Martingale-Like Behavior of Prices and Interest Rates

by

Christopher A. Sims

 This research was supported by NSF grants SES 8309239 and SES 8112026. This paper is an extensive revision of an earlier one with the same title which appeared as a National Bureau of Economic Research discussion paper in 1980. This paper has had more than its share of referees, and though most of them will be happy to be absolved of blame for what is in the paper, the paper is certainly older and wiser for their efforts.

Price changes for a durable good with small storage costs must, in a frictionless competitive market, be in some sense unpredictable -- or so it seems intuitively. After all, if the good were reliably predicted to rise very rapidly in price, one would think the current price should be bid up by speculators eager to cash in on the predicted capital gains, while if it were reliably predicted to fall rapidly in price owners of the good would sell their holdings to avoid the predicted capital losses. These market reactions to predicted price rises or falls should prevent the occurrence of reliably predicted rises or falls. This intuitive idea has sometimes been formalized as the hypothesis that the price P_t of such a good should be a martingale relative to information observable by market participants, i.e. that, if X_t is data which becomes available at t,

$$E\left[P_{t+s} \middle| X_u, u \le t\right] = P_t, \tag{1}$$

for any $s > 0$.

But careful examination of competitive general equilibrium models of behavior under uncertainty shows, as emphasized by R.E. Lucas, Jr. (1978) and by Stephen F. Leroy (1973), among others, that (1) emerges from such models only under extremely restrictive assumptions. The strategy which empirical researchers have sometimes employed -- testing (1) as an econometric specification to obtain evidence on whether a market is functioning as a frictionless competitive market -- is therefore called in to question.

The discrediting of the martingale hypothesis leaves two issues open. Is it then sheer happenstance that econometric tests of (1) often show it to be close to correct? And is the intuitive notion that speculators should eliminate predictable price changes simply fallacious?

This paper puts forward a definition of short-run martingale-like behavior for a stochastic process. Existing continuous time theoretical models of asset prices make them satisfy the definition, and this paper shows that under fairly general assumptions durable goods prices and interest rates will in a frictionless competitive market show such approximate martingale behavior. More precisely, what is shown is that the linear regression of $P_{t+s}-P_t$ on X_t and lagged X_t's, predicted by

(1) to yield an R^2 of zero, instead has an R^2 converging to zero as s goes to zero. Instead of price changes being unpredictable, price changes over small intervals are very nearly unpredictable. [1]

Thus if one wishes to interpret an econometric test of (1) as a test of the importance of frictions and noncompetitive elements in a market, one ought to carry out the test with s "small". Also, the fact that a price is set in a competitive market with few frictions and seems to fit (1) reasonably well for, say, an s of one week, does not mean that (1) should be expected to work well also for an s of one year.

Two central regularity conditions are required to guarantee the result. One is that the price stochastic process have what we define below to be "smooth prospects", meaning that the expected future path of price from t onward is absolutely continuous on the right at t. We show that if this condition fails simple trading schemes can achieve unbounded profits. The other condition is that information "flow steadily" in the sense defined below. Both of these conditions could in principle be violated in continuous time general equilibrium models, and the paper begins with two examples where they fail, to give some insight into how "irregular" an economy or an asset would have to be for these regularity conditions to fail.

The mathematical apparatus of the more abstract part of this paper is similar to that in some literature on the theory of asset markets in continuous time (such as Harrison and Kreps (1979), Ross(1978), Huang (1982),and Duffie(1983)). However, that literature is concerned with deriving conditions on pricing rules from assumptions on behavior of traders when the number of available securities is limited. In this paper we assume the existence of a pricing rule which creates no incentive to open

[1] In the theory of continuous-time stochastic processes, there is a French term "previsible" which names the property that a stochastic process's value at t can be known with arbitrarily high precision at t-s if s is taken small enough. This term is sometimes translated literally as "previsible" in English-language work on the subject (as in, e.g., Williams (1979)) but sometimes instead as "predictable." We will not use "predictable" and "unpredictable" to mean "previsible" and its opposite. We use the terms sometimes informally as in the footnoted paragraph, and sometimes in a precise sense in the phrase "instantaneously unpredictable" defined below.

markets in certain kinds of contingent claims, or equivalently that a rich array of contingent claims are marketed and arbitrage opportunities do not exist. Furthermore, we make certain existence and continuity assumptions directly -- the discount factor for claims to dollars exists and has certain smoothness properties; prices do change -- without deriving them from maximizing behavior.

In effect, we assume that a competitive market equilibrium exists and has "realistic" characteristics. From these assumptions we derive the conclusion of approximate martingale behavior. The harder problem of deriving existence of equilibrium with realistic price behavior from assumptions about individual behavior is sidestepped.

To follow the argument in detail, the reader will have to be familiar with the theory of measure and integration on general spaces. I try, however, to illustrate the main definitions and conclusions with examples, so a reader can grasp what is being asserted without following the argument in detail.

1. Approximate Martingale Behavior: Instantaneous Unpredictability

The approximate martingale property we will show to hold for durable good prices might naturally be called the local martingale property -- except that the term "local martingale" already has a different technical meaning in the theory of stochastic processes. We will call it instead the property of being "instantaneously unpredictable."

Definition: A process P(t) is *instantaneously unpredictable*, or IU at t, if and only if

$$\frac{\mathbf{E}_t\left[\left(P(t+v)-\mathbf{E}_t P(t+v)\right)^2\right]}{\mathbf{E}_t\left[\left(P(t+v)-P(t)\right)^2\right]} \longrightarrow 1 \quad \text{a.s. as } v \longrightarrow 0\ .[2] \qquad (2)$$

[2]Implicit in this definition is that a process does not have the IU property if the conditional expectations in the definition do not exist a.s. Since asset price processes might well have infinite first or second moments, this is unsatisfactory. We could avoid assuming finite moments by working with "local" instantaneous unpredictability, using "local" in the technical sense of that word in martingale

In words, for an instantaneously unpredictable process prediction error is the dominant component of changes over small intervals. Of course, for a martingale with finite second moments, the ratio in (2) is exactly 1.

Notice that an instantaneously unpredictable process P_t has the property that the regression of $P_{t+s}-P_t$ on any variable observable at t (that is, any variable measurable with respect to F_t) has R^2 approaching zero as s goes to zero. In this sense, econometric tests of the martingale hypothesis should show the hypothesis close to true for small s when P_t is instantaneously unpredictable. The martingale hypothesis, equation (1), will not be exactly true, however, and thus might be rejected by statistical tests even though the regression of price change on observable data has very low R^2.

It might appear that for an instantaneously unpredictable process the chance of (1) being rejected at standard significance levels is smaller for smaller values of s. This is not necessarily so, however, because sample size is likely to be related systematically to s. One plausible situation is that in which there is data available for a fixed historical time span and we test (1) with s equal one month using monthly data and (1) with s equal to one week using weekly data, etc. In this case sample size is inversely proportional to the time interval at which data are measured. For many plausible models of asset prices -- e.g. diffusion processes -- the R^2 of a regression of $P_{t+s}-P_t$ on information available at t goes to zero linearly in s. Since the F statistic for the significance of the regression increases linearly in sample size for a given R^2, this case of sample size inversely proportional to s would imply an F statistic tending neither to zero nor infinity as s decreases to zero. Thus there would be no tendency for fine-time-unit data to reject the martingale null hypothesis more often.

On the other hand, it is probably true that researchers are likely to undertake tests of (1) with shorter historical spans of data when they can find data at fine time units, since there will appear to be so many more degrees of freedom in a given historical span when data are available at fine time units. Thus one might expect to

theory (see Jacod [], p.6). However it seems better first to present the main ideas in this less general framework.

find, in examining existing empirical work for a given type of asset price, that studies using fine time unit data do reject (1) less often than studies using coarser time units.

It may help the reader's understanding of instantaneous unpredictability to consider some examples. All these examples will focus on the case where F_t is simply the sigma field generated by past values of the price process P_t itself -- i.e, where the observable information consists only of past and current values of P_t.

Any process whose paths with probability one have nicely behaved derivatives fails to be instantaneously unpredictable. To see this, observe that

$$P_{t+v}-P_t = v\dot{P}_{t+w} \quad ,$$

where the "." over P indicates time differentiation, for some w between 0 and v. Forecasting P_{t+v} as of date t, we could use

$$P_t + v\dot{P}_t,$$

and the resulting expected squared error would certainly be no smaller than that of the minimum variance forecast, E_tP_{t+v}. Thus we can be sure that

$$\mathbf{E}_t\left[(P_{t+v}-\mathbf{E}_tP_{t+v})^2\right] \leq \mathbf{E}_t\left[(P_{t+v}-P_t-v\dot{P}_t)^2\right] = \mathbf{E}_t\left[v^2(P_{t+w}-P_t)^2\right] .$$

But then expression (2) is less than

$$\frac{\mathbf{E}_t\left[v^2(\dot{P}_{t+w}-\dot{P}_t)^2\right]}{v^2\mathbf{E}_t\dot{P}_t^2} = \frac{\mathbf{E}_t\left[(\dot{P}_{t+w}-\dot{P}_t)^2\right]}{\mathbf{E}_t\dot{P}_{t+w}^2} .$$

If the $\dot{P}_t$ process is nicely continuous [3], then at points where $\dot{P}_t$ is nonzero, this expression converges to zero, not one, as v goes to zero. A verbal paraphrase of what we have just argued is that, for a differentiable P_t, the first-order Taylor expansion in t must provide a forecasting formula whose accuracy becomes high at small time intervals, in the sense that forecast error from using the formula becomes small relative to the change predicted by the formula. But our local unpredictability property (2) is the opposite of this -- an instantaneously

[3]Mean square continuity is not quite enough. A sufficient condition would be $E[\max_{|v|<w}(\dot{P}_{t+v}-\dot{P}_t)^2]$ converging to zero with w.

unpredictable P_t has forecast error growing large relative to predicted change even for the best possible prediction formula.

Any P_t which is generated by a univariate first-order linear stochastic differential equation is instantaneously unpredictable. Suppose it is generated by

$$\dot{P}_t = -aP_t + u_t,$$

where u_t is a continuous-time white noise (derivative of a Wiener process).[4] We can verify (2) by noting

$$P_{t+v} = P_t\, e^{-av} + \int_0^v e^{-as} u_{t+v-s}\, ds$$

$$\mathbf{E}_t P_{t+v} = P_t e^{-av}\ .$$

Substituting the expressions above into (2) gives us

$$\frac{\mathbf{E}_t\left[\left(\int_0^v e^{-as} u_{t+v-s} ds\right)^2\right]}{\mathbf{E}_t\left[\left(P_t(1-e^{-av}) + \int_0^v e^{-as} u_{t+v-s} ds\right)^2\right]} = \frac{\frac{1-e^{-2av}}{2a}}{\frac{1-e^{-2av}}{2a} + P_t^2\left(1-e^{-av}\right)^2}$$

Because the second term in the denominator of this expression is $0(v^2)$, while the term common to numerator and denominator is $0(v)$, the expression as a whole converges to one as v goes to zero. It is easy to see that this argument generalizes to the case of any first-order stochastic differential equation of the form

$$dP_t = s(P_t)u_t + m(P_t)dt\ , \qquad (3)$$

when $s(P_t)$ is non-zero and differentiable and m is continuous. Since much work in modern finance theory begins by assuming that asset prices are diffusion processes, such work implicitly assumes the local unpredictability property.

Multivariate diffusions can have components which satisfy local unpredictability almost nowhere. This is obvious when one recalls that a process satisfying a

[4]Fastidious readers may dislike the practice, common in econometric and engineering literature, of writing "stochastic differential equations" with driving random terms without sample paths. They should reinterpret the previous equation to mean that P is a stationary Gaussian process with spectral density $1/|i\omega+a|^2$, where ω is frequency, or equivalently as asserting that P is stationary with stochastic differential $dP(t)=-aP(t)+dW(t)$, where W is a Wiener process.

univariate second-order stochastic differential equation, which will ordinarily have a continuous derivative process, can be written as one component of a two-dimensional first-order stochastic differential equation, e.g.

$$\begin{aligned} dP_t &= Q_t \\ dQ_t &= -aQ_t + u_t \end{aligned} \quad . \qquad (4)$$

The instantaneous predictability arises because the multivariate analog of $s(P_t)$ in (3) is a singular matrix.

Consider a covariance-stationary, linearly regular process (in the terminology of Rozanov (1967)) with moving average representation $P_t=a*u_t$, where "*" indicates convolution. [5] If P is linear, so that best linear prediction coincides with conditional expectation, then P is instantaneously unpredictable if a_s has a right-limit at s=0 which is non-zero, and is nowhere instantaneously unpredictable if it has a right-limit at zero which is zero.

1. The Flow of Information

We begin by specifying a space Ω of states of the world z and a set F of events, subsets of Ω. F is supposed closed under the taking of complements and countable unions and to contain the null set, making F a sigma-field. The flow of information is described by the collection $\{F_t\}$, where each F_t is a sub-sigma field of F and consists of the events which are verifiable at t -- that is, if an event is in F_t, at dates t and later we know whether the event occurred. We assume that an event verifiable at t is also verifiable later, i.e. that $\{F_t\}$ is increasing, or F_t includes F_s for all $t > s$.

The sequence $\{F_t\}$ describes the way information increases over time. The setup often used in macroeconomic theorizing, where information available at t is taken to be current and past values of some vector of real variables X_t, can be accommodated as a special case of this more general specification. The X_t vector is

[5] Here again fastidious readers might prefer that we be explicit: $P(t)=a*u_t$, with u the derivative of the Wiener process W, means that P and W are jointly Gaussian with $\mathrm{cov}\left(P(t),W(s)\right) = \int_{-\infty}^{s} a(t-v)dv$ for any t,s, or equivalently that $P(t) = \int a(t-v)W(dv)$.

thought of as a stochastic process, implicitly depending on the state z as well as on time. Then the events in F_t are taken to be those of the form

$$\left\{ z \;\middle|\; X(s,z) < a \right\},$$

where a is an arbitrary real vector and s is a number less than t, together with all events which can be generated from these by taking countable unions and complements. F_t can then be thought of as the set of events which can be described in terms of characteristics of the history of X_s for s up through t. We assume that F is right-continuous, i.e. that

$$\bigcap_{s>0} F_{t+s} = F_t \,.$$

This means in particular that if X_t is a process adapted to F_t, i.e. a process such that X_t is F_t-measurable for each t, then if X_t is a martingale it is right-continuous.

3. Postulates about Markets.

We assume that there is a market at each date in claims to dollars at dates in the future contingent on any event verifiable by the payment date. This means there is a discount function, $D(t,s,A,z)$, giving the price at date t of one dollar to be paid at s if A contains z, or, if $s<t$, the value at t of one dollar paid at s if A contains z. The event A must be in F_s, of course, so that the condition for the payoff can be verified at the maturity date.

It is an elementary consequence of the absence of arbitrage opportunities that for any sequence $\{A_j\}$ of disjoint sets in F_s we must have $D(t,s,\cap A_j,z) = \sum_j D(t,s,A_j,z)$. If instead, say, the left side of the equation exceeds the right, an arbitrager can issue a security paying a dollar contingent on A_j, using the proceeds to buy all the securities represented on the right. The returns on the purchased securities will exactly pay off the one he sold, yet he will have a profit at t. (Note, though, that if F_s has infinitely many elements this does implicitly assume that infinite collections of contingent securities exist and can be bought and sold.)

It is also reasonable to suppose that D is always finite. We assume there is a probability measure m defined on F and hence also a restriction of m, called m_t, to F_t for each t. Think of m as being the "market" probability measure on states of the world. We require that if m(A)=0, D(t,s,A,z)=0 for all t,z. That is, claims to a dollar contingent on events of "market" probability zero are valueless. While these conditions would hold if m were the "true" probability measure and agents were rational in the sense that they put value zero on securities which pay off only under conditions with probability zero, there is no requirement that m represent truth. The probability statements we will derive take m as the underlying distribution, but any m which puts probability zero only on sets which the market also gives probability zero will do. Thus if we use a "false" econometric model, it will still show the properties we derive in what follows, so long as the econometric model does not rule out as impossible events the market treats as possible.

These conditions are enough to imply that the Radon-Nikodym theorem (see Munroe, p.196) applies and we can write

$$D(0,s,A,w) = \int_A R(s,z)\ m_s(dz) \quad . \qquad (5)$$

We are using the convention that F_0 consists of Ω and the null set, i.e. that events known at t=0 not to have occurred are ignored. This justifies the implication that D(0,s,A,w) does not depend on w. R(s,z)is some random function measurable with respect to m_s -- i.e. a function whose value is not uncertain given information in F_s.

The right-hand side of (5) is the expectation of R(s,z) times the characteristic function of A. It describes the valuation of a simple kind of random payout at s -- one dollar for z in A, nothing otherwise. A further consequence of arbitrage is that the same type of formula must apply to valuation of more complicated patterns of random returns. Obviously arbitrage rules out the possibility that a payment of Q dollars at s conditional on z being in A has any value other than QD(0,s,A,w). Otherwise an arbitrager could buy or sell Q securities each paying one dollar at s contingent on A, while selling or buying one security paying Q dollars at s contingent on A, making a profit.

This, together with our other assumptions, gives us a formula for valuing random payouts which pay Q_i dollars if z is in A_i, i=1,...,k, where the sets A_i are disjoint. Call this a "simple payout pattern." Suppose a security paying H(z) at s has the property that we can bound H(z) above and below by linear combinations of characteristic functions of finite collections of disjoint sets. Arbitrage should preclude the security from being valued outside the range given by the values of securities with simple payout patterns which bound H. This is enough to guarantee that securities with random payouts will be valued as integrals of their payouts if the payouts are measurable. In particular, an F_s-measurable random payout H(z) at s will be valued at time 0 as

$$\int H(z)R(s,z)\; m_s(dz). \qquad (6)$$

Now consider $D(t,s,\Omega,w)$. Because $D(t,s,\Omega,w)$ is an amount paid at t, it is not uncertain at t, so it is m_t-measurable. Yet from the perspective of a date earlier than t, being paid $D(t,s,\Omega,w)$ at t contingent on w being in a set A in F_t is exactly equivalent to being paid a dollar at s contingent on A. In particular, the random payment at t $D(t,s,\Omega,w)$, will be valued at time zero exactly as a payment of a dollar at s contingent on A. Thus

$$\int_A R(t,w)D(t,s,\Omega,w)m_t(dw) = \int_A R(s,z)m_s(dz) \qquad (7)$$

But (7) means, by elementary properties of conditional expectations, that

$$D(t,s,\Omega,w) = \frac{\mathbf{E}_t\left[R(s\,,z)\right]}{R\,(\,t,w)}\,, \qquad (8)$$

where E_t is expectation conditional on F_t.

What we ordinarily call the v-period discount rate, $[1- D(t,t+v,\Omega,z)]/v$, we will denote by $r_v(t,z)$. We will sometimes suppress the dependence on z, referring to $r_v(t)$. Other v-period interest rates could be defined -- $-\log(D(t,t+v,\Omega,z)]/v$ or $[(1/D(t,t+v,\Omega,z))-1]/v$, for example -- and the arguments below would apply to them with slight modification.

The same logic which leads to (8) from (6) gives us a formula for the spot price at t, P(t), for a security which pays a random return H(z) at some date s>t:

$$P(t) = \frac{\mathbf{E}_t\left[H(z)R(s,z)\right]}{R(t)} . \tag{9}$$

We can also consider securities which pay returns in a continuous flow, with the rate of payment q(t,z) at time t being of course F_t-measurable. If in addition to our requirement that payments conditioned on events of zero probability have zero value we impose the requirement that flow payments over intervals of zero length have zero value, we conclude that the spot price Q(t) at t of a security which pays q(s) at time s for all s>t is given by

$$Q(t) = \frac{\mathbf{E}_t\left[\int_t^{\infty} q(s)R(s)\ ds\right]}{R(t)} . \tag{10}$$

The existence of the formula (10) follows from the same Radon-Nikodym derivative type of argument we used above. That the R function in (10) is the same as that appearing in (9) and (8) requires the additional assumption that the value at time 0 of a payment of H(z) dollars spread evenly over an interval of length v containing t approaches that of H(z) dollars paid in a lump at t as v approaches zero.

4. Steady Information Flow and Smooth Prospects

Note that in (8), (9) and (10) we have three formulas for asset prices, each of which gives the price the form $E_t[Z(t)]$, where Z(t) is a random variable not yet observed at t (i.e., not F_t-measurable). Let us define $\Phi_Z(t,s)=E_t[Z(s)]$. If in (2) we replace X(t) by $E_t[Z(t)]$ we can reduce the expression whose limit is taken in (2) to

$$1 - \frac{\left(\Phi_Z(t,t+v)-\Phi_Z(t,t)\right)^2}{\mathbf{E}_t\left[\Phi_Z(t+v,t+v)^2 - \Phi_Z(t,t+v)^2\right]} \tag{11}$$

The asset price given by $E_t[Z(t)]$ is then instantaneously unpredictable at t if and only if the ratio in the second term of (11) goes to zero with v. The claim in this paper, supported by arguments of various sorts below, is that "reasonable regularity conditions" guarantee that the numerator of this ratio is $0(v^2)$, while the denominator is 0(v), and indeed bounded below by a constant times $|v|$ for small v.

For a random variable X, the rate of change with t in the variance of $E_t[X]$ is a measure of the rate of flow of information about X. If we condition the analysis of information flow at a particular time w, we will have $Var_w[E_tX]$ an increasing function of t, up to the time when X is realized (if there is such a time), after which it is constant at $Var_w[X]$. If u>t, a natural measure of the expected flow of information about X between times t and u is $Var_w[E_uX]-Var_w[E_tX]$. This difference is the reduction in the variance of forecast error for X between t and u. When the limit exists, it is then natural to take

$$\lim_{s\to 0}\left\{\frac{Var_t\left[\mathbf{E}_{t+s}X\right]-Var_t\left[E_tX\right]}{s}\right\}$$

to be the " rate of flow of information about X at t". If information flow is homogeneous over neighboring points of time, this limit will exist, but we can do with a slightly more general condition.

Definition: Information flows steadily about the random variable X at t if and only if a.s.

$$\liminf_{h\to 0}\frac{\mathbf{E}_t\left[\left(E_{t+h}X\right)^2-\left(\mathbf{E}_tX\right)^2\right]}{h} > 0. \qquad (12)$$

If W(t) is a Wiener process, $Var_u[E_t[W(s)]]$ is t-u for s>t>=u, s-u for u<=s<t, and 0 for u>=min(s,t), so the rate of information flow about W(s) at t is 1 whenever t is less than s. More generally we would not expect the rate of information flow at t about the value X(s) of a process at s would be independent of s, but if the process is continuous in an appropriate sense we would expect that the rate of information flow at t would be nearly the same for all X(s) as s varies over a small interval.

Definition: If, when we replace X by X(s) in (12), the liminf is bounded away from zero uniformly in s for s in some interval, we will say that information flows steadily and uniformly at t about X(s) in the interval. It should be clear that we have defined steady information flow in such a way that the following proposition is immediate:

Proposition 1: If information flows steadily about Z(s), uniformly over an interval of s values including t and open on the right, then the denominator in (11), $E_t[\Phi_Z(t+v,t+v)^2-\Phi_Z(t,t+v)^2]$ is bounded below by $c|v|$, for some F_t-measurable random variable c.

If Z(t) were a Gaussian stochastic process and the information structure were that values of a Wiener process W(s) (with white noise derivative e(s)) for all s<t are observable at t, we would be able to write Z(t)=a*e(t), where "*" denotes convolution and a is a square-integrable function on the real line. If a(s)=0 for s<0, then Z(t) is realized at or before t, so that $E_t[Z(t)]=Z(t)$. In this case the rate of information flow about Z(t) at t would be zero. More generally, the rate of information flow about Z(t+h) at t is $a^-(h)^2$, where $a^-(h)$ is the left limit of a at h. There is no reason a cannot be zero for large and small values of its arguments; over intervals of s values where it is zero, $E_{t-s}[Z(t)]$ will not be changing at all. Where $E_{t-s}[Z(t)]$ is changing, that is where information is accruing about Z(t), information will be flowing steadily so long as a is bounded away from zero in a neighborhood of s.

The second property we need is as follows:

Definition: When $\{E_t[Z(t+s)]-E_t[Z(t)]\}/s$ is bounded with probability one as s varies over some interval (0,a), we will say that Z has "smooth prospects" at t.

If $E_t[Z(t)]$ is a price P(t), smooth prospects for Z means that $E_t[P(t+s)]-P(t)$ is 0(s). There are several ways to argue that this condition is "reasonable". One is based on economic behavior, and is given in the following proposition.

Proposition 2: If P_t has finite variance and there is a subset S of the interval (a,b) with positive Lebesgue measure such that

$$\min_{t \text{ in } S} \left\{ \frac{\left| E_t\left[P_{t+v}-P_t\right]\right|}{v} \right\} \xrightarrow[v\to 0]{} \infty$$

and

$$\limsup_{v\to 0}\ \max_{t \text{ in } S} \left\{ \frac{\left(\mathbf{E}_t\left[P_{t+v}-P_t\right]\right)^2}{v} \right\}$$

then with positive probability there are investment strategies which achieve unbounded expected return over (a,b) while keeping the variance of the return bounded.

Proof: Cover the interval (a,b) with disjoint intervals of length d. For each such interval containing a point in S, choose a point t in S in the interval. If $E_t[P_{t+s}-P_t]$ is positive, purchase a dollar's worth of the commodity whose price is given by P_t at t and sell it at t+d. If the expected price change is negative, sell short one dollars' worth of the commodity at t and purchase and deliver the commodity at t+d. As the number of intervals increases with d^{-1}, the condition in the theorem guarantees that the expected total return over (a,b) from this strategy can be made arbitrarily high with positive probability if d is small enough.

Now of course there is a risk to the investment in each of the (t,t+d) intervals in the foregoing argument, and the risk over each such interval does not become small relative to the expected gain, in general. In fact if $Var_t[P(t+v)]$ is 0(v), as we have argued is the usual case, the standard error of the deviation of price from its expectation is $0(v^{.5})$, and grows relative to the expected gain on a single interval as v-->0. However, for non-overlapping intervals, the deviations of prices from their conditional expectations are uncorrelated, and the variance of the total gain is the sum of the variances of the gains over the individual subintervals. Thus our assumptions guarantee that the variance of the return from the investment strategy remains bounded while the expected return is unbounded.

The smooth prospects condition rules out foreseen right-discontinuities, but often all discontinuities are treated as left-discontinuities simply as a convention. (That is, at the time of a jump, the value of the process is taken to be the new, rather than the old, value.) More importantly, the assumption rules out the possibility of foreseen variation in the process failing to shrink as we shrink the time unit. If Z(t)=X(t+v), where X is an Ito process

$$dX(t)=a(t)dt+b(t)dW(t)$$

will have smooth prospects so long as a(t+v) has an expectation, because E_t[a(t+v)] will be the right derivative at t of E_t[Z(t)] at t.

It is not hard to construct processes which do not have smooth prospects. In particular, if we generate the marginal distribution of the paths of an ordinary Wiener process, but shift the information structure so that for all t the new F_t is the old F_t for some fixed v, we will have a process without smooth prospects. Putting it another way, martingales always have smooth prospects because their expected future paths are constant, but if we suppose that their actual future paths are known v periods in advance, then, because their actual paths are nondifferentiable, the resulting process does not have smooth prospects. This example is important, because it shows that the smooth prospects property for a process can not necessarily be checked by examining the process's paths -- two processes with exactly the same marginal distribution of paths may differ on the smooth prospects property.

It is true, though, that for a martingale X(t) the best predictor function based on X(s) itself for $s<t$ yields a right- differentiable predicted path, and this property can be checked from the distribution of paths alone. There are processes for which the best predictor based on past X's alone is not right-differentiable. Suppose X is a Gaussian stationary process, with moving average representation X(t)=a*e(t), with e(t) white noise. Assume also that the information set is just current and past X, or equivalently current and past e. If a is right-differentiable, then X(t) is itself differentiable and a′*e(t) is the right-derivative of E_t[X(t+s)] at s=0, implying smooth prospects. But if, say, a(t)=1 on (0,1) and 0 elsewhere, so that X(t)=W(t)-W(t-1) where W(t) is the Wiener process with derivative e(t), then E_t[X(t+s)]=W(t)-W(t+s-1) for $s<1$. Since W(t+s-1) as a function of s traces out part of the path of a Wiener process, X does not have smooth prospects. Another example is provided by processes with $a(t) = e^{-t}t^h$, with h small.

It seems reasonable, given these examples, to think of the smooth prospects property as justified by Proposition 2, i.e. as a regularity condition imposed directly on the P process, rather than as a regularity condition on Z.

At this point we can assemble our definitions into a straightforward theorem.

Theorem 1: If P(t) is an asset price which can be written in the form $E_t[Z(t)]$, where Z(t) is a stochastic process such that Z(t) is not realized at t, then if

i) information flows steadily at t about Z(s) uniformly over an interval of s values including t and open on the right, and
ii) P has smooth prospects at t,

then P is instantaneously unpredictable at t.

Proof: Follows immediately by combining Proposition 1, the definition of smooth prospects, the definition of instantaneous unpredictability, and the decomposition of the defining formula for instantaneous unpredictability as given in (11).

Theorem 1 contains the main idea of this paper. However it is limited in two ways which justify some further discussion. For one, asset price processes do not always appear to have variances, or even expectations, yet the discussion leading to Theorem 1 is based heavily on the first two moments. This approach allows more elementary methods to be used, but the same ideas can be extended to produce a "local" result, in the terminology of modern martingale theory, which is not tied to existence of moments in the same way. The other limitation is that, though the steady information flow about Z assumption is reasonable, we cannot discuss the kinds of pathology in an economic model which might lead to its being violated without going behind Z, to its numerator (future returns) and denominator (the random discount factor R(t)). Natural second-moment assumptions about numerator and denominator do not translate easily into second-moment restrictions on Z, because Z is a nonlinear function of numerator and denominator. To explore the mapping between regularity conditions on the joint behavior of numerator and denominator of Z and regularity conditions on Z itself, we need to apply Ito's formula, the central tool of stochastic calculus.

For each of the cases we have considered, the expected returns component of Z in the numerator is very naturally taken to have smooth prospects and to have smooth information flow concerning it. The R(t) denominator for Z, however, is almost unrestricted by the arbitrage theory. We will first examine how pathological behavior of R might arise and how it might affect our results under some assumptions which make use of Ito’s formula very convenient and which are very common in the finance literature, but are in fact rather restrictive.

5. The Case of Information Generated by Finitely Many Wiener Processes.

Information structures can be characterized by the martingales defined on them. Economic theoretical models are likely most often to be constructed under the convenient assumption that finitely many Wiener processes generate the information structure. A common assumption in theoretical models is that all processes in the model are Ito processes relative to a finite information vector of Wiener processes, meaning that any process X has differential

$$dX(t) = a(t)dt + b(t)dW(t) ,$$

where a(t) is a scalar process and b(t) a vector process, both adapted to $\{F_t\}$. W(t) is a vector of mutually orthogonal Wiener processes.

For such an X, it is easy to check that X is instantaneously unpredictable at t if b(t) is non-zero almost surely, because the smooth prospects property is implicit in the form of dX(t) -- $(d/dv)E_t[X(t+v)]$ at v=0 is just a(t).

Each of the pricing formulas (8)-(10) has the form

$$P(t) = \frac{H(t)}{R(t)} ,$$

where H(t) is expected discounted returns from t onward and R is the random discount factor. These components, under our current regularity assumptions, will be taken to have the form

$$\begin{aligned} dH(t) &= a_H(t)\ dt + b_H(t)\ dW \\ dR(t) &= a_R(t)\ dt + b_R(t)\ dW \end{aligned} ,$$

where a_R, a_H, b_R, and b_H are all stochastic processes and, by Ito’s formula,

$$dP(t) = \left[- \frac{H}{R^2}a_R - \frac{1}{R^2}b_H b_R' + .5\frac{H}{R^3}b_R b_R'\right]dt + \left[\frac{1}{R}b_H - \frac{H}{R^2}b_R\right]dW(t) .$$

So P will be instantaneously unpredictable unless $b_H = Pb_R$ a.s. Because the H process is the conditional expectation of future returns and future returns as of t are not observed at t, it is natural that b_H is non-zero -- otherwise information would not be flowing about future returns at t. If b_R is non-zero, we are then assured that P is IU.

More generally, we will have b_R non-zero. If $b_H = Pb_R$, and both are nonzero, then H and R are both being driven by the same one-dimensional martingale with differential $b_R dW$. If the equality held a.s. for every t, it would imply that short run changes in R and H are perfectly correlated. Unless the returns H and the discount factor R are jointly singular in this sense, it will not be possible to avoid instantaneous unpredictability in P.

Consider the case where $H(t)=E_t[R(T)]$ and P(t) is therefore the discount factor for term T-t. Though in this case H(t) and R(t) both come from the R process in some sense, it is still unlikely that b^H^ and b_R are linearly dependent. If they were linearly dependent for every t and T, the implication would be that prediction of R can be based on a one-dimensional martingale rather than requiring use of all the elements of the W vector individually.

It may help the reader's intuition to consider the special case where log R and log H are jointly Gaussian processes. This means that

$$d\log H = a_H(t)\, dt + b_H(t)\, dW(t) \quad \text{and}$$

$$d\log R = a_R(t)\, dt + b_R(t)\, dW(t)$$

with a_R and a_H Gaussian processes and b_R, and b_H deterministic, though time-varying, vectors.

This implies

$$dP(t) = \left[a_H - a_R + .5\|b_H - b_R\|^2\right] P(t)dt + P(t)\left[b_H - b_R\right] dW \quad .$$

Instantaneous unpredictability emerges, therefore, unless $b_H = b_R$. As we noted above, $b_H = b_R$ amounts to a kind of singularity in the joint distribution of small changes in H and R, with the small changes being nearly perfectly correlated. While such a result might emerge in a model with a single source of stochastic disturbance, more generally, if there are several kinds of real asset, each with its own random variations in return, it will be impossible for H and R to show such exact collinearity for every security.

Now assume log R is in addition a process with stationary increments and consider the case of a discount bill, i.e. $H(t)=E_t[R(T)]$. The stationarity gives us that a_R is stationary and b_R is a constant. Now if log R has stationary increments, we can write it as

$$\log R(t) = \log R_0(t) + \int_0^t c(s)\, dW(t-s) \,, \qquad (13)$$

where c(s) is a vector-valued function square-integrable over finite intervals and R_0 is a function whose path is known at time 0. From (13) we obtain

$$\begin{aligned} \log E_t[R(T)] &= \log H(t) \\ &= \log R_0(t) + .5\int_0^{T-t} c(s)c(s)' \, ds + \int_{T-t}^{T} c(s)\, dW(T-s)\, ds \,, \qquad (14) \end{aligned}$$

where we are using the lognormality of R(T) conditional on $\mathcal{F}_t$ and applying the formula that if log(x) is distributed as N(m,v), then $E[x]=e^{m+.5v}$.

Assuming c is differentiable, we obtain from (14)

$$\begin{aligned} &d\log H(t) \\ &= \left[(\dot{R}_0(t)/R_0(t)) - .5c(T-t)c(T-t)' + \int_{T-t}^{T} c'(s)dW(T-s)\right] dt + c(T-t)dW(t) \;. \qquad (15) \end{aligned}$$

But this means that $c(T-t)=b_H(t)$ identically, and since $b_H=b_R$, a constant, in the case where instantaneous unpredictability fails, this case entails c(s) being a constant vector. If c(s) is constant, log R is a deterministic function of time plus a Wiener process. It is not hard to check that this means that the term structure of interest rates varies only deterministically with time.

The conclusion is that if log R is a Gaussian process with stationary increments, the process of discount factors for dollars delivered at a fixed date in the future can fail to be instantaneously unpredictable only if there is no uncertainty about future interest rates.

Before proceeding we should reemphasize that in this entire section, by assuming that R has a nice stochastic differential, we have been ignoring the possibility of R's failing to have smooth prospects. We have been exploring only the possibility that, despite steady information flow about future returns, steady information flow about Z might fail because variation in R exactly cancels the effects of information flow on

H(t). Our conclusion is that this seems to occur only with a kind of one-dimensionality in information flow.

6. Semimartingales, Local Martingales.

The previous section's discussion of the case where the information structure is generated by finitely many Wiener processes probably covers the cases most likely to emerge from theoretical economic models, because such information structures are both convenient and quite general. It would be disturbing, however, if results depended strongly on this assumption of convenience. Furthermore, it rules out the possibility of discontinuities in H(t). There is no reason to suppose it to be impossible for information to arrive at an instant, causing H(t) to make a discontinuous jump. If the information structure is generated by a set of continuous martingales, such jumps are impossible. One might suppose that such jumps could disturb the instantaneous unpredictability result. It turns out that they can, but only if information about the size or timing of the jumps flows "non-smoothly" in a certain sense.

To be able to consider information structures in which information arrives discontinuously as well as continuously, and to sidestep the convenient but ad hoc assumption that information is generated by finitely many martingales, we need to introduce some definitions and results following Jacod (1979).

Definition: A "stopping time" is a random variable T with values in the extended real line (i.e., its values are real numbers or infinity) with the property that the condition $T>t$ is verifiable at t, i.e. $\{T>t\}$ is in F_t. [6]

Definition: If T is a stopping time and X a stochastic process adapted to F_t, we denote by X^T the stochastic process defined by

$$X^T(t)=X(t) \text{ if } t<T,\ X^T(t)=X(T) \text{ If } t\geq T.$$

X^T is called "X stopped at T."

[6]This is Lipster and Shiryaev's [1977] definition of a "Markov time." They reserve "stopping time" for an a.s. finite Markov time. To make this definition coincide with Jacod's definition of a stopping time, we must assume that *F* is generated by a class of right-continuous processes with left limits.

Definition: If Q is a property of stochastic processes, we say that the process X has the property Q "locally" if and only if there is a sequence of stopping times T(n) converging to infinity a.s. such that $X^{T(n)}$ has property Q for each T(n).

Local martingales are processes which have the martingale property locally in this sense. Local martingales are important in part because martingales themselves are defined by moment properties. Certain modifications of martingale processes whose paths look like those of martingales thus lose the martingale property. For example, suppose X(t) over the interval [0,3] is a Wiener process with parameter S, that is a continuous martingale with independent increments and

$$E_t\left[\left(X(t+s)-X(t)\right)^2\right] = (t-s)S^2 .$$

If at time t=3 we reset S as $|1/X(3)|$ and then let X evolve with this new S, we have generated a process whose paths will just be those of a martingale whose parameter changes at t=3. But since $|1/X(3)|$, as the absolute value of the inverse of a normal random variable, has infinite expectation, $E_t[X_{t+s}]$ is not defined as a Lebesgue integral for t+s>3 and t<3. This process is a local martingale, though: take T(n) =infinity if X(3)<n, otherwise T(n)=3. This T(n) clearly converges a.s. to infinity, while $X^{T(n)}$ is a martingale for each n.

Note that if X has the property Q locally, and if we have data on a single realization of X for a finite time span (s,t), we could never find evidence against X having property Q by examining the data. For any c>0 we can choose a stopped version of X which does have property Q and whose probability of exactly matching the time path of X itself over (s,t) exceeds (1-c). Hence the observed path of X always has a high probability of exactly matching one from a process with property Q.

Definition: A process X is called a "semimartingale" if it can be written in the form X=M+A, where M is a local martingale and A's paths have bounded variation.

Semimartingales are in some senses quite general. With the usual approach to stochastic integration, they are the most general class of processes with respect to which a stochastic integral is defined. Note, though, that if W is a Wiener process, W(t)-W(t-1) is not a semimartingale, and also that if $a(t)=t^{\cdot 2}e^{-t}$ for t>0, 0 for t<0, then X=a*ε, with ε white noise, is not a semimartingale. Semimartingales are processes such that $E_t[X(t+s)]$ does not behave outrageously as a function of s. In both the foregoing examples $E_t[X(t+s)]$ is almost surely not right differentiable in s at s=0. These two examples are also examples we gave of processes without smooth prospects. A semimartingale can fail to have smooth prospects -- for example A can have jump discontinuities at dates known in advance.

Proposition 2: If X is a semimartingale and if its bounded variation component A can be chosen to have paths a.s. absolutely continuous with Radon-Nikodym t-derivative at t, DA(t), existing a.s. at t and a.s. bounded in absolute value in a neighborhood of t by the F_t-measurable random variable B(t), then X has smooth prospects.

Proof: By construction,

$$\mathbf{E}_t\left[X(t+s)-X(t)\right] = \mathbf{E}_t\left[A(t+s)-A(t)\right] = \mathbf{E}_t\left[\int_0^s DA(t+v)\, dv\right].$$

The conclusion then follows directly from the definition of smooth prospects and the boundedness assumption on DA.

Now consider a vector semimartingale process X which can be written

$$X = A + M^c + M^d, \tag{16}$$

where M^c is a local martingale with a.s. continuous paths, A is continuous and of bounded variation, and M^d is a square integrable local martingale with paths of locally integrable bounded variation, which implies it changes only in discontinuous jumps. Such a decomposition exists and is unique if X is what Jacod calls a "left quasicontinuous special semimartingale" and is locally square-integrable.

We wish to impose sufficient regularity on the three components of X that we can guarantee local smooth prospects for F(X), where F is a twice-differentiable function. A generalized version of Ito's formula (given as Jacod's (3.89)) allows us to check the effect of these regularity conditions. We will not attempt to attain the greatest possible generality in our result. Doing so with respect to the M^d component would raise mathematical difficulties.

Obviously to begin we want to insure that X itself has locally smooth prospects.

Condition I: A in (16) has absolutely continuous paths with Radon-Nikodym derivative a.s. bounded in absolute value in a neighborhood of t by an F_t-measurable random variable.

Definition: If M is a local martingale with continuous paths, its "quadratic variation" $<M,M>$ is defined as the unique stochastic process such that $MM' - <M,M>$ is a local martingale.

Note that M is a column vector and $<M,M>$ a matrix-valued process. $<M,M>$ is increasing in the sense that $<M,M>(t) - <M,M>(s)$ is positive semidefinite for $t>s$.

Condition II: M^c has quadratic variation $<M^c,M^c>$ with Radon-Nikodym derivative V, bounded above in a neighborhood of t by the random matrix $V^+(t)$, which is F_t-measurable.

We will take M^d to behave over small intervals after t like a process which takes a jump distributed with probability measure over R^n (n being the dimension of X) given by q(t,.) and probability per unit time of a jump h(t). That is, if Q(t,s,.) is the conditional probability measure for $M^d(t+s)$ given F_t, we require

Condition III: $D_2Q(t,t,C+X(t)) = h(t)q(t,C)$, for C not containing 0, $D_2Q(t,t,\{0\})=-h(t)$, with h bounded in a neighborhood of t and q in a neighborhood of t concentrating probability in a set of jumps of bounded length, both bounds being F_t-measurable.

Condition III restricts jumps in X to be "unpredictably timed." While condition III could be generalized somewhat, jumps occurring at times known in advance undo the smooth prospects property. Even if H and R each have smooth prospects, if R has a discontinuous martingale component with jumps at dates known in advance, H/R will not have smooth prospects. The problem is that even though the jump in R has zero expectation, it generates a jump in 1/R with nonzero expectation, and this generates a jump with both size and timing known in advance in the A component of X/R.

Proposition 4: If $X=A+M^c+M^d$, with the components satisfying conditions I-III, and if F is a twice-differentiable function, then F(X) locally has smooth prospects.

Proof: For ease of notation, we will introduce the stochastic process J(t), where J(t)'s distribution in $\mathbb{R}^n$ is given by q(t,.). Then we apply Ito's lemma in generalized form (Jacod, 3.89), which tells us that

$$dF(t) \quad = \left(DF'DA+.5*D_{ij}FV_{ij}(t)+h(t)\mathbf{E}_t\left[J(t)\right]\right) \; dt$$

$$+ \; DF'dM^c \; + \; dM_F^d \; , \qquad\qquad (17)$$

where M_F^d is the discrete martingale part of F(X). It is easy to see that our conditions I-III are chosen to make the three components of the coefficient of dt in this expression all well defined and bounded in a neighborhood of t, which by proposition 3 yields the result we sought.

Now we need a regularity condition that will guarantee steady information flow for F(X). If M~c~=0, so that M~d~ is the only martingale component of X, the appropriate regularity conditions become messy and technical. Furthermore, we are interested below really only in a two-dimensional X=(H,R) and in the specific F(X)=H/R. With discrete jumps the only source of information flow about X, regularity conditions can be much weaker if only one F is at issue than if all F's have to be considered. So we will assume M~c~ is present, obtaining an easily expressed but overstrong condition.

Condition IV: V(s) from condition II is bounded below in a neighborhood of t by the $\mathcal{F}_t$-measurable, a.s. positive definite matrix $\bar{V}(t)$.

Theorem 2: If X satisfies conditions I-IV, then for any twice-differentiable function F, F(X) is locally IU at t.

Proof: Refer again to the Ito's formula equation (17) used in Proposition 4. From that proposition we know that

$$\left(\mathbf{E}_t\left[F(X(t+s))\text{-}F(X(t))\right]\right)^2,$$

whose behavior depends only on the "dt" portion of (16), locally goes to zero with s^2. We need now to verify that

$$\mathbf{E}_t\left[\left(F(X(t+s))-F(X(t))\right)^2\right],$$

is locally bounded below by bs for some b>0 for small s. But the "dt" portion of (16) contributes only an $0(s^2)$ term to this expression. The M^c component of (16) contributes an 0(s) term bounded below by $cDF(t)'V^-(t)DF(t)s$ for small s. The discrete martingale term contributes a component which may also be 0(s). Might it cancel the contribution of the M^c component? We know it cannot do so because square-integrable purely discrete-jump martingales must be locally uncorrelated with purely continuous-path martingales (Jacod, (2.27a)), which completes the proof.

Corollary: If the bivariate process (H,R) satisfies I-IV, and if R is locally a.s. bounded away from zero on an interval (t,t+s), with s>0, by a positive $\mathcal{F}_t$-measurable random variable, then the asset price H/R is locally IU at t.

Proof: Just an application of the theorem. The bound on R is needed to ensure that H/R is twice-differentiable over the relevant range.

Theorem 2 shows that the presence of randomly timed jump-discontinuities in H,R will not upset the local IU property if that property would have emerged from the continuous martingale component of H,R alone. We have already noted, though, that the IU property does not really depend on the M^c component being present. To see this, consider the case where H and R satisfy

$$\log H = -at + M_1$$
$$\log R = -bt + M_2 \quad ,$$

where M_1 and M_2 are mutually independent martingales which make jumps of +/- 1 at Poisson-distributed times, with the probability of a jump being c_i per unit time for M_i. Then

$$\mathbf{E}_t\left[\frac{P(t+s)-P(t)}{s}\right] = \mathbf{E}_t\left[\frac{\frac{H(t+s)}{R(t+s)} - \frac{H(t)}{R(t)}}{s}\right]$$
$$\xrightarrow[s\to 0]{} -\frac{a}{R(t)} + b\,\frac{H(t)}{R(t)^2} + .5c_2\left[\frac{H(t)}{R(t)+1} + \frac{H(t)}{R(t)-1}\right] .$$

Thus the smooth prospects property is assured. Furthermore we have

$$\mathbf{E}_t\left[\frac{\left(P(t+s)-P(t)\right)^2}{s}\right] = \mathbf{E}_t\left[\frac{\left\{\frac{H(t+s)}{R(t+s)} - \frac{H(t)}{R(t)}\right\}^2}{s}\right] \xrightarrow[s\to 0]{} \frac{c_1}{R(t)^2} +$$

$$.5c_2\left[\left\{\frac{H(t)}{R(t)+1}\right\}^2 + \left\{\frac{H(t)}{R(t)-1}\right\}^2\right],$$

which together with smooth prospects assures the IU property for P.

7. A Simple Equilibrium Example

The asset valuation formulas we have derived from arbitrage conditions are familiar enough in form. Furthermore Huang [1982] and Duffie [1984] have provided examples of equilibrium models in which asset prices are instantaneously unpredictable. It may nonetheless be worthwhile to present such an example here; the example is chosen both to display the connection of R to behavioral variables and to stay close to the type of model which has recently been common in macroeconomics.

Suppose an economy of infinitely lived, identical individuals, each of whom maximizes

$$\mathbf{E}\left[\int_0^{\infty} \log\left(\frac{C(t)M(t)}{P(t)}\right) e^{-rt}\, dt\right], \qquad (18)$$

where C is consumption, M is money balances, and P is the price level. Each individual faces the constraint

$$P(t)C(t) + D_tM(t) + P(t)D_tK(t) = \log K(t) + h(t), \qquad (19)$$

where "D_t" indicates differentiation with respect to t of the term to the right and h(t) is a stochastic process whose past is included in the information set available to agents at t and whose value is the same for each agent. The government sets monetary policy. We will consider the case where it chooses $D_tM(t)=0$, all t.

The first order conditions for the agents' maximization problem are

$$\frac{e^{-rt}}{C(t)} = P(t)L(t) \qquad (20)$$

$$\frac{e^{-rt}}{M(t)} = -D_t^+L(t) \qquad (21)$$

$$-D_t^+\left[P(t)L(t)\right] = \frac{P(t)L(t)}{K(t)}, \qquad (22)$$

where "$D_t^+F(t)$" is short for "the right derivative of $E_t[F(s)]$ with respect to s at s=t" and L(t) is the random Lagrange multiplier applying to the constraint at t. Equation (20) embodies our asset pricing formulas. It is easy to see that one solution to (20) is

$$P(t)L(t) = \mathbf{E}_t\left[\int_t^{\infty} \frac{P(s)L(s)}{K(s)}\, ds\right] . \qquad (23)$$

Any other solution to (20) differs from this one by a martingale, but if we add a martingale to P(t)L(t) it follows from (18) that there must be a nonzero probability of C(t) exploding as e~rt~ no matter how large t becomes. Such a solution is infeasible. Dividing (21) through by L(t), and noting that the rental on capital at date t in a competitive market would be P(t)/K(t), we see that the equation just states that P(t), which is both the price of capital goods and the price of consumption goods in this one-sector model, is the expected discounted present value of returns on capital, with the discount factor being L(t), the random Lagrange multiplier on the budget constraint.

It remains to check that L(t) is nicely behaved and that information flows steadily about discounted future returns. This model (because of the separability of the utility function) has a convenient dichotomy property: the solution for C and K is determined by the budget constraint with $D_tM=0$ inserted in it, together with equations (20) and (22), which reduce to

$$D_t^+\left[\frac{1}{C(t)}\right] = \frac{\left(\frac{1}{K(t)} - r\right)}{C(t)} . \qquad (24)$$

To find a solution to the system, with mutually consistent stochastic processes for C, K, and h, we can postulate a stationary process for C, choosing it so that $D_t^+[1/C(t)]$ is a convenient function of the information set, then use (24) to solve for a (stationary) K, and finally use the budget constraint (19) to solve for h. If we wish to avoid the conclusion that h has a white noise component, we will have to choose the C process in such a way that, though C itself does not have differentiable paths, $C(t)D_t^+[1/C(t)]$ is differentiable with respect to t. This can be accomplished by, e.g., taking

$$\log C = \left[\frac{D+a}{D^2+aD+ab}\right] g\ w(t) , \qquad (25)$$

where w is a continuous time white noise, the D in (25) is the time-differentiation operator, and the rational function of D in brackets is interpreted as a one-sided convolution operator.

It follows from (23) and (20) that C(t) does not have differentiable paths, unless information is not flowing at t about the integral on the right-hand side of (23). Thus we could not in (25) have chosen the denominator polynomial in D to have degree higher by 2 than the numerator without engendering one of the explosive solutions to (23) we ruled out.

But (25) implies that C has the form

$$dC(t) = C(t)\left[\frac{-ab}{D^2+aD+ab}\, w(t)+.5g^2\right] dt + gC(t)dW(t) , \qquad (26)$$

assuming w is the derivative of W, a standard Wiener process.

Now (21) above determines the Lagrange multiplier, i.e. the discount factor, exactly as

$$L(t) = \frac{e^{-rt}}{rW} . \qquad (27)$$

This follows because any other solution to (21) implies a positive probability of negative L. But then (20) implies

$$P(t) = \frac{Mr}{C(t)} , \qquad (28)$$

i.e. that the price level is inversely proportional to C. Returning to (26), we see that a price level inversely proportional to C will certainly have both smooth prospects and a nonzero Wiener component, hence be IU. We could have obtained the same conclusion by looking at (23), noting that L is deterministic, and concluding that unless information is not flowing about the expected future returns to capital, P will have to be IU.

We could also generate a solution where all the uncertainty concerns jumps. To do this we would replace (25) by, say, the assumption that C jumps at Poisson-distributed times with fixed probability per unit time of a jump. To keep C positive and to ensure that investment and hence output is always well-defined, we have to be cautious in the choice of the jump distribution. One choice which works is to make the jumps go to either $C(t)e^{a-Z(t)}$ or $C(t)(1-e^{a})e^{-Z(t)}$ with equal probabilities, and with $Z(t)=(D+n)^{-1}\log C(t)$. (a and n are both positive real constants.) This makes the expected rate of change of C respond negatively to the current level of C, to preserve stability, but at the same time keeps the expected rate of change from jumping when C jumps. This is critical because otherwise (24) implies that K jumps when C jumps, making investment and hence output undefined. But once we have the C process so defined, we can again solve for K from (24) and P from (28). Again we will have the conclusion that PC is a constant. Because the jump times are not foreseen, the implied P process has smooth prospects; and because there is every time period a chance of a jump, information flows steadily, guaranteeing the IU property for P.

Notice that the instantaneous interest rate, $-D_t^+L(t)/L(t)$, will have a lower unconditional expected value for the version of the model with fixed prices than for the version with fixed M. Also note that the correction for inflation required to obtain the real rate of return on capital form the nominal rate on bonds is sensitive. The proper correction is $D_t^+P(t)/P(t)$, and $D_t^+\log P(t)$ does not work. In fact, outside our special cases of $D_tM(t)=0$ or $D_tP(t)=0$, there is in general no correction for expected inflation which will convert the nominal rate to the real rate of return on capital, because in (20) $D_t^+[P(t)L(t)]/P(t)L(t)$ is not the sum of $D_t^+[P(t)]/P(t)$ and $D_t^+[L(t)]/L(t)$ unless the martingale components of P and L are uncorrelated.

In versions of the model with M fixed, the determinism of L makes nominal interest rates constant at r, so bonds of any sort introduced into the model will have constant prices. It is an exercise left to the reader to verify that if monetary policy allows M to vary to keep P constant, bonds will have nonconstant prices which are IU.

6. Conclusion

There are situations where the IU property will certainly fail, but these are mostly intuitively evident: where, say, information relevant to a security's returns is announced at a date fixed in advance, or where the security's price has been pegged at a constant either by taking it as numeraire or, in a model with neutral money, by the government's pegging it with monetary policy. But if information flows steadily about an asset's returns, and if the price of the asset does not show rapid fluctuations over short intervals which allow unbounded returns with bounded variance over finite intervals, then the asset will have an instantaneously unpredictable price at almost all dates. To undo this conclusion requires either that information flow be "lumpy", invalidating the steady flow of information assumption, or that the discount factor R(t) show discontinuities or nondifferentiability whose form is known in advance. Except for the easily identified exceptional cases, neither the real world nor an analytically manageable economic models is likely to generate security prices which fail to be instantaneously unpredictable.

The notion that asset price changes should be unpredictable in a smoothly functioning competitive market is justified -- as an approximation, when the time unit is taken small enough. Tests of the "perfect market hypothesis" do therefore tell us something about how well a market is functioning. The null hypothesis of unpredictability will never be exactly true, however, so our attention should focus on the explanatory power of the regression, rather than on the classical statistical tests of the null hypothesis.

Furthermore, economic theory shows that the accuracy of the perfect market hypothesis is only a short run phenomenon. It will be a statistical null hypothesis hard to reject, even though asset prices changes may be thoroughly predictable at long time horizons. The frequent success of the hypothesis in statistical tests does not justify imposing it as an exact theory on forecasting or policy models.