Available online at www.sciencedirect.com
SCIENCE DIRECT
Japan and the World Economy
16 (2004) 277–301
Japan and the WORLD ECONOMY
www.elsevier.com/locate/econbase

Stochastic income and wealth☆

Martin L. Weitzman*

Department of Economics, Harvard University, Littauer 313, Cambridge, MA 02138, USA

Received 4 April 2003; received in revised form 2 May 2003; accepted 14 April 2004

1. Introduction to the theory of complete accounting under uncertainty

As is well known, the deterministic version of the maximum principle allows us to state rigorously the relationship between deterministic income (essentially the Hamiltonian) and deterministic wealth (essentially the state evaluation function). As a background motivating point of departure, consider the simple story of an infinitely long-lived individual whose sole wealth consists of a bank deposit paying a constant rate of interest. For this simple parable, whether we define and measure income in the spirit of Fisher as being the return on wealth, or in the spirit of Lindahl as being consumption plus net value of investment, or in the spirit of Hicks as being the largest permanently maintainable level of consumption, we get the same answer to the question ‘‘what is income?’’. The theory of the deterministic maximum principle shows is that this fundamental identity of the three seemingly different definitions of income is much broader and goes much deeper than the simple bank-account parable.

We now seek to extend the investigation to cover uncertainty by dealing with the case where net capital accumulation is described by a stochastic diffusion equation in place of a deterministic differential equation. Stochastic diffusion equations introduce mathematical subtleties and complexities, which we will only deal with casually in this treatment.

The introduction of uncertainty ratchets up the required level of analysis another several notches on the scale of mathematical complexity. To treat the subject of this paper fully rigorously and also at the level of generality of a multi-dimensional economic growth problem could easily turn this paper into a book and would be out of all proportion to its

☆ *Note:* This paper is an abridged version of Chapter 7 of my forthcoming book ‘‘Income, Wealth, and the Maximum Principle’’ (Harvard University Press, Spring 2003).

* Tel.: +1 617 495 5133; fax: +1 617 495 7730.
E-mail address: mweitzman@harvard.edu.

0922-1425/$ – see front matter
doi:10.1016/j.japwor.2004.04.002

intended purpose (presuming even that it could be done). As a result of the inherent mathematical complexity of the subject, in this paper we make almost no attempt at being rigorous or general. Instead, to focus as sharply as possible on the core content of the wealth and income version of the maximum principle under uncertainty, we just assume the simplest imaginable stochastic generalization of the basic one dimensional calculus-of-variations prototype economic control problem, along with whatever assumptions it takes to make all of the relevant functions well behaved for the purpose at hand. Then we content ourselves with stating the relevant results in an intuitively reasonable fashion and providing an "economist's proof" consisting of mathematically plausible, but ultimately heuristic, arguments. Economic applications of stochastic diffusion processes constitute what is by now an important area of economics. There exist several books treating this topic at varying levels of mathematical sophistication, which the interested reader is encouraged to consult.[1]

Because it is very easy to get lost in the mathematical details of stochastic diffusion processes, whose rigorous foundation is quite intriguingly sophisticated, we should keep our well-defined goals here sharply in view throughout this paper. These limited are the following. *We are seeking to connect income with wealth* (or welfare) – *under uncertainty.* The major issue here boils down to giving a *convincing wealth-and-income economic interpretation of the* (so-called) *Hamilton–Jacobi–Bellman equation* for the simplest stochastic generalization of the prototype calculus-of-variations economic problem If we keep this limited aim firmly in mind, it can serve as a natural boundary delimiting this paper's treatment of investment and uncertainty from the various treatments of investment and uncertainty available in other papers and books treating other economic applications of stochastic diffusion processes. This paper, then, assumes some prior knowledge of probability theory and stochastic processes; it is more directed at using such a framework to shed light on the stochastic relationship between income and wealth (or welfare) than at rigorously explaining the probabilistic framework itself.

As we will see, the stochastic connection between income and wealth depends critically on some subtle issues of timing, measurement, and information. For example, it will become crucial to specify carefully *what* is being measured as income, *when* is it being measured, and *how* is it being measured. In the deterministic case the appropriate timing and measurement of income flows was so apparent that no special discussion was warranted. With stochastic diffusion processes, all of this can change dramatically, since seemingly slight differences of specification about *when* in the production period the price term of a chain-linked Divisia production index is evaluated (beginning, middle, or end) can give quite different relations between wealth and income. In a stochastic-diffusion economy, it *matters* how index numbers of income or production are constructed.

The workhorse model of this paper treats the one-dimensional firm (or investment opportunity) under uncertainty. In this case direct gain is measurable in money terms as a dividend or payout, the firm's stock market value is observable, and income will refer to *expected true* earnings (dividends plus expected net investment evaluated at opportunity cost). Within this context, the main result of the paper is to state and prove a stochastic

[1] Part II ("Mathematical Background") of Dixit and Pindyck (1994) is particularly recommended as a first introduction to the subject.

wealth and income version of the maximum principle. (The main result will be that forward-looking comprehensively accounted *expected-immediate-future* true earnings must exactly equal the firm's opportunity cost of capital times the observed stock market value of all its outstanding shares.)

It is possible to generalize the stochastic results of this chapter into multi-sector stochastic growth models without differentiability assumptions. However, we make no serious effort to develop formally such generalizations. (This, like a mathematically rigorous version of the present paper, would require a book of its own.) Instead, we content ourselves here with the hope that, from a good intuitive understanding of the most basic stochastic control model, which we endeavor to convey in this paper, there will develop on its own in the reader a sense of what should be the broad outlines of the analogous 'stochasticized' versions of the multi-dimensional maximum principle.

2. The deterministic case in terms of policy functions

We begin by reviewing the simplest case of the deterministic one-capital firm with 'perfectly complete' income accounting. We will recast the solution to this problem in a somewhat non-traditional form to ease the transition to, and emphasize the connection with, the stochastic version, which will be the primary focus of our attention here. Although the main stochastic result generalizes to multi-capital cases with various constraints on the control variables and many other complications, we purposely choose here the simplest imaginable one-dimensional unconstrained version in order to focus as sharply as possible on the core relationship between wealth and income under uncertainty. It is all the more essential to deal with the simplest imaginable case here because, even so, there will be some tricky conceptual (and mathematical) issues involved in the timing and measurement of "expected true earnings," which concept, while it is generalizable, is best examined initially in a pure form that is as free as possible of any conceivable distracting complications.

Although all kinds of generalizations are possible (including consumer portfolio choice formulations), for the sake of having a particularly sharp image with a crisp statement and a vivid storyline we interpret the one-dimensional prototype problem as modeling the dynamic behavior of a firm whose shares are publicly traded. A good specific example to keep in mind might be the optimal extraction of a fixed pool of oil by a publicly-held Hotelling monopolist. Another particularly good example to carry along for viewing in the mind's eye could be a firm described by the q-theory of investment, whose shares are competitively traded. Of course the theory of the stochastic maximum principle covers a much more general situation than just these two applications, but it can help the conceptualization and intuition greatly to keep in mind, as we develop the concepts of this paper, a specific example or two. In particular, it will aid clear thinking throughout what follows to think of investment and capital in *real or physical* terms – i.e., barrels of oil, numbers of trucks, and so forth. This will automatically reinforce the core idea that, even in the multi-sector case, we are assuming that *every investment has an internal shadow accounting price for the firm*, which can vary with changes in the current level of investments or as the background stocks of capital are altered over time.

Let us start with the primitives describing the one-dimensional firm's environment, which are assumed to be given. The control variable is taken to be net investment, I. Essentially, any other control variable (or variables) could be transformed into this reduced-form instrument by a change of variables. For simplicity we are assuming that the choice of I is unconstrained so that, at least in principle, I is allowed to take on any value. The direct gain is represented by the variable G. Since we are interpreting the control problem as being a model of a publicly-owned dynamic firm, the relevant direct gain here is best conceptualized as being a dividend that the firm pays out to its shareholders. G can be expressed as a function of I (for a given level of K) by the equation $G = G(K, I)$. The implied "production possibilities frontier" between G and I, for a given fixed value $K = K(0)$, is depicted as the curve in Fig. 4.

It is important to recognize that here – in this time-autonomous formulation – the firm's "production possibilities frontier" $G = G(K, I)$ does not depend on time explicitly. The time autonomy of the firm's "technology" for trading off dividends against (net) investments for a given level of capital means that all sources of future dividend and investment possibilities are correctly 'accounted for' by changes in capital stocks. It might therefore be said that we have 'complete accounting' here – in the sense that there are no residual forces of growth having their origin in 'unaccounted-for' time-dependent atmospheric changes.

The firm's risk-class-adjusted "opportunity cost of capital" or its "competitive rate of return" is given as the parameter ρ. The firm seeks to maximize, over an infinite horizon, the present discounted value of dividends paid out to its shareholders.

We express the firm's dynamic optimal control problem in a traditional calculus-of-variations form where the control set essentially allows any value of I to be assumed –

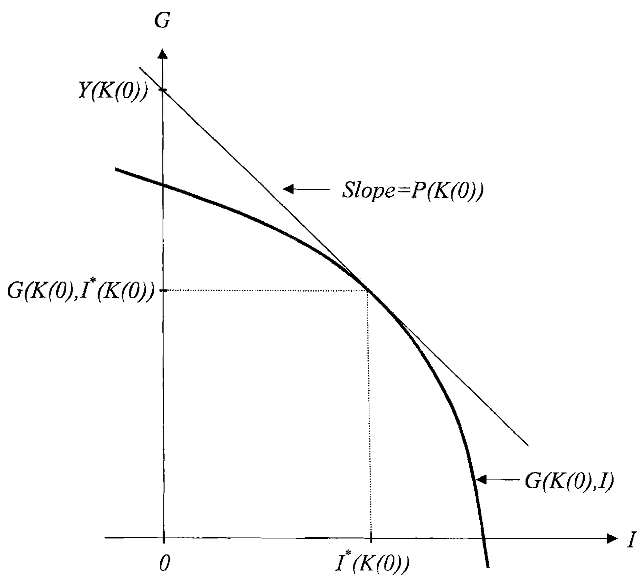

Fig. 4. The firm's true income or earnings.

implying that the firm will not wish to choose "crazy" values of I, even if it can, because of the implicit deleterious effects on G.

Where are we trying to go now with this deterministic model? We are attempting now to rephrase the wealth and income version of the maximum principle in this deterministic model so that the transition to its natural stochastic generalization is made easier and more understandable.

If we try to place it in the finance literature, the wealth and income version of the maximum principle is itself a substantive generalization of the extremely simple 'Gordon model' for calculating the stock market value of a firm with a constant rate of dividend growth. The 'Gordon formula' says that the competitive stock market value of a firm paying dividends (of current value D_0), which will grow exponentially at rate g, is $V = D_0/(\rho - g)$, where ρ is the firm's opportunity cost of capital. This 'Gordon model' can be seen as a very special limiting case of the deterministic wealth and income version of the maximum principle. (In this case, it is as if $G(K, I) = \rho K - I$, $K(0) = D_0/(\rho - g)$, and, in this limiting degenerate case of an optimal control problem where the instrument $I(t)$ is constrained to be equal to $gK(t)$, the efficiency price of investment $p(t)$ is always one.) When we present the *stochastic* wealth and income version of the maximum principle, which we are now leading up to, we will have generalized yet further the scope of a methodology whose simplest conceivable originating example is the 'Gordon model'.

In deterministic control theory it is usual to conceptualize the optimal policy and express the optimality conditions as a function of (the) *time*. What we want to do throughout this paper is to conceptualize the optimal policy and express the optimality conditions as a function of (the) *state*. Let us therefore think of a "policy function" $I(K)$ as expressing the control I as a function of the state variable K. Then an "optimal policy function" $I^*(K)$ expresses the *optimal* control setting as a function of the capital stock K. What exactly do we mean by this notation?

Any policy function $I(K)$ generates a corresponding time trajectory $K(t)$, which satisfies the differential equation

$$\mathrm{d}K(t) = I(K(t))\,\mathrm{d}t, \tag{1}$$

along with the initial condition

$$K(0) = K_0. \tag{2}$$

(For reasons that will soon become apparent, we are using the seemingly arcane notation of (1) here to describe the ordinary differential equation $\dot{K} = I$.) In particular, the *optimal* policy function $I^*(K)$ generates a corresponding time trajectory $K^*(t)$, which also satisfies the relevant versions of (1) and (2), i.e.,

$$\mathrm{d}K^*(t) = I^*(K(t))\,\mathrm{d}t, \tag{3}$$

and

$$K^*(0) = K_0. \tag{4}$$

Waving aside mathematical technicalities and difficulties (as will be our custom throughout this paper), our definition of the *optimal* policy function $I^*(K)$ is that it satisfies

(3), (4) and for *any other* policy function $I(K)$ satisfying (1), (2), it must hold that

$$\int_0^\infty G(K^*(t), I^*(K^*(t)))\, e^{-\rho t}\, dt \geq \int_0^\infty G(K(t), I(K(t)))\, e^{-\rho t}\, dt. \tag{5}$$

Since the calculus of variations form of the problem is a special case of optimal control theory where there are no constraints on the control variable (here net investment), the Hamiltonian formulation works as a special case where the maximized Hamiltonian (with respect to net income) is obtained at an interior solution. (An additional assumption that the control I is constrained to be in a control interval, so that $m(K) \leq I \leq M(K)$, could readily be handled by the model, but we are trying to convey here the basic message as simply and as directly as possible.)

With (5) representing an optimal net investment policy, we can *define* the corresponding accounting price of investment as a function of the capital stock here as

$$P(K) \equiv -G_2(K, I^*(K)). \tag{6}$$

Next *define for all* K_0 the state evaluation function

$$V(K_0) \equiv \int_0^\infty G(K^*(t), I^*(K^*(t)))\, e^{-\rho t}\, dt, \tag{7}$$

where the corresponding trajectories $\{K^*(t)\}$ satisfy conditions (1)–(5) above. In words, $V(K_0)$ is the state evaluation function because it represents the value of an optimal policy expressed parametrically as a function of the initial condition $K(0) = K_0$.

Of course we know from other considerations of optimal control theory that it is *also true* (wherever $V(K)$ is differentiable) that

$$P(K) = V'(K). \tag{8}$$

For reasons that will become apparent presently, we choose here to emphasize the less-customary current-rate-of-transformation accounting-price interpretation of Eq. (6), by making it into a definition of $P(K)$, while Eq. (8) then appears here in the form of a derived result or theorem. (The more customary order, of course, is exactly the opposite – namely, to define $P(K)$ first by (8), and then to derive Eq. (6) as a theorem.) With the primary interpretation being (6), $P(K)$ can then be viewed more conspicuously as a theoretically observable shadow or accounting price representing the current marginal rate of transformation between investments and dividends, when the capital stock is at level K.

In order to motivate the definition of the firm's true income or earnings, it is convenient to invent and apply, in a thought experiment, the fiction of an ''ideal accountant.'' Suppose the firm is currently at the level of capital stock K. This means the firm is located on its production possibilities frontier at the point where $I = I^*(K)$ and where $G = G(K, I^*(K))$, and that the firm is acting *as if it* is using $P(K)$ as an internal shadow price to represent the appropriate current marginal rate of transformation of I into G. It is assumed that the ideal accountant knows all of this information. Thus, the ideal accountant observes the current dividend flow $G(K, I^*(K))$, the current investment flow $I^*(K)$, and also knows the firm's corresponding current internal efficiency price of investment in terms of dividends, $P(K)$.

In this situation, a 'not unreasonable' definition of the firm's true income or true earnings (as a function of its capital stock) is

$$Y(K) \equiv G(K, I^*(K)) + P(K)I^*(K). \tag{9}$$

Formula (9) represents a 'not-unreasonable' definition of income because it is a direct analogue of the intuitive idea that *income is measured as consumption plus the value of net investment, where the accounting price of investment is taken as its opportunity cost in terms of marginal consumption foregone*.

The geometric relationship between G, I, P and Y, for the initial value $K = K(0)$, is depicted in Fig. 4.

The ''wealth and income version'' of the maximum principle applied to this situation is the *theorem* that, under the stated assumptions of the model,

$$\rho V(K) = Y(K), \tag{10}$$

which holds for all K.

Let us now review very carefully the intended operational meaning of theorem (10) as a description of a dynamic optimizing firm whose shares are publicly traded. The variable ρ measures the firm's ''opportunity cost of capital'' or its ''competitive rate of return'', and is assumed to be known. (The simplest example to have in mind here is an environment of risk neutrality, where p represents the available risk-free return-although more general interpretations are possible via so-called ''risk-neutral evaluation.'') The state evaluation function $V(K)$ is observable as the competitive market value of all shares in this firm, when the capital stock is K. Essentially, $V(K)$ must equal this competitive share value because it equals the (maximized) present discounted value of all future dividends that will be paid out by the firm, which a ''share'' entitles the owner to have.

The left hand side of Eq. (10), $\rho V(K)$, represents the flow of return payments that the holders of shares of this firm could expect to obtain on alternative comparable investments made elsewhere in the economy. The right hand side is genuine income, $Y(K)$, as defined by formula (9).

Eq. (10) is a genuine theorem relating two independently measured concepts. There is nothing in the least degree tautological or circular about (10). True earnings are measured by the ideal accountant, who need know absolutely nothing about the stock market value of the firm or its competitive rate of return. (It is completely irrelevant to the ideal accountant even whether the firm is privately held or it issues shares that are publicly traded.) The cost of capital times the stock market value of shares is noted by a stock market observer, who may know absolutely nothing about the true currently accounted earnings of the firm – except insofar as it manifests itself through the share price. The theorem embodied in (10) says that, in the above model, these two independently measured entities must theoretically be equal.

If earnings could be comprehensively and accurately measured, with all investments evaluated at true opportunity costs, then the ratio of true earnings to share valuation should equal exactly the firm's competitive rate of return on capital. All of the information about the future that is embodied in competitive asset pricing is *in principle* also captured by current perfectly–comprehensively accounted income. Thus, we have here a theory of ideal income accounting explaining quite sharply why someone attempting to understand asset pricing might be *very* interested, at least in principle, in examining true economic earnings. The essential idea here is that ideally measured income and wealth are, at least in principle, two independently observable sides of the same coin. While financial economists who specialize in theories of asset pricing occasionally make passing or indirect reference to earnings or income, the centrality of the kind of tight conceptual connection revealed by (10) – which is at least useful as a theoretical organizing principle for thinking about what earnings or income is ideally *supposed to be* measuring or representing – seems not to have been grasped in the finance literature.

Of course Eq. (10) is only a *theoretical* result about the relationship between ideal comprehensively accounted income and wealth. The real world of accounting may actually be very different from the theoretical idealizations of this model. However, the overarching point here is that there *is* a theoretically tight connection between an idealized measure of income and wealth. A tight theoretical relation like this can serve as a valuable starting point for focusing our thinking about some important possible financial (or even welfare) connections between stock values and flow values.

3. Stochastic income and wealth

Thus far, the precise connection between ideally accounted comprehensive income and perfectly competitive share evaluation has been developed for what is essentially a deterministic dynamic firm. But since asset pricing is widely believed to be associated with uncertainty in a genuinely essential way, it will be interesting to see what happens to the relation (10) when the underlying dynamic model of the firm is augmented by uncertainty.

We now follow a long modeling tradition by introducing uncertainty as a stochastic diffusion process appended to the accumulation Eq. (1). (The analytical power, which it can bring to bear on many dynamic optimization problems, is one of the reasons that many modelers like stochastic diffusion processes in the first place.) This particular form of uncertainty is far from innocuous and in a sense will drive the strong results that come out of the model. Modeling uncertainty as a continuous time stochastic diffusion process in the accumulation of capital stocks certainly has ample precedence in the economics and finance literature, and, one might argue, such a formulation at least approximates some realistic situations. As we will show, this formulation in terms of a stochastic diffusion process will turn out to be a good starting point for thinking about the introduction of uncertainty into the basic model-because with this formulation we will still be able to obtain a striking connection, in the spirit of (10), between ideally measured stochastic income and stochastic wealth.

By defining the deterministic problem in the unusually roundabout form that we just did – in terms of policy functions – we have greatly eased our transition to an analysis of the proper stochastic generalization. We now introduce genuine uncertainty by replacing (1) with its stochastic generalization

$$\mathrm{d}K(t) = I(K(t))\,\mathrm{d}t + \sigma(K(t))\,\mathrm{d}Z(t), \tag{11}$$

and (3) by its corresponding stochastic generalization

$$\mathrm{d}K^*(t) = I^*(K^*(t))\,\mathrm{d}t + \sigma(K^*(t))\,\mathrm{d}Z(t). \tag{12}$$

In the above problem, $\mathrm{d}Z$ represents the stochastic differential of a simple Wiener process, while the function $\sigma(K)$ represents the standard deviation per unit time of 'unintended' or 'unexpected' capital accumulation at stock level K. The control variable is the unconstrained net investment level I. In (11), the level of I is chosen at time t by the policy function $I(K(t))$, conditional on observing the state variable $K(t)$. Similar comments apply to (12) for $I^*(K^*(t))$.

Were this a paper whose main theme was about applications of stochastic optimal control theory to economics or finance, then much more effort would be spent motivating and explaining the stochastic diffusion Eq. (11) (or (12)). Even so, it is worth noting that almost all such books for economists on the subject (which I know of) do *not* define rigorously the exact meaning of the underlying stochastic diffusion equations or the exact meaning of an optimum. Nor is the treatment of the dual optimality conditions in such books fully rigorous. Such is the inherent mathematical complexity of a rigorous treatment, that it is far more 'economical' for most economists to rely heavily on intuitively plausible explanations and heuristic proofs. Here, we shortchange even *this* approach to the subject by relying on *other* texts, which are more centered directly on the economics or finance applications of stochastic diffusion processes, to provide such a more-detailed heuristic background explanation of stochastic diffusion processes themselves. As we have noted, even these books typically fall far short of full mathematical rigor, but they at least try more than we are able or willing to do in this paper.

The ultimate purpose of this paper is to develop a stochastic wealth and income version of the maximum principle for the simplest possible one-sector dynamic stochastic model, which generalizes what has been done previously for the deterministic case. For such a purpose, we can rely on others providing the background elements of stochastic diffusion processes, while we concentrate our 'detailed heuristics' here on those mathematical aspects that are particularly relevant for defining the concept of stochastic income and relating it to stochastic wealth.

Stochastic diffusion processes, such as what we are using here in the basic model, are mathematically quite tricky. Such processes are continuous almost everywhere but they fluctuate so violently, per unit time in the limit, that they are differentiable almost nowhere. The *variances* of such processes are of order $\mathrm{d}t$ in time, so that the standard deviations are of order $\sqrt{\mathrm{d}t}$ – which means we must be very careful in taking differentials of stochastic functions because we must consider the impact of the second-order terms of a Taylor series expansion in standard deviations to ensure that we have retained all relevant first-order terms.

Our one modest attempt here at an ''explanation'' of the meaning of (11) (or (12)) is to give an *as-if* heuristic story, which will be useful later in interpreting the concept of stochastic income. In fact, this ''explanation'' is so contradictory from a strictly rigorous mathematical viewpoint that it cannot even be carried out meaningfully. Nevertheless, the *as-if* story is exceedingly useful as a first introduction, because it conveys, albeit in misleadingly simple terms, the ''spirit'' of what remains after the real mathematicians, like

Wiener, Kolmogoroff, Itô, and Stratonovich, have performed their magic with rigorous definitions and proofs.

Suppose we want to conceptualize the diffusion process (11) as the limit of a finite-difference version (with infinitesimally small step sizes), which is analogous to the very useful usual way of thinking about an ordinary differential equation as being the limit of what happens to a finite difference story – in the limit as the step size goes to zero. (*Everyone begins* by intuiting the stochastic diffusion story along these lines, even the great mathematicians who provided a precisely rigorous description in order to "patch up" the parts where the *as-if* story goes wrong.) Let the step size be h, which is to be interpreted as some very small positive number. Then the *as-if* story that Eq. (11) is trying to tell us goes something like this. *Given* $K(t)$ at time t, the situation is *as if* there is a fair-coin-flipping probability of one half that the coin comes up "heads," in which case the realized value of K at time $t + h$ is

$$K^1(t+h) = K(t) + I(K(t))h + \sigma(K(t))\sqrt{h} + \mathrm{O}(h^{3/2}) \tag{13}$$

while the situation is *as if* there is a fair-coin-flipping probability of one half that the coin ends up "tails," in which case the realized value of K at time $t + h$ is

$$K^2(t+h) = K(t) + I(K(t))h - \sigma(K(t))\sqrt{h} + \mathrm{O}(h^{3/2}), \tag{14}$$

where the expression $\mathrm{O}(h^{3/2})$ stands for all terms of order 3/2 or higher in h.

Note the informational timing in the above description, which turns out to be a crucial aspect of the *as-if* story. First, at time t the value $K(t)$ of the state variable is observed. Almost simultaneously, at almost that same time instant of t but just *immediately after* the value $K(t)$ of the state variable is observed, the value of the control variable $I(K(t))$ is chosen. The observation of $K(t)$ and the subsequent choosing of $I(K(t))$ *both occur before* the fair coin is flipped. This "observing and then choosing" at time instant t is subsequently followed by a period of length h, during which the fair coin is flipped and the outcome is observed at time $t + h$. From (13) and (14), the actual realized value $K(t + h)$ of the state variable at time $t + h$ may then be conceptualized, for small enough h, *as if it has taken on the value* $K(t) + I(K(t))h$ *plus or minus the standard deviation* $\sigma(K(t))\sqrt{h}$.

The other appropriate change that must be noted in going over to a stochastic diffusion generalization of a prototype-economic control problem is that the deterministic criterion defining the optimal investment function (5) must be changed to its appropriate expected-value version. As usual here omitting mathematically significant details, our definition of the *optimal* policy function $I^*(K)$ is that it satisfies the stochastic diffusion process (12), (4) and for *any other* policy function $I(K)$ satisfying the stochastic diffusion process (11), (2), it must hold that

$$E\left[\int_0^\infty G(K^*(t), I^*(K^*(t)))\,\mathrm{e}^{-\rho t}\,\mathrm{d}t\right] \geq E\left[\int_0^\infty G(K(t), I(K(t)))\,\mathrm{e}^{-\rho t}\,\mathrm{d}t\right]. \tag{15}$$

The notation $E[\cdot]$ refers to the expectation of the random variable contained within the square brackets. In the case of (15), what is inside the square brackets is a relatively simple example of a so-called *stochastic integral*. For any given realized stochastic trajectory of $\{K(t)\}$, (or of $\{K^*(t)\}$) the integrals within the square brackets of (15) can be understood in the usual (Riemann–Stieltjes) sense. Thus, we are allowed to conceptualize each sample path of realized $\{K(t)\}$ (or $\{K^*(t)\}$) as yielding a corresponding realized value of the stochastic integrals appearing inside the square brackets of (15). Since what is inside the square brackets is, by this procedure, a well-defined random variable, there is no problem with interpreting its expected value in the usual way.

It will be useful to note for later reference that the stochastic version of the 'prototype economic control problem' (15) embodies the stochastic analogue of the idea that all sources of future growth are correctly 'accounted for' via changes in capital stocks. With the stochastic version, however, future capital and capital changes are unknown at the present time. As usual in such problems, it is mathematically trivial here to extend the model to include atmospheric time-dependent stochastic diffusion shocks – provided we are allowed to 'account for' them properly by knowing their correct efficiency prices. (For this mathematically trivial extension of the theory, we merely treat the variable that is being shocked over time 'as if' it were just another capital-stock-like state variable – and then apply the existing time-free theory.)

We now want to make use of a particular application of a famous result from the theory of stochastic diffusion processes, which application we will here call "Itô's Expectation Formula." This result is a specific case of what is known in the stochastic-diffusion literature more generally as "Itô's Lemma." We do not really need the more general form here, and there is no sense cluttering up the presentation of this paper with it. Any book dedicated to stochastic diffusion processes, including books on applications to economics and finance, will contain a full discussion of Itô's Lemma, consistent with the level of abstraction of the book.

What we are calling here "Itô's Expectation Formula" can be seen as a generalization of a right-direction version of Taylor's Theorem. Consider any function $F(K)$ that has everywhere a continuous second derivative. The result we need for this paper, in the notation of this paper, is the following.

Lemma (*Itô's expectation formula*).
For $h > 0$,

$$E[F(K(h))] = F(K(0)) + \{F'(K(0))I(K(0)) + \frac{1}{2}F''(K(0))\sigma^2(K(0))\}h + O(h^{3/2}), \tag{16}$$

where the expression $O(h^{3/2})$ stands for all terms of order 3/2 or higher in h. From staring at (16), *it should become fairly evident the sense in which* (*the right-handed version of*) *Taylor's Theorem about a first-order deterministic approximation can be seen as a particular instance of Itô's expectation formula applied to the special case $\sigma^2(K(0)) = 0$.*

Itô's expectation formula looks very simple, but this is a deception. Behind the scenes, Itô had to define precisely what he meant by a stochastic integral, whose integrator is the outcome of a stochastic diffusion process (here $\{K(t)\}$) having unbounded variation within any time interval $[0, h]$. The issues involved in a rigorous definition of an Itô stochastic integral, which lie just below the surface of (16), are mathematically quite sophisticated. We limit our discussion to an attempt to convey the basic idea verbally of why the timing implicit in an Itô stochastic integral is compatible in an essential way with stochastic optimal control theory. We also try to indicate preliminarily why we may want to express projected stochastic income as the expectation of a differently defined "Stratonovich stochastic integral," even though Itô's is the concept compatible with the information-timing sequence implicit in the kind of stochastic-diffusion optimal-control problem that forms the backbone of this paper. (Whenever we employ such terms as "Itô stochastic integral" or "Stratonovich stochastic integral" throughout this paper, it is essentially for background, motivational, or heuristic purposes not really requiring a rigorous mathematical definition.)

In the deterministic dynamic model of a firm's optimal policy, there was no need to make a distinction between infinite horizon optimal plans completely specified now, without revision, and infinite horizon optimal plans when revision is later allowed – because in the absence of uncertainty no revision is required. But here, in the presence of uncertainty, we specify that an optimal program allows instantaneous updating based on new information and that the "optimality" is with respect to maximizing the present discounted *expected* value of gains given the current state and with future updating allowed. Essentially, we are working with a concept of information and timing that is compatible only with an "Itô stochastic integral" – meaning, heuristically, that in the limiting process defining the integral, all relevant functions are evaluated at the *left endpoint* of time subintervals of form $[t, t + \mathrm{d}t]$. Such kind of stochastic integration concept is completely appropriate for measuring functions in the context of a dynamic decision process of "non-anticipating control" – meaning the firm decides $I(t)$ at time instant t based on observed $K(t)$, but only *after* this decision is made is the value of $\mathrm{d}Z(t)$ revealed and allowed to exert its influence throughout the vanishingly small time interval $(t, t + \mathrm{d}t]$.

Although the treatment of stochastic diffusion processes presented here has been heuristically rudimentary to the point of casualness, it is only fair to warn the reader that some truly deep mathematical issues underlie a genuinely rigorous mathematical formalization. Speaking generally and for the most part, there usually is not a complete need for an economist to plumb down to the full depths of mathematical rigor to understand the essential take-home economic messages that emerge from such stochastic control models in economics or finance. However, it turns out that for the particular application in which we are interested, the treatment of stochastic income will force us to rethink at least one basic issue that goes to the core of what a stochastic diffusion process means as a representation of economic dynamics. This basic issue has to do with some delicate matters of timing and information, involving ex post and ex ante measurement of index numbers, that seem not to have been encountered or noticed so far in other economic applications of stochastic processes, but which come to the fore in trying to define the concept of stochastic income operationally in terms of the thought experiment by which it is supposed to be accounted.

The stochastic control problem represented by (15) is describing a situation where, in the discrete version whose limit it is, the action taken at time t, at the beginning of a (vanishingly small) period of length $\mathrm{d}t$, could depend on the knowledge of the current state $K(t)$, but not on any knowledge of the random future state $K(t + \mathrm{d}t)$. If time is truly continuous, the two times t and $t + \mathrm{d}t$ cohere in the limit as $\mathrm{d}t$ vanishes. To rule out any kind of seeing-ahead or clairvoyance property, the mathematical limiting process must be carefully modeled so as not to allow current choices to depend on information available even the tiniest instant ahead.

The limiting process that rules out clairvoyance, by in effect forcing strategies to be continuous from the left while the uncertainties are continuous from the right, gives rise in a natural way to the construction of an Itô stochastic integral. In fact, this is one major reason why the particular Itô way of constructing a stochastic limiting process is the only concept of stochastic integration ever encountered by most economists. We would not need to spend so much time belaboring the point except that when it comes time to give a proper definition of how to measure "expected stochastic income," it *will not* be based on the interpretation of an Itô stochastic integral of the value of capital accumulated within a (vanishingly small) period of time.

As we will presently argue in more detail, for measuring stochastic income a different kind of stochastic integral than Itô is required, which is based on the "trapezoid rule" of averaging the price values at both endpoints of each tiny time interval, rather than evaluating it, as Itô does, only at the left endpoint of the limiting interval used to define the stochastic integral. Thus, it turns out, one famous kind of stochastic integral, called an Itô integral, is appropriate (behind our heuristics) *for formulating and evaluating the firm's stochastic optimal control problem*, while another, also very well known, competing kind of stochastic integral, called a Stratonovich integral, is more appropriate (behind our heuristics) *for measuring stochastic income*.

We can readily provide an intuitive explanation for why (16) holds, which is at the same level of heuristic description as our treatment of (13) and (14). If we make use of the interpretation in the story that outcomes (13) and (14) are *as-if* determined as a result ("heads" or "tails") of the random flipping of a fair coin, by expanding $F(K(h))$ as a *second order* Taylor series expansion evaluated at time $t = 0$, we can tell a *new as-if story* about *it*. The *new as-if story* told "in the spirit of Itô" is that for the *given* $K(0)$ at time $t = 0$, the situation is *as if* there is a fair-coin-flipping probability of one half (case 1: the coin comes up "heads") that the realized value at time $t = h$ of the stochastic function $F(K(h))$ is

$$F^1(K(h)) = F(K(0)) + F'(K(0))I(K(0))[K_1(h) - K(0)] + \tfrac{1}{2}F''(K(0)) \times [K_1(h) - K(0)]^2 + \mathrm{O}^3, \tag{17}$$

while the situation is *as if* there is a fair-coin-flipping probability of one half (case 2: the coin comes up "tails") that the realized value at time $t = h$ of the stochastic function $F(K(h))$ is

$$F^2(K(h)) = F(K(0)) + F'(K(0))I(K(0))[K_2(h) - K(0)] + \tfrac{1}{2}F''(K(0)) \times [K_2(h) - K(0)]^2 + \mathrm{O}^3, \tag{18}$$

where the notation "O^3" stands for all terms of third or higher order in $[K(h) - K(0)]$.

Combining (17) and (18) into an expected value expression, we then have

$$E[F(K(h))] = F(K(0)) + \tfrac{1}{2}F^1(K(h)) + \tfrac{1}{2}F^2(K(h)) + \mathrm{O}^3. \tag{19}$$

Now substitute from (13) (for $t = 0$) into (17), and from (14) (for $t = 0$) into (18). Then combine the resulting expressions for (17) and (18) in terms of h into the formula (19), cancel all redundant terms, and collect all terms of order 3/2 or higher in h. The resulting equation is (16), the Lemma we want.

We will now reinforce our intuition about the exact timing that is involved in describing the expectation of an Itô stochastic integral by generating heuristically a very important condition that must be met by the optimal policy function $I^*(K)$. This famous condition is known as the "*Hamilton–Jacobi–Bellman equation*," hereafter abbreviated as "*HJB*." Some version of *HJB* shows up in a very high fraction of stochastic control problems, ranging from physics to finance.

We begin our development *of HJB* by constructing the state evaluation function here, in this stochastic setting, as an obvious generalization of the deterministic case. *Define for all* K_0 the state evaluation function

$$V(K_0) \equiv E\left[\int_0^\infty G(K^*(t), I^*(K^*(t)))\,\mathrm{e}^{-\rho t}\,\mathrm{d}t\right], \tag{20}$$

where the corresponding trajectories $\{K^*(t)\}$ satisfy conditions (2), (4), (11), (12), (15). In words, $V(K_0)$ is the state evaluation function because it represents the *expected value* of an optimal policy expressed parametrically as a function of the initial condition $K(0) = K_0$.

We now proceed heuristically from (20) to *HJB*. From the basic dynamic programming principle of optimality, we have for any positive h that (20) can also be expressed recursively as

$$V(K(0)) = E\left[\int_0^h G(K^*(t), I^*(K^*(t)))\,\mathrm{e}^{-\rho t}\,\mathrm{d}t\right] + \mathrm{e}^{-\rho h}\,E[V(K^*(h))]. \tag{21}$$

Combining (15), (20), (21), we have that for any policy function $I(K)$ satisfying (11) and (2), it must hold for all positive h that

$$V(K(0)) \geq E\left[\int_0^h G(K(t), I(K(t)))\,\mathrm{e}^{-\rho t}\,\mathrm{d}t] + \mathrm{e}^{-\rho h}\,E[V(K(h))\right]. \tag{22}$$

In (21), (22) and what follows, the stochastic variables within the square brackets are being understood "in the spirit of Itô." This means, loosely speaking, that for any time instant t the state $K(t)$ is observed and then the policy $I(K(t))$ is chosen at the left end of time instant t *just before* the random variable $\mathrm{d}Z(t)$ becomes known at the right end. "In the spirit of Itô" signifies that the timing of a stochastic process is to be understood in the order that "$\mathrm{d}Z$ moves just ahead" of the state and control variables. The timing/information sequence we are postulating for an Itô process may here be written symbolically as:

$$\begin{aligned} &K(t) \mapsto I(K(t)) \mapsto \mathrm{d}Z(t) \mapsto K(t+\mathrm{d}t)[= K(t) + I(K(t))\mathrm{d}t + \sigma(K(t))dZ(t)] \\ &\quad \mapsto I(K(t+\mathrm{d}t)) \mapsto \mathrm{d}Z(t+dt) \mapsto \cdots \end{aligned} \tag{23}$$

The main point about being "in the spirit of Itô" here is this. *The timing/information convention described by the sequence (23) must be followed whenever expressing any Taylor series expansions of random variables "in the spirit of Itô."*

From the fundamental theorem of the calculus and Taylor's Theorem, the following result is just a particular application here of a condition that must hold for *any* well-defined stochastic integral:

$$E\left[\int_0^h G(K(t), I(K(t)))\,\mathrm{e}^{-\rho t}\,\mathrm{d}t\right] = G(K(0), I(K(0)))h + \mathrm{O}(h^2). \tag{24}$$

We now apply Itô's expectation formula (16) to the state evaluation function $V(K)$. Setting $F(K) = V(K)$ in (16), we obtain

$$E[V(K(h))] = V(K(0)) + \{V'(K(0))I(0) + \tfrac{1}{2}V''(K(0))\sigma^2(K(0))\}h + \mathrm{O}(h^{3/2}). \tag{25}$$

It should be evident that the second-order Taylor series expansion (25) is being carried out and evaluated "in the spirit of Itô."

We know that

$$\mathrm{e}^{-\rho h} = 1 - \rho h + \mathrm{O}(h^2). \tag{26}$$

Combine (26) with (25) and collect all terms of order higher than 3/2 to obtain the expression

$$\begin{aligned}\mathrm{e}^{-\rho h} E[V(K(h))] &= (1 - \rho h)\{V(K(0)) + V'(K(0)I(K(0)) \\ &\quad + \tfrac{1}{2}V''(K(0))\sigma^2(V(K(0)))\}h + \mathrm{O}(h^{3/2}).\end{aligned} \tag{27}$$

Next, plug (27) and (24) into (22). Then expand out the resulting expression and consolidate all terms of order 3/2 or higher. We thereby obtain

$$V \geq V - \rho V h + G h + V' I h + \tfrac{1}{2}V''\sigma^2 h + \mathrm{O}(h^{3/2}), \tag{28}$$

where all functions are evaluated "in the spirit of Itô" at time $t = 0$ when the capital stock is $K(0)$.

Now cancel V from both sides of the inequality (28), divide by positive h, and go to the limit as $h \to 0+$. Then rewrite the resulting expression out fully as

$$\rho V(K(0)) \geq G(K(0), I(K(0))) + V'(K(0))I(K(0)) + \tfrac{1}{2}V''(K(0))\sigma^2(K(0)). \tag{29}$$

The above procedure carried out "in the spirit of Itô" converted the *inequality* (22) into the equivalent *inequality* (29). If we apply exactly the analogous procedure to the *equality* (21), we obtain the equivalent *equality*

$$\rho V(K(0)) = G(K(0), I^*(K(0))) + V'(K(0))I^*(K(0)) + \tfrac{1}{2}V''(K(0))\sigma^2(K(0)). \tag{30}$$

Combining (30) with (29), we then obtain the famous *Hamilton–Jacobi–Bellman equation* (*HJB*) for this situation:

$$\rho V(K(0)) = \max_I [G(K(0), I) + V'(K(0))I] + \tfrac{1}{2}V''(K(0))\sigma^2(K(0)). \tag{31}$$

The question we are now wanting to ask is 'what is the proper *economic interpretation of HJB as a relation between income and wealth*?'. The answer is not immediately obvious.

From (30) (or (31)), the term

$$G(K(0), I^*(K(0))) + V'(K(0))I^*(K(0)) \tag{32}$$

certainly *looks like* true current income at properly accounted prices, but what does the term

$$\tfrac{1}{2}V''(K(0))\sigma^2(K(0)) \tag{33}$$

stand for? The conventional answer is very mechanical. In some technical sense, term (33) represents the expected loss of value from the concavity of the value function – because Jensen's Inequality plays a significant role when the underlying stochastic process is fluctuating so violently per unit time in the limit as a Wiener process does when the time interval is made infinitesimally small.

In this "explanation," the Jensen's-Inequality term (33) is being "explained" by the *as-if* risk aversion of a shadow central planner who "owns" the state evaluation function $V(K)$. For such a story, $V''/2$ stands for the "price of risk," while σ^2 represents the "quantity of risk." But if the uncertainty in the model is firm-specific in the first place, why cannot such "risk" be diversified away by shareholders in the usual way – merely by holding just a *small amount* of this firm's stock in a portfolio? *Why should any kind of "risk adjustment" be necessary here*? In this Jensen's-Inequality way of looking at the world there is a fundamental inconsistency with the basic principles of finance. Surely we can tell a better story for an economic interpretation than this!

The second derivative of the state evaluation function is an endogenously derived construct, rather than an exogenously given primitive. Where do we look in a stochastic economy to find $V''(K(0))$ (or, for that matter, $\sigma^2(K(0))$)? The Jensen's-Inequality story may be technically correct as a pure mathematical description, in some narrow sense, but what does it mean operationally in terms of measurement? Which index number might we have our ideal accountant calculate in order to 'account for' the Jensen's-Inequality term (33)? *What is the economic interpretation of the HJB condition (31) in terms of some well-defined thought experiment linking stochastic income with stochastic wealth*?

We are now ready to begin to confront formally the issue of stating rigorously the stochastic analogue of the deterministic "wealth and income" Eq. (10) and explaining its relation to *HJB*. For convenience in seeing some useful analogies, we employ the same notation and use a conceptually similar apparatus to what was developed to explain the deterministic case, only henceforth it is intended that *all notation refers to the stochastic solution*.

Because the control variable I is being treated as if it could in principle take any value, the solution of the *HJB* Eq. (28) must satisfy the marginalist-interior first-order condition

$$V'(K) = -G_2(K, I^*(K)). \tag{34}$$

Following the methodology of the deterministic case, let us first *define* here the primary relationship to be the accounting price of investment in state K – i.e.,

$$P(K) \equiv -G_2(K, I^*(K)). \tag{35}$$

Then, as with the deterministic case, we are conceptualizing the derived secondary relationship

$$P(K) = V'(K) \tag{36}$$

as being in the form of a theorem obtained by combining the definition (35) with the inference (36). (If the standard deviation depends upon I as well as K, so that we must write $\sigma(K, I)$, it is not possible to define so simply as by using (35) the accounting price of investment goods. Nevertheless, the theory goes through if accounting prices are defined by (36), with the only other modification being that I is then chosen to maximize *expected* immediate-future income, rather than actually realized current income.)

As before in the deterministic case, the variable ρ here measures this stochastic firm's risk-neutral-evaluated opportunity cost of capital or its competitive rate of return, and is assumed to be observable (or at least it can be calculated from applicable financial-market considerations and data). The state evaluation function $V(K)$ is again observable in the stochastic case here as the competitive stock market value of all shares in the firm, when the capital stock is K.

Essentially, $V(K)$ must be the competitive share value by the usual arguments – because it equals the (maximized) expected present discounted value of all (here stochastic) future dividends that will be paid out by the firm, and which a ‘‘share’’ entitles the owner to have. The expression $\rho V(K)$ represents the flow of returns that the holders of shares of this firm could expect to obtain on competitively equivalent alternative investments made elsewhere in the economy. Thus, what is going to appear on the left hand side of the stochastic generalization of Eq. (10) has essentially the same interpretation as in the previous deterministic case – it is interpretable as the firm's cost of capital times the stock market value of its shares, which is recorded by a stock market observer who may know absolutely nothing directly about the true earnings of the firm.

The more challenging issue here is to define rigorously what is the appropriate proxy for expected true earnings on the right hand side of the appropriate stochastic generalization of Eq. (10). We are groping here to find for this simple stochastic model a rigorous definition of the concept of ‘‘projected earnings.’’ We again employ the useful fiction of the ‘‘ideal accountant.’’ Only now, in this stochastic setting, the ideal accountant is endowed with even more powerful abilities and is being asked to take on an even more daunting thought-experimental task.

Ideally measured earnings in the near future are themselves a random variable. Stated heuristically, the ideal accountant will be asked to evaluate true future earnings along all realizable stochastic trajectories throughout a given near-future period – and then to estimate their expected value per unit time (in the limit as the length of the future period goes to zero). Intuitively, therefore, we are searching for a reasonable definition of the expected value of forward-looking future stochastic income over what might be called loosely the ‘‘*immediate future*.’’ The ideal accountant is effectively asking himself: ‘based on current information, what do I *expect* that I will have measured as the true income of this firm over the immediate future?’. The ideal accountant is non-anticipating, in the strict mathematical sense that he is *using* only current information to ‘estimate’ or ‘project’ *expected* true income over the immediate future, which is a legitimate construction to make on the basis of non-anticipating current information.

Let us first begin by fixing the length of the near future interval at h, which is some arbitrarily given positive number that will eventually be made to go to zero in the limit. What we define rigorously as "expected future income" over the time interval (0, h] is intended to be a stochastic generalization of the right hand side of formula (9). Up to this point in the development of the argument we could essentially finesse the issue, but now we must seriously face up to this issue of what exactly we intend to mean here by a stochastic generalization of true income (or earnings). We now pose this issue very sharply in terms of which price index of investment goods to use.

For any chosen value of the parameter λ, where $0 \leq \lambda \leq 1$, define the *weighted-average accounting price* of capital accumulated throughout the interval [0, h] to be the random variable

$$P_\lambda(K^*((h)) \equiv (1-\lambda)P(K(0)) + \lambda P(K^*((h)). \tag{37}$$

What we are now going to ask the ideal accountant to measure as forward-looking *immediate-future expected income* (per unit time, now, at time $t = 0$) is an expected-value generalization of the deterministic case, taking here the form

$$EY(\lambda) \equiv G(K(0), I^*((K(0))) + \lim_{h \to 0+} \frac{1}{h} E[P_\lambda(K^*((h))(K^*(h) - K(0))], \tag{38}$$

where the relevant underlying stochastic process is

$$\mathrm{d}K^*(t) = I^*(K^*(t))\,\mathrm{d}t + \sigma(K^*(t))\,\mathrm{d}Z(t), \tag{39}$$

and $K(0)$ is given as an initial condition.

The critical question before us then becomes very specific: 'what value of λ should we instruct the ideal accountant to use?'. The choice will very much matter for the evaluation of the second term on the right hand side of (38). Everything is now a stochastic random variable in this setup. The efficiency price of investment (in terms of dividends) and the capital being accumulated are both changing stochastically throughout the time interval [0, h]. There is an efficiency price of investment at the *beginning* of the interval, $P(K(0))$, and there is a (different) efficiency price of investment at the *end* of the interval, $P(K(h))$. If $\lambda = 0$, it corresponds to using the *initial* price to evaluate the capital that has been accumulated throughout the interval [0, h]. If $\lambda = 1$, it corresponds to using the *final* price to evaluate the capital that has been accumulated throughout the interval [0, h]. If $\lambda = 1/2$, it corresponds to using the *average* price throughout the interval to evaluate the capital that has been accumulated throughout the interval [0, h]. In a deterministic setup (i.e., $\sigma^2(K(0)) = 0$), it makes no difference what value of λ is chosen in definition (38). However, in a stochastic-diffusion setting, this choice can potentially make a significant difference.

The case $\lambda = 0$ corresponds to making the random variable within the square brackets of (38) be effectively an *Itô stochastic integral of* $P(K(t))\,\mathrm{d}K(t)$ between $t = 0$ and $t = h$ – meaning the price function $P(K)$ is to be understood as being evaluated at the *left endpoint* of the tiny subintervals that, in the limit, define this Itô stochastic integral. The case $\lambda = 1/2$ corresponds to making the random variable within the square brackets of (38) be effectively a *Stratonovich stochastic integral of* $P(K(t))\,\mathrm{d}K(t)$ between $t = 0$ and $t =$ h – meaning the price function $P(K)$ is to be understood as being based upon the "trapezoid rule'' of *averaging* the price function values at both endpoints of the tiny subintervals that, in the limit, define this Stratonovich stochastic integral. In any event, we come back again to the very-specifically posed critical question here: 'which value of λ should we instruct the ideal accountant to use to evaluate the capital that is being accumulated stochastically within the interval?'.

I think that most economists would agree on the answer to this question. We should use some measure of the *average* efficiency price that occurs within the interval – i.e., we should take the *average* of the price that occurs at the beginning of the interval and the price that occurs at the end of the interval. This is essentially a natural extension of the idea that it is best to measure a finite-difference approximation of a Divisia index of the value of capital accumulation by using the average investment price, which corresponds to the midpoint of an interval, for evaluating the amount of capital that has been accumulated during the interval.

It turns out that – with the unboundedly violent fluctuations that occur in realized capital stock changes per tiny subinterval as the tininess of the subinterval approaches zero, which degree of violence is inherent in the nature of a Wiener process – it *matters* at which point of the tiny subinterval we evaluate the stochastically changing price (when we define rigorously, as a limiting process, the stochastic integral implicitly contained within the square brackets of (38)). If we want to base the efficiency price on the ''trapezoid rule'' of having the ideal accountant average the price values at both endpoints of the tiny subintervals, which intuitively seems more plausible than picking either of the extremes of the price at the beginning left endpoint or the price at the ending right endpoint, and which is consistent with best-practice real-world construction of Divisia indexes, then the concept of integration we should be using to evaluate (38) is the Stratonovich stochastic integral corresponding to the case $\lambda = 1/2$. Once we buy into the notion – whether intuitively or based on a formal result from index number theory – that it is better for comparisons to use some *average* of Laspeyres and Paasche price indices, rather than using either price index alone, then by logic we should also buy into the notion of using $\lambda = 1/2$ for measuring expected income.

The Stratonovich stochastic integral of $P(K(t))\,\mathrm{d}K(t)$ corresponding to $\lambda = 1/2$ is *symmetric backwards and forwards*, which is an essential property for a Divisia-like index of real capital accumulation to possess in the setting of a stochastic diffusion process. If, in a thought experiment, we decumulated capital symmetrically by retracing our steps backwards along the same realization of the stochastic trajectory along which we accumulated it, when we arrived back at the initial state we would then want our index of the total value of real net accumulated capital to register *zero*. This is precisely the forwards–backwards symmetry that characterizes the Stratonovich stochastic integral, because the Stratonovich path integral is always zero around any closed loop. However, in the above thought experiment, the path-dependent Itô stochastic integral of $P(K(t))\mathrm{d}K(t)$, corresponding to $\lambda = 0$, would yield some non-zero forwards–backwards round-trip value of net total capital accumulated, when there actually has been zero net total capital accumulated. As we will later show formally, the *expected value* of the difference-from-zero biased-measurement error of such an Itô ($\lambda = 0$) 'round trip evaluation' of net capital accumulated, per unit time, into the immediate future and back, exactly accounts for the Jensen's-Inequality term (30).

Thus, it turns out, our stochastic integral of choice for formulating and evaluating the stochastic optimal control problem is unquestionably the Itô integral, which is most appropriate for *this* particular application – but our stochastic integral of choice for measuring expected true income in (38) is the Stratonovich stochastic integral corresponding to $\lambda = 1/2$, which represents the undeniably appropriate concept for *this* particular application. Such a distinction matters only for the evaluation of the "accumulated-capital" part of expected income – because *both* the amount of capital accumulated *and* the accounting price of capital accumulated are *jointly* changing stochastically, which with a continuous diffusion process must be accounted for instantaneously, even within the tiniest measurement interval.

Therefore, to make a long story short, our mathematically rigorous definition here of the firm's *ideally measured immediate-future expected stochastic income is*

$$EY\left(\frac{1}{2}\right) \equiv G(K(0), I^*(K(0))) + \lim_{h \to 0+} \frac{1}{h} E\left[\frac{P(K(0)) + P(K(h))}{2}(K(h) - K(0))\right], \tag{40}$$

which, it is readily confirmed, corresponds to the special $\lambda = 1/2$ 'Stratonovich-case' of (38).

Because the outcome will hinge on subtle questions of timing in what might be called "stochastic-diffusion index number theory," it has been absolutely crucial here to specify carefully *what* is being measured as the properly accounted net value of capital accumulated over a near-future period, when is it being measured, and *how* it is being measured. To make more vivid the image of a forward-looking trapezoidal accountant working on the right-hand-side of the stochastic wealth and income version of the maximum principle, let us give him a name. Let us call our ideal accountant 'Mr. Strat'. From an economist's standpoint, Mr. Strat is the world's greatest accountant. Mr. Strat is the leading authority on the theory and practice of projecting forward a firm's expected true economic earnings. Expression (40) is the expected immediate-future value of the firm's true income, representing exactly what Mr. Strat *expects* he will be measuring as complete income over the immediate future. In this sense, Mr. Strat's expected immediate-future income measurement (40) can be interpreted as the mathematical formalization, which is appropriate to this stochastic setting, of the idea of a firm's "projected true earnings."

The stochastic generalization of the deterministic wealth-and-income statement (10) is the following theorem, which is the main result of this paper.

Theorem (*Stochastic wealth and income version of the maximum principle*).

$$\rho V(K(0)) = EY(\tfrac{1}{2}). \tag{41}$$

In asserting that expected income is the return on expected wealth, it is directly apparent that (41) *represents a stochastic generalization of* (10). *We will "prove" the stochastic wealth and income version of the maximum principle here in two steps. First, applying Itô's expectation formula to* (38) *and* (37) *will give us a useful general result holding for any* λ. *Second, we will then obtain* (41) *by just plugging into this general result the specific parameter value* λ =1/2.

Define, then, the function

$$F(K) \equiv [(1-\lambda)P(K(0)) + \lambda P(K)][K - K(0)], \tag{42}$$

and note that

$$F(K(0)) = 0. \tag{43}$$

Taking first and second derivatives of $F(K)$ from (42), and then evaluating at, $K = K(0)$, we obtain

$$F'(K(0)) = P(K(0)), \tag{44}$$

and

$$F''(K(0)) = 2\lambda P'(K(0)). \tag{45}$$

Plugging (43), (44), and (45) into the right hand side of (16), we then obtain from Itô's expectation formula the result that

$$E[F(K^*((h))] = \{P(K(0))I^*((K(0)) + \lambda P'(K(0))\sigma^2(K(0))\}h + O(h^{3/2}). \tag{46}$$

Now divide both sides of (46) by positive h and go to the limit $h \to 0+$, thereby obtaining

$$\lim_{h\to 0+} \frac{1}{h} E[F(K^*(h))] = P(K(0))I^*(K(0)) + \lambda P'(K(0))\sigma^2(K(0)). \tag{47}$$

Applying definitions (37) and (42)–(47) then yields

$$\lim_{h\to 0+} \frac{1}{h} E[P_\lambda(K^*(h))(K^*(h) - K(0))] = P(K(0))I^*(K(0)) + \lambda P'(K(0))\sigma^2(K(0)). \tag{48}$$

Applying Eq. (48) to the definition (38) turns the latter expression into

$$EY(\lambda) \equiv G(K(0), I^*(K(0))) + P(K(0))I^*(K(0)) + \lambda P'(K(0))\sigma^2(K(0)). \tag{49}$$

Since (36) must hold as an identity for all K, by differentiating it we obtain

$$P'(K) = V''(K), \tag{50}$$

which must also hold for all K. Now use (36) and (50) to compare (49) with the *HJB* condition (30). We have then just shown that for all λ satisfying $0 \le \lambda \le 1$, the relation must hold that

$$\rho V(K(0)) = EY(\lambda) + (\tfrac{1}{2} - \lambda)V''(K(0))\sigma^2(K(0)). \tag{51}$$

The stochastic wealth and income version of the maximum principle (41) follows immediately from plugging $\lambda = 1/2$ into Eq. (51).

Now all of the pieces should fit. *If the accounting is done correctly* (which includes using the proper price index number corresponding to $\lambda = 1/2$) *expected income is the return on expected wealth*. The key economic interpretation of *HJB* is provided directly by the "stochastic wealth and income version of the maximum principle," which is Eq. (41) (along with the definition (40)).

Using (51), we can analyze formally the important concept of the *measurement error or index-number bias* introduced by various values of λ. To do so, let us examine what is measured as the accounting value of the capital accumulated (as investment) along any round-trip forwards–backwards trajectory. At time $t = 0$, the stochastic trajectory $\{K^*(t)\}$ begins at $K = K(0)$. Throughout the time interval $[0, h]$ the trajectory $\{K^*(t)\}$ grinds out a particular stochastic path realization. At time $t = h$, the stochastic trajectory arrives at $K = K^*(h)$.

Throughout the time interval $[h, 2h]$ let us imagine a hypothetical trajectory, which is exactly the same as the realized trajectory of time interval $[0, h]$, *except that time is running backwards*. For any time t with $0 \leq t \leq h$, the hypothetical trajectory $\{\tilde{K}(h+t)\}$, which we are analyzing as a thought experiment throughout the time interval $[h, 2h]$, is related to the stochastic trajectory $\{K^*(t)\}$, which was actually realized throughout time interval $[0, h]$, by the equation

$$\tilde{K}(h+t) = K^*(h-t). \tag{52}$$

What would be measured by an ideally instructed accountant, given any value of λ, as the value (per unit time) of capital accumulated along such a forwards–backwards-symmetric path in the full time interval $[0, 2h]$? From (52) and the definition (37), it would be

$$\frac{P_\lambda(K^*(h))[K^*(h) - K(0)] + P_{1-\lambda}(K^*(h))[K(0) - K^*(h)]}{2h}. \tag{53}$$

Now it should be quite intuitive that a *truly ideal* stochastic index of the value of capital accumulated throughout the time interval $[0, 2h]$ should register that (53) is exactly *zero*. When zero net capital has *actually* been accumulated, the properly accounted *value* of total net accumulated capital *should also be zero*. It is readily seen that the value of (53) is *zero for* $\lambda = 1/2$ *and is non-zero for any other value of* λ.

The degree of expected measurement error in the sense of immediate-future expected index-number bias is then captured by the limiting expected value of expression (53), which "should" be zero but instead is

$$B(\lambda) \equiv \lim_{h\to 0} \frac{1}{2h}\{P_\lambda(K^*(h))[K^*(h) - K(0)] + P_{1-\lambda}(K^*(h))[K(0) - K^*(h)]\}. \tag{54}$$

Using arguments very similar to what was used to derive (52), it can readily be shown that the definition (54) reduces to

$$B(\lambda) = (\frac{1}{2} - \lambda)V''(K(0))\sigma^2(K(0)). \tag{55}$$

From formula (55), it follows immediately that the Jensen's-Inequality term (33) can be interpreted as representing precisely the expected index-number bias or measurement error (per unit time) from using the "incorrect" Itô-like beginning-of-period price corresponding to $\lambda = 0$ instead of the "correct" Stratonovich-like mid-period average price corresponding to $\lambda = 1/2$. When capital accumulation is "correctly accounted" ($\lambda = 1/2$), there is no measurement bias and correctly accounted expected income is exactly the return on correctly accounted expected wealth.

4. Discussion and conclusion

Note that the stochastic wealth and income version of the maximum principle (41) can readily be interpreted as giving a direct finance-theoretic asset-evaluation interpretation to the Hamilton–Jacobi–Bellman Eq. (31). The result (41) means that, for this dynamic stochastic model of the firm, *HJB* is trying to tell us that the information content of stock market evaluation (times the competitive rate of return) is captured completely by forward-looking expected earnings in the immediate future, provided that these earnings have been ideally accounted.

Notice that it is *not current* true earnings that line up exactly with the current stock market evaluation. That was true for the deterministic case, but it no longer holds in a genuinely stochastic environment where $\sigma^2(K(0)) > 0$. Rather, in this model it is forward-looking expected true earnings in the *immediate future* that must equal the current stock market evaluation. The stock market, so to speak, now knows current true earnings and is already looking ahead to projected true earnings in the immediate future. For this model, the entire informational content of competitive stock market asset evaluation (times ρ) is exactly mirrored in expected true earnings for the immediate future. Therefore, to the extent that the stock market value of a firm is correlated with its expected future true earnings for a stochastic diffusion model like this, in principle all of the correlation is explained by the *immediate* future, and none by periods further ahead than that.

The assumption of ideal comprehensive accounting is very strong here, but it is also able to purchase a very strong focal-point result as a polar case. Perhaps a more balanced way of conceptualizing the possible significance of (the multi-capital version of) this basic result for the theory of asset pricing is in its inverse form as a kind of representation theorem.

Asset pricing theory essentially consists of different versions of representing today's asset prices as the expected value of 'something' tomorrow. The economic interpretation of the Hamilton–Jacobi–Bellman equation, which is given by the stochastic wealth and income version of the maximum principle (41), gives us a new 'something' to look at and ponder. From (a multi-capital version of) (41), the current stock market evaluation of a firm's shares must have a *representation as the immediate-future expectation of income* – or at least of an income-like linearly weighted expression, whose quantities are changes in state variables affecting the firm's ability to pay future dividends, and whose weights for evaluating these changes in capital are shadow accounting prices representing current rates of transformation into dividends. Under ideally complete accounting, this expected-income-like term is exactly the expected immediate-future true earnings of the firm. Otherwise, and in reality of course, this expected-income-like term will be captured by projected immediate-future earnings only to the degree that the firm's accounting system is relatively complete and accurate in being able to assess changes in all those continuously changing capital-like state variables that are relevant for the firm's ultimate dividend-paying ability.

Even in a world of imperfect accounting, such a new representation theorem may be useful for asset pricing theory. Knowing the precise theoretical link between asset value and earned income may lead to new ways of conceptualization and measurement. Also, such a result may serve the theory of income accounting as a kind of a guiding beacon lighting the way toward better practice by providing a rigorous conceptual foundation and by suggesting what activities to include as investments and how best to price them.

For an interpretation of the main result of this model in terms of national income accounting, we need to modify the setting appropriately. The relevant prices of capital or investment goods in the setting of a national economy would be observableas the result of competitive market processes operating both in a static and in a dynamic sense. The static interpretation is the usual one: it is as if every agent is currently optimizing (the agent's static net utility or static net profits) for the given current prices. The dynamic interpretation in the stochastic economy would involve a dynamic competitive *rational-expectations* equilibrium in price uncertainty – meaning, essentially, a self-reinforcing equilibrium in expectations where no agents, even when carrying out their optimal actions, can *expect* to make pure profits over time. (While the details are omitted here, what translates into the 'zero expected pure profits' condition of a rational-expectations decentralized dynamic competitive equilibrium is precisely the wealth and income version of the maximum principle, (41), which is based implicitly on the case $\lambda = 1/2$.)

Even for a hypothetical situation of perfectly comprehensive accounting, in a stochastic setting the national income statistician is recording as current *NNP* some just-observed or just-measured index of the value of current economic activity, rather than the expected value of immediate-future income. Therefore, in the rational expectations competitive dynamic equilibrium of a stochastically evolving economy, current comprehensive *NNP* is an accurate barometer of expected wealth (or welfare) only to the extent that current income is an accurate barometer of expected immediate-future income. A nation's expected future welfare is not completely reflected by *last year's realized NNP*, but, rather, it is mirrored completely accurately in *next year's expected NNP*. The theory is telling us to smooth out the unanticipated shocks of the immediate past, which have distorted last year's recorded *NNP*, by projecting ahead to next year's expected *NNP*. This makes the welfare interpretation of present comprehensive *NNP* somewhat more complicated than in the deterministic case, but at least we understand, in the spirit of the stochastic diffusion model of this paper, what exactly are the theoretical relationships between all of the relevant concepts.

The relative brevity of this paper is illusory because the treatment of stochastic diffusion processes here has been so extremely compressed. Had we really tried carefully or rigorously to explain optimal control theory for stochastic diffusion processes, this would have been a book instead of a paper. While cursory, the treatment of this paper should be sufficient to illustrate how the basic principles may be extended into the domain of uncertainty, at least for some classes of stochastic processes.

5. Bibliographic notes

As was indicated in the text, there are a number of books whose aim is to explain the optimal control of stochastic diffusion processes to economists. Overall, I think the best introduction is contained in Dixit and Pindyck (1994). A slightly more rigorous approach is taken in Merton (1990). Yet another step up in rigor is the treatment in Malliaris and Brock.

Fully rigorous, yet simplified and introductory, treatments of stochastic integrals and stochastic diffusion-differential equations, along with nice discussions of applications, are contained in Gard (1988) and von Weizsäcker and Winkler (1990). These two books also include a good discussion of the modeling issues involved in the "Itô versus Stratonovich" debate. In this connection, the original book by Stratonovich (1968) may profitably be consulted.

A fully rigorous treatment of stochastic-diffusion optimal control theory is developed in the books of Fleming and Rishel (1975) and Krylov (1980).

The literature on the theory of complete accounting under uncertainty is sparse. A pioneering article laying out the basic issues is Aronsson and Löfgren (1995), which is also presented as Chapter 8 in Aronsson, Johansson, and Löfgren (1997). Otherwise, this chapter is reporting new work.