

where K_0 is given.

As is well known, the first-order duality conditions characterizing the solution of the optimal growth problem have the interpretation of describing a dynamic competitive equilibrium of this neoclassical one-sector economy. Therefore, to characterize the asymptotically-approached steady state solution of the optimal growth problem is essentially the same thing as describing the limiting behavior of the corresponding competitive equilibrium.

The usual expression for the instantaneous real interest rate at time t is

$$r(t) = \left.\frac{\partial F}{\partial K}\right|_{K(t),A(t)}\ , \tag{8}$$

but the deterministic concept of the marginal product of capital does not generalize so easily to the case of uncertainty. It is therefore more useful for the purposes of this paper to conceptualize the instantaneous real interest rates at time t as being given by the dual-equivalent (for the deterministic case) formula

$$r(t) \;=\; \frac{-1}{U'(C(t))} \lim_{h \to 0+} \frac{e^{-\rho h}\, U'(C(t+h)) \;-\; U'(C(t))}{h}, \tag{9}$$

which, by contrast with (8), generalizes readily to an expected-value analogue in a stochastic setting. In a deterministic setting, (9) is just the familiar formula from optimal growth theory:

$$r(t) \;=\; \rho \;-\; \frac{\dot{U}'(C(t))}{U'(C(t))}. \tag{10}$$

A well known result from this deterministic case is that, going to the limit of a steady state along an optimal trajectory,

$$\lim_{t \to \infty} r(t) \;=\; \rho \;+\; \eta g\,. \tag{11}$$

Equation (11) is a famous formula from optimal growth theory, which is the foundation for understanding the influence of long-run rates of technological progress upon long-run interest rates. The economy approaches (and remains in) a steady state where all quantities (output, consumption, investment, capital) grow at the rate of fixed-factor-augmenting technological progress g. If, as seems reasonable, the taste parameters ρ and η are treated as being essentially constant, then, from (11), the steady-state marginal product of capital is linear in the growth rate g.

For a given fixed g in (11), the steady-state yield curve, which depicts the term structure of forward interest rates, is perfectly flat over time. Thus, the ultimate justification for using historic trend values of past real interest rates to discount future events may be interpreted in this deterministic neoclassical framework as the implicit belief that future rates of technological progress (and economic growth) will match those of the historic past. The question we want to address now is: What happens when the residual evolves as a stochastic process? The next section of the paper introduces uncertainty in the rate of growth of technological progress, which, when integrated into the neoclassical optimal growth model, will alter the term structure of interest rates in a fundamental way.

3. A Noisy-Signal Model of Uncertain Technological Progress

By far the most critical issue to be faced in any prediction of long term economic growth and interest rates is the specification of the future rate of growth of so-called "knowledge capital." (The term "knowledge capital" is here taken as being synonymous with the "state of technological progress.") While the state of knowledge capital can theoretically be modeled as part of a process of endogenous economic growth, and while such models can be very useful for the conceptual insights they may yield, endogenous growth theories are silent about predicting the future. When all is said and done about endogenous growth theories, the central fact remains that predicting the rate of growth of technological progress (or knowledge capital) involves forecasting a residual "measure of our ignorance" from the past into the future. So the approach of this paper begins with a Solow-style model of atmospheric technological change – not so much because it provides a completely satisfactory description of the growth process, but because placing the black box here focuses directly on the irreducible residual stochastic process that must be specified for any long run prediction of future economic growth.

The growth rate of technological change cannot be observed directly. Instead, it must be estimated indirectly as a Solow residual by using the famous growth-accounting formula, which here takes the form

$$\frac{dA}{A} = \frac{1}{1-S_k}\left[\frac{dY}{Y} - S_k\frac{dK}{K}\right], \qquad \textbf{(12)}$$

where

$$S_k \equiv \frac{K\frac{\partial Y}{\partial K}}{Y} \tag{13}$$

is the measured "competitive share of capital" and all variables are being evaluated at some particular observation time-point t. The point to be made here is that the number being called the "rate of growth of technological progress" for any particular time is the end product of the calculation of a residual at that time, and it therefore accumulates all of the statistical uncertainty that is inherent in each stage of such an indirect measurement process. In other words, what is actually measured by growth accounting at any time as the left-hand side of equation (12) is likely to be a very noisy signal of the underlying true trend rate of technological progress. Thus, when it comes time to model the growth of technological progress as a stochastic process, it is probably appropriate to think in terms of a relatively large standard error.

In most of what follows, it will streamline notation immensely to be able to work with the natural logarithm of A, rather than with A itself. Introducing the symbol

$$a(t) \equiv \ln A(t) \ , \tag{14}$$

we then have

$$\frac{dA(t)}{A(t)} = da(t) \ , \tag{15}$$

meaning that at any time the differential of a represents the rate of growth of technological progress.

At this point in the paper, uncertainty is introduced. For the sake of being able to derive useful results, we impose now what might be called a "Gauss-Kalman structure" of uncertainty, learning, and forecasting. The central problem of estimation and prediction for this paper is posed by postulating the following analytically-tractable structure. Suppose that the realized values of

$\{da(t)\}$ evolve as the coupled pair of stochastic diffusion equations:

$$da(t) \;=\; g(t)\;dt \;+\; \sigma_a\;dZ_a(t)\;, \tag{16}$$

and

$$dg(t) \;=\; \sigma_g\;dZ_g(t)\;. \tag{17}$$

In the above specification (16), (17), the stochastically-evolving random variable $g(t)$ represents the not-directly-observable underlying "true" trend growth rate of technological progress at time t, while $da(t)$ is the corresponding actual realized value of the growth rate of the Solow residual that is indirectly measured by an observer just immediately after time t. The stochastic differentials $dZ_a(t)$ and $dZ_g(t)$ refer to two independently distributed Wiener processes. The parameters σ_a and σ_g are presumed to be known constants representing the standard deviations of the stochastic diffusion processes (16) and (17).

The system (16), (17), along with the definition (14), is intended to portray the evolution of technological knowledge over very long periods of time – as measured in centuries, or perhaps even millennia. In this spirit, it is analytically convenient to assume that the decision maker at time t has a historical record of all previous measurements of rates of growth of technological progress – i.e., the decision maker at time t knows all past values of the time series

$$\{da(s)\}_{s\le t}\;. \tag{18}$$

Suppose, for the sake of argument, that the decision maker at time t does not know the underlying true trend rate $g(t)$ as being anything other than some realization of a normally distributed probability density

$$g(t) \sim N(\mu(t), V(t))\;. \tag{19}$$

It can be shown[4] that observing all past values of (18) and assuming the postulated

---

[4] See, for example, Liptser and Shirayev (2000) Theorem 12.1 (page 22), for which the result cited in this paper is a special case.

architecture (16), (17) implies the remarkably strong conclusion for all t that

$$\mu(t) = \int_{-\infty}^{t} \theta \, e^{-\theta(t-s)} \, da(s) , \tag{20}$$

and

$$V(t) = \sigma_a \sigma_g , \tag{21}$$

where

$$\theta \equiv \frac{\sigma_g}{\sigma_a} . \tag{22}$$

Conditions (20), (21), (22) mean that the economy is in what might be called a state of "informational equilibrium." At time t the observer has a probabilistic estimate of the true underlying trend rate $g(t)$, as given by (19). At time $t+dt$, the precision of this estimate is simultaneously both *decreased* (because, from (17), the unobserved true value has drifted from $g(t)$ to $g(t)\pm\sigma_g\sqrt{dt}$) and *increased* (because, by (16), we now have one more sample point, namely $da(t+dt)=g(t+dt)dt\pm\sigma_a\sqrt{dt}$, with which to estimate $g(t+dt)$ more accurately). In a steady state of informational equilibrium, these two effects offset each other exactly. (It can be shown that just such an informational steady state is approached in the limit of a Bayesian learning process as the number of observations approaches infinity. The driving force behind the analytical tractability of formulas (20), (21), (22) is based, ultimately, on the analytical tractability of the Bayesian updating formula when sampling from a normal distribution whose mean is itself a normal distribution.)

In a state of informational equilibrium, the rate at which the system is resolving "old uncertainties" by learning more over time about the underlying structure is being just nullified by the arrival of "new uncertainties" about changes over time in that very structure. What might be said to be the *irreducible uncertainty* of the system in this kind of informational equilibrium has an unchanging variance over time, given by formula (21). In the long term informational equilibrium of this Gaussian setup the "best estimate" of $g(t)$ – namely, its mean, $\mu(t)$ – changes stochastically over time in accordance with (20), while the variance of this "best estimate" remains constant at the value V determined by (21).

Note from (21) that the derived standard deviation $\sqrt{V}$ of the "best estimate" of the true trend value is the symmetric geometric mean of both of the more-primitive underlying standard deviations σ_a and σ_g. Combining (19), (20), (21), (22), as $\sigma_g \to 0^+$ there is zero residual uncertainty in the system, and the "best estimate" is known with the precise accuracy of the law of (infinitely) large numbers to be the average of all past observed values corresponding to the limiting case $\theta \to 0^+$. Analogously, as $\sigma_a \to 0^+$ there is also zero residual uncertainty, and the "best estimate" is known to be exactly the very last observation of the data point $da(0^-)$, which corresponds to the limiting case $\theta \to +\infty$. In this sense, the irreducible uncertainty of information about the underlying trend rate of growth of technological progress might be said to arise here from the *combination* of a noisy signal wobbling around a randomly walking true value.

From (16), the underlying trend rate of technological progress has the potential to become negative (but only at some time in the very remote future if σ_x is sufficiently small). The existence of such a possibility is appropriate for a model aspiring to portray very long term growth trends (although we might want to be cautious about actual applications of such an idea to predict distant-future interest rates) because, over the ten or so millennia since the Neolithic revolution, humankind has endured many cycles of boom and bust. There is little in the approximately four-thousand-year-old written record of civilization (or, for that matter, in what we can infer of the human condition before that time) to suggest that societies cannot experience declines from wars, pestilence, earthquakes, overpopulation, limitations of carrying capacity, running out of resources, running out of new ideas, or just plain bad luck of many other kinds. Actually, if the specification did *not* allow negative trend growth under any circumstances, such an omission would leave the model open to the legitimate complaint of a doomsday critic, who thinks that we cannot keep growing at the same rates as we have been for the last century or two, and who could argue that the equations have been rigged to exclude the very scenario (of a worst-case "dark age") whose implications are actually the most important in deciding the interest rates to be used for discounting distant future events, such as global warming.

4. Using a Stochastic Optimal Growth Model to Forecast Interest Rates

To summarize the Gauss-Kalman model of technological progress presented in the previous section, at any time t the decision maker knows the observable state variable $\mu(t)$ (from (20)), but does not know the exact value of the hidden state variable $g(t)$. From the viewpoint of a representative agent making decisions in such an economy at time zero, there are four state variables – namely K, A, μ, and g – of which only the first three are observable, while the fourth is hidden but is known to evolve by the stochastic process (17).

Any feasible policy can be based only upon the three observable state variables K, A, and μ. At time t, the known state of the system is summarized by the triplet $(K(t), A(t), \mu(t))$. At the present time zero, the inherited initial conditions are that the three state variables $K(0)$, $A(0)$, and $\mu(0)$ are known exactly (as the given values K_0, A_0, and μ_0, respectively), while the initial value of the fourth state variable $g(0)$ is known only as a probability distribution satisfying (19), for t=0, and whose variance is given by (20).

Let

$$C = \Psi(K,A,\mu) \qquad \textbf{(23)}$$

be any policy function prescribing consumption as a function of the observable state variables K, A, and μ. Given any such policy function (22), the economy evolves stochastically according to the four simultaneous differential/diffusion equations of state

$$\dot{K}(t) = F(K(t),A(t)) - \Psi(K(t),A(t),\mu(t)) \ , \qquad \textbf{(24)}$$

$$da(t) = g(t)\ dt + \sigma_a\ dZ_a(t) \ , \qquad \textbf{(25)}$$

$$d\mu(t) \;=\; \frac{\sigma_g}{\sigma_a}\,[da(t) \;-\; \mu(t)dt] \;, \tag{26}$$

$$dg(t) \;=\; \sigma_g \; dZ_g(t) \;, \tag{27}$$

with the three deterministic initial conditions

$$K(0)=K_0 \;, \quad A(0)=A_0 \;, \quad \mu(0)=\mu_0 \;, \tag{28}$$

and the probabilistic initial condition

$$g(0) \sim N(\mu_0,\sigma_a\sigma_g) \;. \tag{29}$$

By way of explanation, the differential equation (24) is just formula (6) from the deterministic case transposed to the notation of the stochastic version. The diffusion equation (25) is merely (16) combined with the definition (14). The diffusion equation (26) is obtained by differentiating (20) and making use of (22). Finally, the diffusion equation (27) is exactly the same as (17). The only other condition possibly requiring comment is (29), which is a rewritten version of (19), incorporating (28) and (21).

We are now ready to state formally the relevant stochastic optimal growth problem for this economy.

An *optimal* policy function

$$C \;=\; \Psi^*(K,A,\mu) \tag{30}$$

is a policy function such that, for any *other* policy function (22), it must hold that

$$E[\int_0^\infty U(\Psi^*(K(t),A(t),\mu(t)))\; e^{-\rho t}\; dt] \;\;\geq\;\; E[\int_0^\infty U(\Psi(K(t),A(t),\mu(t)))\; e^{-\rho t}\; dt] \;, \tag{31}$$

where the corresponding trajectories generated by both policy functions must satisfy the feasibility conditions of the form (23)-(28). (The operator notation $E[X]$ stands for the expected value of the random variable X.)

From basic dynamic-finance theory, the competitive-equilibrium instantaneous goods interest rate at time t in such a stochastic economy is a straightforward probabilistic generalization of (9), namely

$$r(t) = \frac{-1}{U'(C^*(t))} \lim_{h \to 0^+} \frac{e^{-\rho h} E[U'(C^*(t+h))] - U'(C^*(t))}{h} , \qquad \textbf{(32)}$$

where for convenience the abbreviated notation being employed is

$$C^*(t) \equiv \Psi^*(K(t), A(t), \mu(t)) . \qquad \textbf{(33)}$$

Suppose that now, at the present time zero, we are considering an investment opportunity promising to pay one small extra unit of the single homogeneous good at some future time t. The relevant discount factor for converting a marginal unit of the good at time t back to a marginal unit of the good at time zero is

$$D(t) = E[\exp(-\int_0^t r(s)ds)] \qquad \textbf{(34)}$$

The forward-yield interest rate, to be used for discounting marginal changes in goods at time t back to time zero is then

$$R(t) = \frac{-\ln D(t)}{t} . \qquad \textbf{(35)}$$

The primary goal of this paper is to provide a usable analytical expression for the forward-yield interest rate $R(t)$, defined by (35), in terms of basic parameters of the underlying real economy. Although relatively easy to state, this problem is very difficult to solve in the general case. Unfortunately, there is no way of getting a usable form of $R(t)$ without placing a lot more structure on the problem. We have already imposed a considerable amount of structure on the uncertainty part of the problem (in the form of the Gaussian specification of the last section of the paper). What remains to be done is an analogous further structuring of the utility and production functions. We want to provide analytically tractable, yet reasonable, functional forms for the utility and production functions. To the details of these important issues of model specification we now turn.

5. Solving the Model with Bernoulli Utility and Gompertz Production

The stochastic optimal growth problem posed in the previous section involves four state variables (one of which is known only as a probability distribution). As perhaps will be appreciated, such a control problem is analytically intractable in the general case. What is required is a particular specification of utility and production functions that is reasonable, yet will allow us to cut cleanly through the morass of four diffusion/differential equations by exploiting powerfully the strong decompositional properties of some ultra-simple optimal policy rule.

Just such magic trick is provided by the combination of utility and production functional forms about to be introduced. It will become readily apparent that the results of the paper generalize, but in the more general case it is impossible to find a closed-form expression for the optimal solution (or for the term structure of interest rates), and one is stuck with the frustrating task of proving messy theorems about the qualitative behavior of a general solution.

Concerning the instantaneous utility function, the previously-imposed isoelastic form (2) is henceforth further restricted to the parameter value $\eta = 1$. This corresponds to the famous "Bernoulli utility function"

$$U(C) = \ln(C), \qquad \textbf{(36)}$$

which was historically the first explicitly written utility function, being used by Daniel Bernoulli in 1727 to resolve the St. Petersburg paradox.

The production side is characterized by the (average) productivity of capital being loglinear in capital (at a given state of technological knowledge), meaning

$$\frac{F(K,A)}{K} = \alpha - \beta \ln(K/A) \tag{37}$$

for non-negative parameters α and β, which corresponds to the constant-returns-to-scale production function

$$F(K,A) = \alpha K - \beta K \ln(K/A). \tag{38}$$

The form (38) was first introduced by Benjamin Gompertz in 1825 and is widely used throughout fisheries economics, in which context, without loss of generality, α is commonly taken as zero while A stands for the "carrying capacity" of the fishery. Technological progress may then be seen as either augmenting the productivity of fixed background factors (like labor) or, equivalently here, increasing the carrying capacity of the system, or both.

Although it was originally intended to describe the arithmetic growth of a biological population (corresponding to zero harvesting or consumption), the "Gompertz function" (38) can lay claim to being the first explicitly written production function in history, as it precedes by thirteen years Pierre Verhulst's quadratic "function logistique" and by over a century the famous geometrically-weighted functional form introduced formally in 1927 by Charles Cobb and Paul Douglass (although Wicksell, Wicksteed and Walras were also apparently aware of it).[5] While expression (11) is perhaps unfamiliar to most non-fisheries economists, it represents a perfectly legitimate production function, which, as will be seen presently, possesses some extraordinarily tractable analytical qualities when combined with Bernoulli utility (and with modeling the growth of A as a Gauss-Kalman stochastic diffusion process of the type specified in the previous section of the paper).

The parameter β measures the *strength of diminishing returns*, since, for the production function (38), β always equals the difference between the average product of capital and the marginal product of capital. This statement is readily confirmed by verifying of expression (38)

---

[5] See entry in New Palgrave.

that it satisfies the equation

$$\frac{F(K,A)}{K} - \frac{\partial F}{\partial K} = \beta \tag{39}$$

for all K and A.

To appreciate why β may play a significant role in the analysis that follows, use the definition (14) of $a(t)$ to rewrite (38) as

$$Y_t(K) = [\alpha + \beta a(t) - \beta \ln K]\ K\ , \tag{40}$$

which means that the term $\beta a(t)$ is acting like a "shift operator" over time, offsetting the diminishing returns to capital accumulation caused by the term $-\beta \ln K$. Suppose, further, that there is no uncertainty and that $A(t)$ grows exponentially at the constant rate g. In this case, (40) becomes

$$Y_t(K) = [\alpha + \beta a(0) + \beta g t - \beta \ln K]\ K\ . \tag{41}$$

Equation (41) means that β may be viewed as a measure of the dependence or "sensitivity" of the growth process to the rate of technological progress g. If $\beta = 0$, there exists a linear technology that can be replicated indefinitely, without any need for technological progress. (In this case the interest rate is precisely α, always.) However, in the much more relevant case where $\beta > 0$ there is diminishing returns to the accumulation of capital, and without the offsetting effect of technological progress economic growth would eventually grind to a halt. (In this limiting stationary state without technological progress, the interest rate would be ρ.) The existence of technological progress growing at rate $g > 0$ offsets the diminishing returns to capital accumulation by augmenting the hidden fixed factor ("labor" or other "fixed background factors" or "carrying capacity" – depending upon the interpretation). The important point here is that any effect of changes in g upon the growth process are transmitted *via* β in the form of a multiplicative package βg. Hence the importance of β for the analysis that follows.

Rather than going right to the main result of the paper, we approach that goal gradually by trying to build up intuition about the properties of an optimal solution via a series of three lemmas.

The ability of the Bernoulli-Gompertz combination to perform a "magic trick" (in enabling a stochastic optimal growth problem to be decomposed and easily solved) derives essentially from the following strong property. Let $A(t)$ be some given positive function of t. Consider the deterministic Bernoulli-Gompertz time-dependent optimal growth problem to maximize

$$\int_0^\infty \ln(C(t))\; e^{-\rho t}\; dt \tag{42}$$

subject to the constraint

$$\dot{K}(t) \;=\; \alpha K(t) \;-\; \beta K(t)\ln(K(t)/A(t)) \;-\; C(t) \tag{43}$$

and with the initial condition

$$K(0) \;=\; K_0 \;. \tag{44}$$

The following lemma then drives the analytical tractability of the rest of the paper.

Lemma 1. For *any* positive function $A(t)$, the optimal solution of the deterministic control problem (42)-(44) is the myopic consumption rule

$$C^*(t) \;=\; (\beta+\rho)\; K(t) \;. \tag{45}$$

Sketch of proof: As is well known from deterministic optimal growth theory, a necessary (and, with convexity, sufficient) condition for the consumption instrument to be optimal with a logarithmic utility function (i.e.,a utility function of the iso-elastic class (12) with $\eta=1$) is that it satisfy everywhere along its trajectory the first-order condition

$$\frac{\dot{C}}{C} \;=\; \frac{\partial F}{\partial K} - \rho \;. \tag{46}$$

Plugging (45) into (43), and calculating the marginal product of capital for the Gompertz production function, it is readily confirmed by direct inspection that the policy (45) satisfies the first-order condition (46). (The only possible loose end in this proof is the transversality condition, which can be handled by a separate argument, omitted here to save space.) □

Since the optimal consumption policy given by (45) is independent of $\{A(t)\}$, it must also be independent of any stochastic process driving $\{A(t)\}$. It is this feature that allows the very complicated-looking stochastic problem (23)-(31) to be decomposed and solved relatively easily. The next lemma describes the interest rate at any time along an optimal trajectory as a function of the state variables at that time. The point of departure is the general expression (32) for the equilibrium interest rate at any time along an optimal trajectory.

Lemma 2. At time t along an optimal trajectory, the interest rate defined in (32) is given by the formula

$$r(t) \;=\; \alpha - \beta + \beta \ln A(t) - \beta \ln K(t) \quad [= \frac{\partial F}{\partial K}\Big|_{K(t),A(t)}\,] \,. \tag{47}$$

Sketch of proof: From (45), the marginal utility of consumption along an optimal trajectory at time t for the Bernoulli utility function is

$$U'(C^*(t)) \;=\; \frac{1}{(\beta+\rho)K(t)} \,, \tag{48}$$

whereas at time $t+h$ it is

$$U'(C^*(t+h)) \;=\; \frac{1}{(\beta+\rho)K(t+h)} \,. \tag{49}$$

Combining (48) with (49) and differentiating, we have

$$\lim_{h \to 0^+} \frac{e^{-\rho h}\; U'(C^*(t+h)) \,-\, U'(C^*(t))}{h} \;=\; \frac{-1}{(\beta+\rho)K(t)}\,(\rho + \frac{\dot{K}(t)}{K(t)}) \,. \tag{50}$$

From the optimal consumption policy formula (45) and from the definition (38) of a Gompertz production function, along an optimal trajectory it must hold that

$$\dot{K}(t) \;=\; \alpha K(t) \;-\; \beta K(t)\;\ln(K(t)/A(t)) \;-\; (\beta+\rho)K(t)\;. \tag{51}$$

The desired conclusion (47) then follows from plugging (48), (50), and (51) into (32), and consolidating all loose terms. □

Lemma 2 shows that formula (8) holds for an optimal policy, but the reader should be warned that such a certainty-equivalence-like conclusion is not general in a stochastic setting. It occurs here only because of the special Bernoulli-Gompertz structure. Thus, the Bernoulli-Gompertz logarithmic specification induces an extremely tractable certainty-equivalent-type result having very similar useful properties to what emerges from the more familiar linear-quadratic optimal control problem.

The final lemma gives the reduced-form stochastic evolution of equilibrium interest rates over time.

Lemma 3. Interest rates (the formula for which is given by Lemma 2) evolve over time in the stochastic Bernoulli-Gompertz economy according to the pair of diffusion equations

$$dr(t) \;=\; \beta[\hat{r}(t)-r(t)]\;dt \;+\; \beta\sigma_a\;dZ_a(t)\;, \tag{52}$$

and

$$d\hat{r}(t) \;=\; \sigma_g\;dZ_g(t)\;, \tag{53}$$

and with the additional initial condition

$$\hat{r}(0) \;=\; g(0)+\rho\;. \tag{54}$$

Sketch of proof: Differentiating (47) yields

$$dr(t) \;=\; \beta\;da(t) \;-\; \beta\frac{\dot{K}(t)}{K(t)}dt\;. \tag{55}$$

Plugging (47) into (51) and rearranging terms yields

$$\frac{\dot{K}(t)}{K(t)} = r(t) - \rho \ . \tag{56}$$

Now substitute from (56) into (55), thereby obtaining

$$dr(t) = \beta[da(t) + \rho dt - r(t)dt] \ . \tag{57}$$

By making use of (16), equation (57) may then be rewritten as

$$dr(t) = \beta[g(t)dt + \rho dt - r(t)dt + \sigma_a dZ_a(t)] \ . \tag{58}$$

Define

$$\hat{r}(t) \equiv g(t) + \rho \ , \tag{59}$$

so that

$$d\hat{r}(t) = dg(t) \ . \tag{60}$$

The initial condition (54) follows immediately from definition (59). Substituting (59) into (58) yields (52). And, finally, combining (60) with (17) gives (53). □

From (53), the same noisy-signal problem that plagues the exact identification of the underlying rate of technological progress $g(0)$ in this economy at the current time zero also plagues the exact identification of the underlying trend interest rate $\hat{r}(0)$. Indeed, from (29) and the definition (59), the initial condition $\hat{r}(0)$ is known only as the probability distribution

$$\hat{r}(0) \sim N(\bar{r}(0), \sigma_a \sigma_g) \ , \tag{61}$$

where

$$\bar{r}(0) \equiv \mu_0 + \rho \ . \tag{62}$$

Armed with the three lemmas, we are now ready to exhibit the equilibrium term structure of interest rates in this stochastic Bernoulli-Gompertz economy. The general formula looks formidable at first glance and depends upon the initial condition for the current interest rate at time zero, $r(0)$. However, it is relatively easy to characterize the limiting behavior, which is what we are mostly interested in anyway for long term discounting.

6. The Equilibrium Term Structure of Interest Rates

The following proposition is the main result of the paper.

Theorem. With the model of this paper, the forward-yield interest rate defined by (35) is

$$R(t) \;=\; \bar{r}(0) + (r(0)-\bar{r}(0))(\frac{1-e^{-\beta t}}{\beta t}) - \frac{1}{2}\sigma_a^2 \int_0^1 (1-e^{-\beta t x})^2 dx$$

$$- \frac{1}{2}\sigma_a\sigma_g \,[1-(\frac{1-e^{-\beta t}}{\beta t})]^2\, t - \frac{1}{2}\sigma_g^2 \,\{\int_0^1 [x - (\frac{1-e^{-\beta x t}}{\beta t})]^2 dx\}\, t^2 . \qquad \textbf{(63)}$$

Sketch of proof: We begin by substituting the dummy variable τ for the dummy variable t, and then integrate the diffusion equation (53) to express it in the integral form

$$\hat{r}(\tau) \;=\; \hat{r}(0) + \sigma_g Z_g(\tau). \qquad \textbf{(64)}$$

Using (64), rewrite (52) as the linear diffusion equation,

$$dr(\tau) \;=\; -\beta r(\tau)d\tau + f(\tau)d\tau\ , \qquad \textbf{(65)}$$

where the forcing function is

$$f(\tau) \;=\; f_1(\tau) + f_2(\tau) + f_3(\tau)\ , \qquad \textbf{(66)}$$

with

$$f_1(\tau)d\tau \;\equiv\; \beta\hat{r}(0)d\tau\ , \qquad \textbf{(67)}$$

and

$$f_2(\tau)d\tau \equiv \beta\sigma_g Z_g(\tau)d\tau \ , \qquad \textbf{(68)}$$

and

$$f_3(\tau)d\tau \equiv \beta\sigma_a dZ_a(\tau) \ . \qquad \textbf{(69)}$$

The integrated solution of (65) is

$$r(\tau) = r(0)e^{-\beta\tau} + \int_0^\tau f(s)\ e^{-\beta(\tau-s)}ds \ . \qquad \textbf{(70)}$$

Next, define the random variable

$$X(t) \equiv \int_0^t r(\tau)\ d\tau \ . \qquad \textbf{(71)}$$

From (65)-(71), we then have

$$X(t) = X_0(t) + X_1(t) + X_2(t) + X_3(t) \ , \qquad \textbf{(72)}$$

where

$$X_0(t) \equiv \int_0^t r(0)e^{-\beta\tau}d\tau \ , \qquad \textbf{(73)}$$

and

$$X_1(t) \equiv \int_0^t [\int_0^\tau f_1(s)\ e^{-\beta(\tau-s)}ds]d\tau \ , \qquad \textbf{(74)}$$

and

$$X_2(t) \equiv \int_0^t [\int_0^\tau f_2(s)\, e^{-\beta(\tau - s)} ds] d\tau \ , \tag{75}$$

and

$$X_3(t) \equiv \int_0^t [\int_0^\tau f_3(s)\, e^{-\beta(\tau - s)} ds] d\tau \ . \tag{76}$$

Because $X(t)$ is a weighted sum of normally distributed random variables, it is itself a normally distributed random variable. The remainder of the proof consists of calculating its mean and variance by calculating the means and variances of its four independently-distributed component sub-terms from (72). This part of the proof essentially amounts to a rather lengthy and messy series of brute-force calculations based on the theory of linear differential equations with a forcing function. Only the most basic steps are presented here.

From (73), the mean of $X_0(t)$ is

$$m_0(t) = r(0)\, \frac{1 - e^{-\beta t}}{\beta} \ , \tag{77}$$

while its variance is

$$V_0(t) = 0 \ . \tag{78}$$

Turning now to $X_1(t)$, by (74) and (67) we have

$$X_1(t) = \hat{r}(0)\ [t - \frac{1 - e^{-\beta t}}{\beta}] \ , \tag{79}$$

and it is immediately apparent from combining (79) with (61) that $X_1(t)$ is normally distributed with mean

$$m_1(t) = \bar{r}(0)\ [t - \frac{1-e^{-\beta t}}{\beta}] \tag{80}$$

and variance

$$V_1(t) = \sigma_a \sigma_g\ [t - (1-e^{-\beta t})/\beta]^2 \ . \tag{81}$$

Next, combine (68) with (75) to obtain

$$X_2(t) = \sigma_g \int_0^t [\int_0^\tau Z_g(s)\ \beta e^{-\beta(\tau-s)} ds] d\tau \ . \tag{82}$$

Use integration by parts and the fact that $Z_g(0)=0$ to evaluate the integral within the square brackets of the right hand side of (82) as

$$\int_0^\tau Z_g(s)\ \beta e^{-\beta(\tau-s)} ds = Z(\tau) - \int_0^\tau e^{-\beta(\tau-s)} dZ_g(s) \ . \tag{83}$$

Substituting (83) into (82) yields the expression

$$X_2(t) = \sigma_g \int_0^t [\int_0^\tau (1-e^{-\beta(\tau-s)}) dZ_g(s)] d\tau \ . \tag{84}$$

Changing the order of integration, (84) becomes

$$X_2(t) = \sigma_g \int_0^t [\int_s^t (1-e^{-\beta(\tau-s)})\ d\tau]\ dZ_g(s) \ , \tag{85}$$

which, carrying out the integration within the square brackets of the right hand side, becomes

$$X_2(t) \;=\; \sigma_g \int_0^t [(t-s) \;-\; (1-e^{-\beta(t-s)})/\beta] \; dZ_g(s) \;, \tag{86}$$

Making the change of variables $y=t-s$ in (86), it becomes clear that $X_2(t)$ is normally distributed with mean

$$m_2(t) \;=\; 0 \tag{87}$$

and variance

$$V_2(t) \;=\; \sigma_g^{\,2} \int_0^t [y \;-\; (1-e^{-\beta y})/\beta]^2 \; dy \;. \tag{88}$$

Turning, finally, to $X_3(t)$, from (76) and (69) it may be expressed as

$$X_3(t) \;=\; \sigma_a \int_0^t [\int_0^\tau \beta e^{-\beta(\tau-s)} \; dZ_a(s)] d\tau \;. \tag{89}$$

Change the order of integration in (89) to rewrite it as

$$X_3(t) \;\equiv\; \sigma_a \int_0^t [\int_s^t \beta e^{-\beta\tau} \; d\tau] e^{\beta s} \; dZ_a(s) \;, \tag{90}$$

which may immediately be simplified to

$$X_3(t) \;=\; \sigma_a \int_0^t [1 \;-\; e^{-\beta(t-s)}] \; dZ_a(s) \;. \tag{91}$$

Making the change of variables $y=t-s$ in (91), it becomes clear that $X_2(t)$ is normally distributed with mean

$$m_3(t) = 0 \tag{92}$$

and variance

$$V_3(t) = \sigma_a^2 \int_0^t (1-e^{\beta y})^2 \, dy \, . \tag{93}$$

From its definition (72), the random variable $X(t)$ is normally distributed with mean

$$m(t) = m_0(t) + m_1(t) + m_2(t) + m_3(t) \tag{94}$$

and variance

$$V(t) = V_0(t) + V_1(t) + V_2(t) + V_3(t) \, . \tag{95}$$

Therefore, the random variable

$$W(t) = \exp(-X(t)) \tag{96}$$

is log-normally distributed and from the relevant theory has mean

$$E[W(t)] = \exp(-m(t)+V(t)/2) \, . \tag{97}$$

Combining (34) with (97) yields

$$D(t) = \exp(-m(t)+V(t)/2) \, . \tag{98}$$

Combining (98) with the definition (35), we then have

$$R(t) = \frac{1}{t} \left[m(t) - \frac{V(t)}{2}\right] . \tag{99}$$

The final expression (63) is obtained from (99) by writing out the explicitly-derived formulas for the four sub-components of (94), (95), and by introducing the change of variables $x = y/t$. □

6. Analysis and Implications

To be presented at NBER meeting. Some points to be covered:

Interpretation of the five terms of (63).

Meaning and significance of possible negative interest rates.

What do results mean?

Suggestion that declining interest rates can be an important feature of long-term discounting. May be deeply embedded in any model of evolving uncertainty.

Distinction between evolutionary uncertainty and stationary uncertainty. Reconciliation and final comments on applicability to discounting.