

NBER WORKING PAPER SERIES

IMPLEMENTING OPTION PRICING
MODELS WHEN ASSET RETURNS
ARE PREDICTABLE

Andrew W. Lo
Jiang Wang

Working Paper No. 4720

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 1994

We thank Petr Adamek, Lars Hansen, John Heaton, Chi-fu Huang, Ravi Jagannathan, Barbara Jansen, René Stulz, and especially Bruce Grundy and the referee for helpful suggestions, and seminar participants at Boston University, Northwestern University, the Research Triangle Econometrics Workshop, the University of Texas at Austin, the University of Chicago, Washington University, UCLA, the Wharton School, and Yale University for their comments. Financial support from the Research Program in Computational Finance is gratefully acknowledged. A portion of this research was conducted during the first author's tenure as an Alfred P. Sloan Research Fellow. This paper is part of NBER's research program in Asset Pricing. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

IMPLEMENTING OPTION PRICING
MODELS WHEN ASSET RETURNS
ARE PREDICTABLE

ABSTRACT

Option pricing formulas obtained from continuous-time no-arbitrage arguments such as the Black-Scholes formula generally do not depend on the drift term of the underlying asset's diffusion equation. However, the drift is essential for properly implementing such formulas empirically, since the numerical values of the parameters that do appear in the option pricing formula can depend intimately on the drift. In particular, if the underlying asset's returns are predictable, this will influence the theoretical value and the empirical estimate of the diffusion coefficient σ . We develop an adjustment to the Black-Scholes formula that accounts for predictability and show that this adjustment can be important even for small levels of predictability, especially for longer-maturity options. We propose a class of continuous-time linear diffusion processes for asset prices that can capture a wider variety of predictability, and provide several numerical examples that illustrate their importance for pricing options and other derivative assets.

Andrew W. Lo
Sloan School of Management
M.I.T.
Cambridge, MA 02142-1347
and NBER

Jiang Wang
Sloan School of Management
M.I.T.
Cambridge, MA 02142-1347

1 Introduction

There is now a substantial body of evidence in the recent empirical literature which shows that financial asset returns are predictable to some degree.¹ Despite the current lack of consensus as to the sources of such predictability—some attribute it to time-varying expected returns, perhaps due to changes in business conditions, while others argue that predictability is a symptom of inefficient markets or irrational investors—there seems to be a growing consensus that predictability is a genuine feature of many financial asset returns.

In this paper, we investigate the impact of asset return predictability on the prices of an asset's options. A comparison between the polar cases of perfect predictability [certainty] and perfect unpredictability [the random walk] suggests that predictability must have an effect on option prices, although what that effect might be is far from obvious.

However, in the continuous-time no-arbitrage pricing framework of Black and Scholes (1973) and Merton (1973), and in the martingale pricing approach of Cox and Ross (1976) and Harrison and Kreps (1979), option pricing formulas are shown to be functionally independent of the drift of the price process. Since the drift is usually where predictability manifests itself—it is, after all, the conditional expectation of [instantaneous] returns—this seems to imply that predictability is irrelevant for option prices.²

The source of this apparent paradox lies in our attempt to link the properties of finite holding-period returns, e.g., predictability, to the properties of infinitesimal returns, e.g., the instantaneous volatility which determines option prices, without properly fixing the appropriate quantities. In particular, while it is true that changes in predictability arising from the drift cannot affect option prices under the Black-Scholes assumption that the volatility σ of instantaneous returns is fixed, fixing σ implies that the unconditional variance of finite

¹See, for example, Bessembinder and Chan (1992), Bekaert and Hodrick (1992), Campbell and Hamao (1992), Chan (1992), Breen, Glosten, and Jagannathan (1989), Chen (1991), Chen, Roll and Ross (1986), Chopra, Lakonishok, and Ritter (1992), DeBondt and Thaler (1985), Engle, Lilien, and Robbins (1987), Fama and French (1988a, 1988b, 1990), Ferson (1989, 1990), Ferson, Foerster, and Keim (1993), Ferson and Harvey (1991a, 1991b), Ferson, Kandel, and Stambaugh (1987), Gibbons and Ferson (1985), Harvey (1989b), Jegadeesh (1990), Keim and Stambaugh (1986), Lehmann (1990), Lo and MacKinlay (1988, 1990, 1992), and Poterba and Summers (1988).

²Predictability can also manifest itself in the diffusion coefficient, in the form of stochastic volatility with dynamics that depend on predetermined economic factors. However, since predictability is more commonly modeled as part of the conditional mean, we shall focus primarily on the drift.

holding-period returns will change as predictability changes. But since the unconditional variance of returns is usually fixed for any given set of data—for example, the annual standard deviation of the return on the market is typically assumed to be 20 percent—fixing σ and varying predictability can yield counterfactual implications for the data.³

The resolution of this apparent paradox lies in the observation that if we fix the unconditional variance of the “true” [finite holding-period] asset return process, i.e., the data, then as more predictability is introduced via the drift, the population value of the diffusion coefficient must change so as to keep the unconditional variance constant. Therefore, although the option pricing *formula* is unaffected by changes in predictability, option *prices* do change. In this respect, ignoring predictability in the drift is tantamount to a committing a specification error that can lead to incorrect prices just as any other specification error can [see Merton (1976b), for example].

But why should the unconditional variance be fixed? One answer is provided by the fact that the marginal distribution of asset returns is a more fundamental or primitive object than the joint distribution of asset returns and other economic variables. Therefore, a logical sequence of investigation is to first match the marginal distribution of returns, and then focus on the implications for the joint and conditional distributions. This is the approach typically taken in studies of the predictability of asset returns: when regressors are added to or subtracted from a forecasting equation, the conditional moments of returns change—affecting the joint distribution of returns and predictors—but the unconditional moments of the marginal distribution of returns, e.g., mean, variance, skewness, kurtosis, etc., remain the same as long as the data is fixed.⁴

Of course, when choosing among several competing specifications of the data, we hope to select the specification that matches most closely *all* of its properties, i.e., its finite-

³In fact, we argue more generally below that *all* of the unconditional moments of the marginal distribution of returns are “fixed”, in the sense that empirical estimates of values are readily obtained from the data, hence any economic or statistical model of predictability must be calibrated to these values to be of empirical relevance. But there is a compelling reason for focusing first on the unconditional variance of returns: any sensible comparative static analysis of predictability must keep fixed the unconditional variance of the variable to be predicted, since this is the benchmark against which the predictive power of a forecast is to be measured.

⁴This is also the approach taken in the growing “calibration” literature begun by Mehra and Prescott (1985). More recent examples include Abel (1992), Cecchetti et al. (1991), Heaton and Lucas (1992a,b), Kandel and Stambaugh (1988, 1990), and Weil (1989).

dimensional distributions.⁵ But since our most basic understanding of and intuition for the data comes from its marginal distribution, at the very least we shall require that any plausible specification must match the marginal distribution's unconditional moments.⁶ This is tantamount to fixing the mean, variance, skewness, etc. at the "true" values.

Alternatively, from a purely empirical standpoint, the unconditional *sample* moments of the data are fixed at a given point in time since we have only one historical realization of each asset return series. The specification search that we undertake can almost always be viewed as an attempt to fit a statistical model to these fixed sample moments.

By fixing the unconditional moments of the data, we show that changes in predictability generally affect the population value of the diffusion coefficient, and this in turn will affect option prices. However, the particular effect on option prices will depend critically on the particular form of predictability inherent in the drift. For example, if the drift depends only on exogenous time-varying economic factors, then an increase in predictability unambiguously decreases option values. But if the drift also depends upon lagged prices, then an increase in predictability can either increase or decrease option values, depending on the particular specification of the drift.

We derive explicit pricing formulas for options on assets with predictable returns, and show that even small amounts of predictability can have a large impact on option prices, especially for longer-maturity options. For example, under the standard Black-Scholes assumption of a geometric random walk for stock prices, the price of a one-year at-the-money call option on a \$40 stock with a daily return volatility of 2% per day is \$6.908. However, under a trending Ornstein-Uhlenbeck price process—which implies serially-correlated returns—we show that a daily first-order autocorrelation coefficient of -0.20 and a daily return-volatility of 2% per day would yield an arbitrage-free option price of \$7.660, an increase of about 11% [see Section 3.3 and Table 1a]. Of course, the particular adjustment to option prices is wholly determined by the specification of the drift, and we propose several

⁵Although the finite-dimensional distributions do not completely determine a continuous-time stochastic process, for our purposes they shall suffice. More rigorously, the concepts of separability and measurability must be introduced to complete the definition of continuous-time processes—see, for example, Doob (1953, Chapter II.2).

⁶For convenience, we shall refer to the unconditional moments of the marginal distribution of returns as simply the "unconditional moments". These moments are not to be confused with unconditional "co-moments" which are moments of the joint distribution of returns, not of the marginals.

specifications that can account for a broad variety of predictability in asset returns, and illustrate the importance of these adjustments with several numerical examples.

In Section 2 we provide a brief review of the Black-Scholes option pricing model to clarify the role of the drift, and to emphasize the distinction between the data-generating process [DGP] and the “risk-neutralized” process for the underlying asset’s price. The implications of this distinction for option prices are developed in Section 3, where we present an adjustment for the Black-Scholes volatility parameter σ that accounts for the most parsimonious form of predictability: autocorrelation in asset returns. To account for more general forms of predictability, we propose two classes of linear diffusion processes in Sections 4 and 5, the bivariate and multivariate trending Ornstein-Uhlenbeck processes, respectively. In Section 6 we show how the parameters of these predictable alternatives can be estimated with discretely-sampled data by recasting them in state-space form and using the Kalman filter to obtain the likelihood function. We consider several extensions and qualifications in Section 7, and conclude in Section 8.

2 Option Prices and the Drift

Much of the success and growth of the market for options and other derivative assets may be linked to the pricing and hedging techniques pioneered by Black and Scholes (1973) and Merton (1973). The fundamental insight of the Black-Scholes and Merton approach is the dynamic investment strategy in the underlying asset and riskless bonds that replicates the option’s payoff exactly. In particular, if the underlying asset’s price process $P(t)$ satisfies the following stochastic differential equation:

$$dP(t) = \alpha(\cdot)P(t)dt + \sigma P(t)dW \quad (2.1)$$

$$d \log P(t) \equiv dp(t) = \mu(\cdot)dt + \sigma dW \quad (2.2)$$

and trading is frictionless and continuous, then the no-arbitrage condition yields the following restriction on the call option price C :

$$\frac{1}{2}\sigma^2 P^2 \frac{\partial^2 C}{\partial P^2} + rP \frac{\partial C}{\partial P} + \frac{\partial C}{\partial t} = rC \quad (2.3)$$

where r is the instantaneous risk-free rate of return.⁷ Given the two boundary conditions for the call option, $C(P(T), T) = \text{Max}[P(T) - K, 0]$ and $C(0, t) = 0$, there exists a unique solution to the partial differential equation (2.3), the celebrated Black-Scholes formula:

$$C(P(t), t; K, T, r, \sigma) = P(t)\Phi(d_1) - Ke^{-r(T-t)}\Phi(d_2) \quad (2.4)$$

where:

$$d_1 \equiv \frac{\log(P(t)/K) + (r + \frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t}} \quad (2.5)$$

$$d_2 \equiv \frac{\log(P(t)/K) + (r - \frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t}}, \quad (2.6)$$

and $\Phi(\cdot)$ is the standard normal cdf. Although it is well-known that the Black-Scholes formula does not depend on the drift α , it is rarely emphasized that $\alpha(\cdot)$ need not be a constant, as in the case of geometric Brownian motion, but may be an *arbitrary* function of P and other economic variables.⁸ This remarkable fact implies that the Black-Scholes formula is applicable to a wide variety of price processes, processes that exhibit complex patterns of predictability and dependence on other observed and unobserved economic factors [see, for example, the processes described in Sections 4 and 5 below].

The second and more modern approach to pricing options is to construct an equivalent martingale measure, which is always possible if prices are set so that arbitrage opportunities do not exist. Under the equivalent martingale measure all asset prices must follow martingales, thus the price of an option is simply the conditional expectation of its payoff at maturity.

More specifically, the martingale pricing method explicitly exploits the fact that the pricing equation is independent of the drift. Since the drift $\alpha(\cdot)$ of $P(t)$ does not enter into the pricing equation (2.3), for purposes of pricing options it may be set to any arbitrary

⁷That C is a function only of P and t , twice-differentiable in P , and once-differentiable in t are properties that can be derived from the replicating strategy, and need not be assumed *a priori*. See Merton (1973) for further details.

⁸This was first observed by Merton (1973), and is also explicitly acknowledged by Jagannathan (1984) and Grundy (1991), but is far too often overlooked in textbook derivations of the Black-Scholes formula. Of course, $\alpha(\cdot)$ must still satisfy some regularity conditions to ensure the existence of a solution to the stochastic differential equation (2.1).

function without loss of generality (subject to some regularity conditions). In particular, under the equivalent martingale measure in which all asset prices follow martingales, the option's price is simply the present discounted value of its expected payoff at maturity, where the expectation is computed with respect to the *risk-neutralized process* $P^*(t)$:

$$dP^*(t) = rP^*(t)dt + \sigma P^*(t)dW \quad (2.7)$$

$$d \log P^*(t) \equiv dp^*(t) = \left(r - \frac{\sigma^2}{2} \right) dt + \sigma dW . \quad (2.8)$$

Although the risk-neutralized process is not empirically observable, it is nevertheless an extremely convenient specification for evaluating the price of an option on the stock with a DGP given by $P(t)$.

The two approaches to pricing derivative assets show that as long as the diffusion coefficient for the log-price process is a fixed constant σ , then the Black-Scholes formula yields the correct option price regardless of the specification and arguments of the drift. More generally, it may be shown that for any derivative asset which can be priced purely by arbitrage, and where the underlying asset's log-price dynamics is described by an Itô diffusion with constant diffusion coefficient, the derivative pricing formula is functionally independent of the drift, and is determined purely by the diffusion coefficient and the contract specifications of the derivative asset.

But the fact that the drift plays no role in determining a derivative's pricing formula belies its importance in the formula's implementation. Because the risk-neutral distribution and the true distribution of the DGP are linked, predictability *can* and generally *does* have an influence on the pricing of derivative assets, despite the fact that only the parameters of the risk-neutral distribution appear in derivative pricing formulas. As we show in the next sections, when we change the true distribution of the DGP, e.g., change predictability, the risk-neutral distribution also changes, and these changes will generally affect the prices of the corresponding derivative assets.

3 Predictability and the Black-Scholes Formula

Although the same symbol σ is used in both the risk-neutralized process P^* and the DGP P , both the theoretical value and the empirical estimate of σ are determined solely by the DGP, not by the risk-neutralized process, and both will be affected by the functional form of the drift. Predictability in the drift can be safely ignored when *deriving* the option pricing formula, but we shall argue below that it must be addressed explicitly for any given DGP.

In Section 3.1, we consider the most parsimonious form of predictability—autocorrelated asset returns—and show how it affects σ directly in the specific case of a trending Ornstein-Uhlenbeck process for log-prices. We also provide a simple adjustment to the Black-Scholes formula that can account for it. More general and empirically plausible log-price processes, with considerably more flexible forms of predictability, are presented in Sections 4 and 5.

3.1 The Trending O-U Process

In distinguishing between the risk-neutral and true distributions of an option's underlying asset return process, Grundy (1991, p. 1049) observes that the Black-Scholes formula still holds for an Ornstein-Uhlenbeck log-price process, and we shall begin with a slight generalization of his example to illustrate the link between predictability and option prices. Specifically, let the log-price process $p(t)$ satisfy the following stochastic differential equation:

$$dp(t) = \left(-\gamma(p(t) - \mu t) + \mu \right) dt + \sigma dW \quad (3.1)$$

$$\text{where } \gamma \geq 0, p(0) = p_0, t \in [0, \infty).$$

Unlike the geometric Brownian motion of the original 1973 derivation of the option pricing model by Black and Scholes, which implies that log-prices follow an arithmetic random walk with independently and identically distributed Gaussian increments, this log-price process is the sum of a zero-mean stationary autoregressive Gaussian process—an Ornstein-Uhlenbeck process—and a deterministic linear trend, so we call this the “trending O-U” process. Rewriting (3.1) as:

$$d(p(t) - \mu t) = -\gamma(p(t) - \mu t) dt + \sigma dW \quad (3.2)$$

shows that when $p(t)$ deviates from its trend μt , it is pulled back at a rate proportional to its deviation where γ is the "speed of adjustment".⁹ For notational convenience, we shall work with the detrended log-price process $q(t)$ for the remainder of this paper where $q(t) \equiv p(t) - \mu t$. From (3.2), we have:

$$dq(t) = -\gamma q(t)dt + \sigma dW \quad (3.3)$$

and $q(0) = q_0 = p_0$.

To develop further intuition for the properties of (3.3), consider its explicit solution:

$$q(t) = e^{-\gamma t} q_0 + \sigma \int_0^t e^{-\gamma(t-s)} dW(s) \quad (3.4)$$

from which we can obtain the unconditional moments and co-moments of continuously-compounded τ -period returns $r_\tau(t) \equiv p(t) - p(t-\tau) = \mu\tau + q(t) - q(t-\tau)$:¹⁰

$$E[r_\tau(t)] = \mu\tau \quad (3.5)$$

$$\text{Var}[r_\tau(t)] = \frac{\sigma^2}{\gamma} [1 - e^{-\gamma\tau}] \quad , \quad \tau \geq 0 \quad (3.6)$$

$$\text{Cov}[r_\tau(t_1), r_\tau(t_2)] = -\frac{\sigma^2}{2\gamma} e^{-\gamma(t_2-t_1-\tau)} [1 - e^{-\gamma\tau}]^2 \quad , \quad t_1 + \tau \leq t_2 \quad (3.7)$$

$$\text{Corr}[r_\tau(t), r_\tau(t+\tau)] \equiv \rho_\tau(1) = -\frac{1}{2} [1 - e^{-\gamma\tau}] \quad (3.8)$$

Since (3.1) is a Gaussian process, the moments (3.5) - (3.7) completely characterize the finite-dimensional distributions of $r_\tau(t)$.

⁹Note that $\gamma > 0$ ensures the stationarity of $p(t)$.

¹⁰Since we have conditioned on $q(0) = q_0$ in defining the detrended log-price process, we must be more precise about what we mean by an "unconditional" moment. If q_0 is assumed to be stochastic and drawn from its stationary distribution, then an unconditional moment of a function of $q(t)$ may be defined as the expectation of the corresponding conditional moment [conditional upon q_0], where the expectation is taken with respect to the stationary distribution of q_0 . Alternatively, if $q(t)$ is stationary, as it is in (3.3), the unconditional moment may be defined as the limit of the corresponding conditional moment as t increases without bound. We shall adopt this definition of an unconditional moment throughout the remainder of the paper. In deriving the unconditional moments of τ -period returns, we have assumed that $\gamma > 0$. In the special case that $\gamma = 0$, the moments can be obtained by simply taking the limit $\gamma \rightarrow 0$ in (3.5) - (3.8).

Unlike the arithmetic Brownian motion or random walk which is nonstationary and often said to be “difference-stationary” or a “stochastic trend”, the trending O-U process is said to be “trend-stationary” since its deviations from trend follow a stationary process. An implication of trend-stationarity is that the unconditional variance of τ -period returns has a finite limit as τ increases without bound—in this case $\frac{\sigma^2}{\gamma}$ —in contrast to the case of a random walk in which the unconditional variance increases linearly with τ . This difference between the trending O-U process and the random walk also exists for the conditional variance of returns. In fact, for the trending O-U process we have

$$\text{Var}[r_\tau(t)|p(t)] = \frac{\sigma^2}{2\gamma} [1 - e^{-2\gamma\tau}] \quad , \quad \gamma > 0 \quad (3.9)$$

which does not equal $\sigma^2\tau$. The conditional variance of τ -period returns also has a finite limit (when $\gamma > 0$) as $\tau \rightarrow \infty$.¹¹

While trend-stationary processes are often simpler to estimate, they have been criticized as unrealistic models of financial asset prices since they do not accord well with the common intuition that longer-horizon asset returns exhibit more risk, or that price forecasts exhibit more uncertainty as the forecast horizon grows. However, if the source of such intuition is empirical observation, it may well be consistent with trend-stationarity since it is now well known that for any finite set of data, trend-stationarity and difference-stationarity are virtually indistinguishable [see, for example, Campbell and Perron (1991) and the many other “unit root” papers cited in their references]. Nevertheless, in Section 5 we shall provide a generalization of the trending O-U process that contains stochastic trends, in which case the variance of returns will increase with the holding period τ .

Note that the first-order autocorrelation (3.8) of the trending O-U increments is always less than or equal to zero, bounded below by $-\frac{1}{2}$, and approaches $-\frac{1}{2}$ as τ increases without bound. These shall prove to be serious restrictions for many empirical applications, and will motivate the alternative processes introduced in Sections 4 and 5, which have considerably more flexible autocorrelation functions. However, as an illustration of the impact of serial correlation on option prices the trending O-U process is ideal.

¹¹Observe that the conditional variance (3.9) is conditioned on only the current price $p(t)$. Of course, more general conditioning information can be used, with potentially different implications. We discuss this further in Section 4.2.

3.2 Relating Unconditional Moments to Parameters

Despite the differences between the trending O-U process and an arithmetic Brownian motion, Grundy (1991) points out that both data-generating processes yield the same risk-neutralized price process (2.7), hence the Black-Scholes formula still applies to options on stocks with log-price dynamics given by (3.1). This may seem paradoxical, especially since the Black-Scholes formula is independent of the parameter γ which determines the degree of autocorrelation in returns. After all, autocorrelation is a simple form of predictability, and we have argued in the introduction that predictability should have some impact on option prices.

The paradox is readily resolved by observing that the two data-generating processes (2.2) and (3.1) must fit the same price data—they are, after all, two competing specifications of a single price process, the “true” DGP.¹² Therefore, in the presence of autocorrelation, (3.1), the *numerical value* for the Black-Scholes input σ will be different than in the case of no autocorrelation, (2.2).

To be concrete, denote by \bar{r}_τ , $s^2(r_\tau)$, and $\rho_\tau(1)$ the unconditional mean, variance, and first-order autocorrelation of $r_\tau(t)$, respectively, which may be defined without reference to any particular data-generating process.¹² Moreover, the numerical values of these quantities may also be fixed without reference to any particular data-generating process. All competing specifications for the true data-generating process must match these moments at the very least to be plausible descriptions of that data [of course, the best specification is one that matches *all* the moments, in which case the true data-generating process will have been discovered]. For the arithmetic Brownian motion, this implies that the parameters (μ, σ^2) must satisfy the following relations:

$$\bar{r}_\tau = \mu\tau \quad (3.10)$$

$$s^2(r_\tau) = \sigma^2\tau \quad (3.11)$$

$$\rho_\tau(1) = 0. \quad (3.12)$$

¹²Of course, it must be assumed that the moments exist. However, even if they do not, a similar but more involved argument may be based on location, scale, and association parameters.

From (3.11), we obtain the well-known result that the Black-Scholes input σ^2 may be estimated by the sample variance of continuously-compounded returns r_τ . However, in the case of the trending O-U process, the parameters (μ, γ, σ^2) must satisfy:

$$\bar{r}_\tau = \mu\tau \quad (3.13)$$

$$s^2(r_\tau) = \frac{\sigma^2}{\gamma} [1 - e^{-\gamma\tau}] \quad , \quad \tau \geq 0 \quad (3.14)$$

$$\rho_\tau(1) = -\frac{1}{2} [1 - e^{-\gamma\tau}] \quad (3.15)$$

Observe that these relations must hold for the *theoretical* or *population* values of the parameters if the trending O-U process is to be a plausible description of the DGP. Moreover, while (3.13) – (3.15) involve population values of the parameters, they also have implications for estimation. In particular, under the trending O-U specification, the sample variance of continuously-compounded returns is clearly *not* an appropriate estimator for σ^2 .

Holding the unconditional variance of returns fixed, the particular value of σ^2 now depends on γ . Solving (3.14) and (3.15) for γ and σ^2 yields:

$$\gamma = -\frac{1}{\tau} \log(1 + 2\rho_\tau(1)) \quad (3.16)$$

$$\sigma^2 = s^2(r_\tau) \gamma (1 - e^{-\gamma\tau})^{-1} = \frac{s^2(r_\tau)}{\tau} \cdot [\gamma\tau (1 - e^{-\gamma\tau})^{-1}] \quad (3.17)$$

which shows the dependence of σ^2 on γ explicitly.

In the second equation of (3.17), σ^2 has been re-expressed as the product of two terms: the first is the standard Black-Scholes input under the assumption that arithmetic Brownian motion is the data-generating process, and the second term is an adjustment factor required by the trending O-U specification. Since this adjustment factor is an increasing function of γ , as returns become more highly (negatively) autocorrelated, options on the stock will become more valuable *ceteris paribus*. Specifically, substituting (3.16) into (3.17) and simplifying yields σ^2 as an explicit function of $\rho_\tau(1)$:

$$\sigma^2 = \frac{s^2(r_\tau)}{\tau} \cdot \frac{\log(1 + 2\rho_\tau(1))}{2\rho_\tau(1)} \quad , \quad \rho_\tau(1) \in (-\frac{1}{2}, 0] \quad (3.18)$$

where the restriction that $\rho_r(1) \in (-\frac{1}{2}, 0]$ is equivalent to the restriction that $\gamma \geq 0$.

More generally, suppose that returns of one holding period τ_1 are used to obtain the unconditional variance $s^2(r_{\tau_1})$, and returns of another holding period τ_2 are used to obtain the first-order autocorrelation coefficient $\rho_{\tau_2}(1)$. Since the data-generating process is defined in continuous time, this poses no problems for deriving the restrictions on the parameters (μ, γ, σ^2) , and manipulating those restrictions yields the following version of (3.18):

$$\sigma^2 = \frac{s^2(r_{\tau_1})}{\tau_1} \cdot \frac{\tau_1}{\tau_2} \cdot \frac{\log(1 + 2\rho_{\tau_2}(1))}{[1 + 2\rho_{\tau_2}(1)]^{\tau_1/\tau_2} - 1} \quad (3.19)$$

$$= \frac{s^2(r_{\tau_1})}{\tau_1} \cdot \Delta(\tau_1, \tau_2, \rho_{\tau_2}(1)) \quad (3.20)$$

$$\Delta(\tau_1, \tau_2, \rho_{\tau_2}(1)) \equiv \frac{\tau_1}{\tau_2} \cdot \frac{\log(1 + 2\rho_{\tau_2}(1))}{[1 + 2\rho_{\tau_2}(1)]^{\tau_1/\tau_2} - 1}, \quad \rho_{\tau_2}(1) \in (-\frac{1}{2}, 0]. \quad (3.21)$$

Without loss of generality and as a convenient normalization, let $\tau_1 = 1$ and $\tau_2 = \tau$ so that the first-order autocorrelation coefficient $\rho_r(1)$ is defined over the holding period τ , which in turn is measured in units of the holding period used to measure the unconditional variance of returns $s^2(r_1)$, thus:

$$\sigma^2 = \frac{s^2(r_1)}{\tau} \cdot \frac{\log(1 + 2\rho_r(1))}{([1 + 2\rho_r(1)]^{1/\tau} - 1)}, \quad \rho_r(1) \in (-\frac{1}{2}, 0]. \quad (3.22)$$

This expression provides a simple adjustment for the Black-Scholes input σ^2 using the unconditional variance $s^2(r_1)$ of returns sampled at unit intervals, and the first-order autocorrelation $\rho_r(1)$ of returns sampled at τ -intervals.

Returning to the simpler relation (3.18) between σ^2 and the first-order autocorrelation coefficient, and holding fixed the unconditional variance of returns $s^2(r_\tau)$, observe that the value of σ^2 increases without bound as the absolute value of the autocorrelation increases from 0 to $\frac{1}{2}$.¹³ This implies that a specification error in the dynamics of $p(t)$ can have dramatic consequences for pricing options. We shall quantify the magnitudes of such consequences in

¹³We focus on the absolute value of the autocorrelation to avoid confusion in making comparisons between results for negatively autocorrelated and positively autocorrelated asset returns. For example, whereas in this case an increase in the absolute value of autocorrelation increases the option's value, in Section 4.2 we provide an example of a positively autocorrelated asset return process for which an increase in autocorrelation decreases the option's value. These two cases are indeed polar opposites, and for important reasons. But without focusing on the absolute value of the autocorrelation, they seem to be in agreement: in both cases, the option price is an decreasing function of the algebraic value of the autocorrelation.

Sections 3.3 and 4.2 below.

3.3 Implications for Option Prices

Expression (3.18) provides the necessary input to the Black-Scholes formula for pricing options on an asset with the trending O-U dynamics. In particular, if the unconditional variance of daily returns is $s^2(r_1)$, and if the first-order autocorrelation of τ -period returns is $\rho_\tau(1)$, then the price of a call option is given by:

$$C_{ou}(P(t), t; K, T, r, \sigma) = P(t)\Phi(d_1) - Ke^{-r(T-t)}\Phi(d_2) \quad (3.23)$$

where:

$$d_1 \equiv \frac{\log(P(t)/K) + (r + \frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t}} \quad (3.24)$$

$$d_2 \equiv \frac{\log(P(t)/K) + (r - \frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t}} \quad (3.25)$$

$$\sigma^2 = \frac{s^2(r_1)}{\tau} \cdot \frac{\log(1 + 2\rho_\tau(1))}{([1 + 2\rho_\tau(1)]^{1/\tau} - 1)}, \quad \rho_\tau(1) \in (-\frac{1}{2}, 0] \quad (3.26)$$

which is simply the Black-Scholes formula with an adjusted volatility input.¹⁴ In particular, the adjustment factor multiplying $s^2(r_1)/\tau$ in (3.26) is easily tabulated [see Table 3 and the discussion in Section 7.2], hence in practice it is a simple matter to adjust the Black-Scholes formula for negative autocorrelation of the form (3.15): multiply the usual variance estimator $s^2(r_1)/\tau$ by the appropriate factor from Table 3 and use this as σ^2 in the Black-Scholes formula.

Note that for all values of $\rho_\tau(1)$ in $(-1/2, 0]$, the factor multiplying $s^2(r_1)/\tau$ in (3.26) is greater than or equal to one, and increasing in the absolute value of the first-order autocorrelation coefficient. This implies that option values under the trending O-U specification are always greater than or equal to options under the standard Black-Scholes specification, and that option values are an increasing function of the absolute value of the first-order autocorrelation coefficient. These are purely features of the trending O-U process and do

¹⁴Since $s^2(r_1)$ is the unconditional variance of daily returns, observe that σ^2 is also measured in days, as is the time-to-maturity $T-t$ and the interest rate r .

not generalize to other specifications of the drift, as we shall see below.

To gauge the empirical relevance of this adjustment for autocorrelation, Tables 1a-c report a comparison of Black-Scholes prices under arithmetic Brownian motion and under the trending Ornstein Uhlenbeck process for various holding periods, strike prices, and autocorrelations for a hypothetical \$40 stock. Table 1a reports option prices for values of daily autocorrelations from -5% to -45%, and Tables 1b and 1c report prices for weekly and monthly autocorrelations of the same numerical values. For all three tables, the unconditional standard deviation of daily returns is held fixed at 2% per day. The Black-Scholes price is calculated according to (2.4), setting σ equal to the unconditional standard deviation. The trending O-U prices are calculated by solving (3.14) and (3.15) for σ given τ and the return autocorrelations $\rho_\tau(1)$ of -0.05, -0.10, -0.20, -0.30, -0.40, and -0.45, and using these values of σ in the Black-Scholes formula (2.4). In Table 1a, $\tau = 1$; in Tables 1b and 1c, $\tau = 7$ and $364/12$, respectively.

The first panel of Table 1a shows that even extreme autocorrelation in daily returns does not affect short-maturity in-the-money call options prices very much. For example, a daily autocorrelation of -45% has no impact on the \$30 7-day call; the price under the trending O-U process is identical to the standard Black-Scholes price of \$10.028. But even for such a short maturity, differences become more pronounced as the strike price increases; the at-the-money call is worth \$0.863 in the absence of autocorrelation, but increases to \$1.368 with an autocorrelation of -45%.

However, as the time to maturity increases, the remaining panels of Table 1a show that the impact of autocorrelation also increases. With a -10% daily autocorrelation, an at-the-money 1-year call is \$7.234, and rises to \$10.343 with a daily autocorrelation of -45%, compared to the standard Black-Scholes price of \$6.908. This pattern is not surprising, given the autocorrelation (3.15) of the trending O-U process, which declines with the length of the holding period so that longer-horizon returns are more highly negatively autocorrelated and, therefore, depart more severely from the standard Black-Scholes paradigm than shorter-horizon returns.

More formally, since the Black-Scholes formula applies to both arithmetic Brownian motion and the trending O-U process, the impact of a specification error in the drift can be related to the sensitivity of the Black-Scholes formula to changes in volatility σ . This sen-

sitivity is measured by the derivative of the call price with respect to σ , and is often called the option's "vega":

$$\frac{\partial C}{\partial \sigma} = P(t)\sqrt{T-t}\Phi'(d_1) \quad (3.27)$$

where d_1 is defined in (2.5). From (3.27), we see that for shorter-maturity options, changes in σ have very little impact on the call price, but longer-maturity options will be more sensitive.

This is also apparent in the patterns of Tables 1b and 1c, which are similar to those in Table 1a but much less striking since the same numerical values of $\rho_\tau(1)$ are now assumed to hold for weekly and monthly returns, respectively. As Table 1 shows, the impact of a -45% autocorrelation in monthly returns is considerably less than the same autocorrelation in daily returns.

In contrast to Table 1a where an at-the-money 1-year call increases from \$6.908 to \$10.343 as the autocorrelation decreases from 0 to -45% , in Table 1c the same option increases from \$6.908 to only \$7.018. We shall see in Table 3 of Section 7.2 that this is a symptom of all diffusion processes, since the increments of any diffusion process becomes less autocorrelated as the differencing interval declines. In particular, Table 3 will show that the impact of a -45% autocorrelation in monthly returns is considerably less than the same autocorrelation in daily returns. Indeed, from (3.15), a -45% autocorrelation in monthly returns implies an autocorrelation of -0.97% in daily returns. Therefore, the importance of autocorrelation for option prices hinges critically on the degree of autocorrelation for a given return horizon τ and, of course, on the data-generating process which determines how rapidly this autocorrelation decays with τ . For this reason, in the next section we introduce several new stochastic processes that are capable of matching more complex patterns of autocorrelation and predictability than the trending O-U process.

4 The Bivariate Trending O-U Process

An obvious deficiency of the trending Ornstein-Uhlenbeck process as a general model of asset prices is the fact that its returns are negatively autocorrelated at all lags, which is inconsistent with the empirical autocorrelations of many traded assets. For example, Lo and MacKinlay

(1988, 1990) show that equity portfolios tend to be positively autocorrelated at shorter horizons, while Fama and French (1988) and Poterba and Summers (1988) find negative autocorrelation at longer horizons. Moreover, since the trending O-U's drift depends only on $q(t)$, it leaves no role for other economic variables to play in determining the predictability of asset returns.

To address these shortcomings, we propose the "bivariate trending O-U" process in the following sections. Although it is a special case of a bivariate linear diffusion process, and is therefore extremely tractable, it exhibits a surprisingly wide variety of autocorrelation patterns [see, for example, Figure 1]. Moreover, as its name suggests, the bivariate trending O-U process allows the log-price process to depend upon a second process, which may be interpreted as a time-varying expected return factor that may or may not be observable.

4.1 Properties and Unconditional Moments

Let the detrended log-price process $q(t) \equiv p(t) - \mu t$ satisfy the following pair of stochastic differential equations:

$$dq(t) = -(\gamma q(t) - \lambda X(t))dt + \sigma dW_q \quad (4.1)$$

$$dX(t) = -\delta X(t)dt + \sigma_x dW_x \quad (4.2)$$

where $\gamma \geq 0$, $\delta \geq 0$, $q(0) = q_0$, $X(0) = X_0$, $t \in [0, \infty)$.

W_q and W_x are two standard Wiener processes such that $dW_q dW_x = \kappa dt$, and $X(t)$ is another stochastic process which may or may not be observable. For reasons that will become apparent below, we shall call this system the "bivariate trending O-U" process.

The bivariate system (4.1) - (4.2) contains several interesting special cases. For example, when $\lambda = 0$ it reduces to the univariate trending O-U process of Section 3.1, in which asset returns are always negatively autocorrelated. When $\gamma = 0$, the drift of the detrended log-price process is $\lambda X(t)$, which is stochastic and mean-reverting to its unconditional mean of zero. In the more general case when $\gamma \neq 0$, the detrended log-price process may be rewritten

as:

$$dq(t) = -\gamma \left(q(t) - \frac{\lambda}{\gamma} X(t) \right) dt + \sigma dW_q \quad (4.3)$$

which shows that $q(t)$ is mean-reverting to a stochastic mean $\frac{\lambda}{\gamma} X(t)$, with "speed of adjustment" γ .

Since (4.1) - (4.2) is a system of linear stochastic differential equations, (q, X) is Gaussian given its initial value (q_0, X_0) at $t = 0$, and has the following explicit solution:

$$q(t) = e^{-\gamma t} q_0 + \frac{\lambda}{\gamma - \delta} [e^{-\delta t} - e^{-\gamma t}] X_0 + \int_0^t e^{-\gamma(t-s)} \sigma dW_q(s) + \frac{\lambda}{\gamma - \delta} \int_0^t [e^{-\delta(t-s)} - e^{-\gamma(t-s)}] \sigma_x dW_x(s) \quad (4.4)$$

$$X(t) = e^{-\delta t} X_0 + \int_0^t e^{-\delta(t-s)} \sigma_x dW_x(s) \quad (4.5)$$

where $t > 0$ and $q_0 = p_0$. Conditional upon $\{q_0, X_0\}$, $q(t)$ and $X(t)$ are jointly normally distributed.¹⁵

From (4.4) - (4.5) we can readily derive the properties of the asset return series that (q, X) generates. To do this, observe that when $\gamma > 0$ and $\delta > 0$, both $q(t)$ and $X(t)$ are stationary and their first two unconditional moments are:

$$\begin{aligned} E[q(t)] &= E[X(t)] = 0 \\ \text{Var}[q(t)] &= \frac{\sigma^2}{2\gamma} + \frac{\lambda^2 \sigma_x^2}{2\gamma\delta(\gamma+\delta)} + \frac{\lambda\kappa\sigma\sigma_x}{\gamma(\gamma+\delta)} \\ \text{Var}[X(t)] &= \frac{\sigma_x^2}{2\delta} \\ \text{Cov}[q(t), X(t)] &= \frac{1}{\gamma+\delta} \left(\kappa\sigma\sigma_x + \frac{\lambda\sigma_x^2}{2\delta} \right) \end{aligned}$$

The unconditional moments of continuously-compounded τ -period returns then follow from

¹⁵Even if $\{q_0, X_0\}$ are stochastic, as long as they are drawn from their stationary joint distribution, $\{q(t), X(t)\}$ is still jointly normally distributed.

(4.4) - (4.5):¹⁶

$$E[r_\tau(t)] = \mu\tau \quad (4.6)$$

$$\text{Var}[r_\tau(t)] = 2\text{Var}[q(t)] \left[(1-e^{-\gamma\tau}) - \frac{\lambda}{\gamma-\delta}\beta_{qx}(e^{-\delta\tau} - e^{-\gamma\tau}) \right] \quad (4.7)$$

$$\text{Cov}[r_\tau(t+\tau), r_\tau(t)] = -\text{Var}[q(t)] \left[\frac{\lambda}{\gamma-\delta}\beta_{qx} \left[(1-e^{-\delta\tau})^2 - (1-e^{-\gamma\tau})^2 \right] + (1-e^{-\gamma\tau})^2 \right] \quad (4.8)$$

$$\rho_\tau(1) = -\frac{(1-e^{-\gamma\tau})^2 + \frac{\lambda}{\gamma-\delta}\beta_{qx} \left[(1-e^{-\delta\tau})^2 - (1-e^{-\gamma\tau})^2 \right]}{2 \left[(1-e^{-\gamma\tau}) - \frac{\lambda}{\gamma-\delta}\beta_{qx}(e^{-\delta\tau} - e^{-\gamma\tau}) \right]} \quad (4.9)$$

where $\beta_{qx} \equiv \text{Cov}[q(t), X(t)]/\text{Var}[q(t)]$ and $\rho_\tau(1)$ is the first-order autocorrelation function of τ -period returns.

As in the case of the univariate trending O-U process, the bivariate process is trend-stationary, the variance of its increments approaches a finite limit of $2\text{Var}[q(t)]$ and the first-order autocorrelation $\rho_\tau(1)$ of τ -period returns approaches $-\frac{1}{2}$ as τ increases without bound. Both of these restrictions are relaxed in the multivariate version of Section 5.

To see that the bivariate trending O-U process can capture more complex patterns of autocorrelation than its univariate counterpart, consider the behavior of its first-order autocorrelation function as a function of the holding period τ for the special case where $\lambda = \gamma$. As τ increases without bound, $\rho_\tau(1)$ approaches $-\frac{1}{2}$ as it must for the continuously-compounded τ -period return of any stationary process. As τ decreases to 0, $\rho_\tau(1)$ also approaches zero as it must for any diffusion process, since diffusions have locally independent increments by construction. For small τ , we have:

$$\rho_\tau(1) \approx -\frac{\tau}{2} \left(\gamma - \frac{\delta\beta_{qx}}{1-\beta_{qx}} \right) \quad (4.10)$$

which can be either positive or negative, depending on whether β_{qx} is greater than or less than $\frac{\gamma}{\gamma+\delta}$. Therefore, when $\beta_{qx} > \frac{\gamma}{\gamma+\delta}$, the bivariate trending O-U process will display an

¹⁶If $\gamma = 0$ or $\delta = 0$, the unconditional moments of $q(t)$ and $X(t)$ may not exist. However, the unconditional moments of returns are always well-defined, and may be obtained by taking the appropriate limits in the following results.

autocorrelation pattern that matches the empirical findings of both Lo and MacKinlay (1988) and Fama and French (1988) simultaneously: positive autocorrelation for short horizons, and negative autocorrelation for long horizons. Some other examples of first-order autocorrelation functions of the bivariate trending O-U process are given in Figure 1.

A closely-related quantity that may help to develop further intuition for the bivariate trending O-U process is the general autocorrelation function $\rho_\tau(k)$, defined as the correlation between two τ -period continuously-compounded returns that are $(k-1)\tau$ periods apart, i.e.,

$$\rho_\tau(k) \equiv \frac{\text{Cov}[r_\tau(t+k\tau), r_\tau(t)]}{\text{Var}[r_\tau(t)]}$$

Observe that the first-order autocorrelation function $\rho_\tau(1)$, defined in (3.8), is indeed a special case of this more general definition. In the case of the bivariate trending O-U process, the autocorrelation function is given by:

$$\rho_\tau(k) = e^{-(k-1)\gamma\tau} \rho_\tau(1) - \frac{\lambda}{\gamma-\delta} [e^{-(k-1)\delta\tau} - e^{-(k-1)\gamma\tau}] \theta(\tau) \quad (4.11)$$

where

$$\theta(\tau) \equiv \frac{\beta_{qx}(1-e^{-\delta\tau})^2}{2 \left[(1-e^{-\gamma\tau}) - \frac{\lambda}{\gamma-\delta} \beta_{qx}(e^{-\delta\tau} - e^{-\gamma\tau}) \right]}$$

4.2 Predictability vs. Autocorrelation

We have argued in Section 3 that the numerical value of the Black-Scholes input σ depends on our assumption about the data generating process when we have discretely-sampled data. In the particular case of the univariate trending O-U process of Section 3.1, the numerical value of σ increases with the absolute value of the return autocorrelation, given a fixed numerical value for the unconditional variance of returns. However, in the case of the bivariate trending O-U process, there is no longer such a simple relation between autocorrelation and σ .

For example, consider the special case of the bivariate trending O-U process in which $\gamma = 0$, hence $\lambda X(t)$ is the drift of the detrended log-price process, and the system reduces

to:

$$dq(t) = \lambda X(t)dt + \sigma dW_q \quad (4.12)$$

$$dX(t) = -\delta X(t)dt + \sigma_x dW_x. \quad (4.13)$$

For simplicity, also let $\kappa = 0$ so that dW_q and dW_x are statistically independent. In this special case, asset returns are positively autocorrelated at all leads and lags. We may calculate the unconditional variance and autocorrelation of returns by taking the limit of $\gamma \rightarrow 0$ in (4.7) and (4.9). Then, for any holding period τ we have:

$$s^2(r_\tau) = \sigma^2 \frac{1}{1 - \sigma_{qx}^*} \left[\tau - \frac{\sigma_{qx}^*}{\delta} (1 - e^{-\delta\tau}) \right] \quad (4.14)$$

$$\rho_\tau(1) = \frac{\frac{\sigma_{qx}^*}{\delta} (1 - e^{-\delta\tau})^2}{2\tau \left[1 - \frac{\sigma_{qx}^*}{\delta\tau} (1 - e^{-\delta\tau}) \right]}, \quad \text{where } \sigma_{qx}^* \equiv \frac{1}{1 + \left(\frac{\delta}{\lambda}\right)^2 \frac{\sigma_x^2}{\sigma_q^2}}.$$

Observe that $0 < \sigma_{qx}^* \leq 1$, and that σ_{qx}^* is an increasing function of λ . Since $\rho_\tau(1)$ is an increasing function of σ_{qx}^* , it is also an increasing function of λ . By increasing λ while holding fixed the unconditional variance of returns, we can see the effects of increasing autocorrelation on the Black-Scholes input σ . Re-arranging (4.14) yields:

$$\sigma^2 = \frac{1 - \sigma_{qx}^*}{1 - \frac{\sigma_{qx}^*}{\delta\tau} (1 - e^{-\delta\tau})} \cdot \frac{s^2(r_\tau)}{\tau} \quad (4.15)$$

which shows that an increase in the return autocorrelation (through an increase in λ) is accompanied by a decrease in σ and a corresponding decrease in the Black-Scholes call option price.¹⁷ Increasing return autocorrelation in this case has precisely the opposite effect on option prices than in the case of the univariate trending O-U process, in which an increase in the absolute value of the return autocorrelation [recall that in this case, the autocorrelation is always nonpositive] increases the numerical value of σ , increasing option prices.

While increasing autocorrelation can either increase or decrease option prices, depending on the particular specification of the drift, the special case (4.12) - (4.13) does illustrate a

¹⁷It is easy to show that the expression $(1 - \sigma_{qx}^*) / \left[1 - \frac{\sigma_{qx}^*}{\delta\tau} (1 - e^{-\delta\tau}) \right]$ decreases as σ_{qx}^* increases. It then follows that increasing λ will increase its value since σ_{qx}^* will increase.

general relation between option prices and *predictability*. To see this, we must first define predictability explicitly. Perhaps the most common definition is the R^2 coefficient, or the fraction of the unconditional variance of the dependent variable that is “explained” by the conditional mean or predictor. Higher R^2 s are generally taken to mean more predictability, and this interpretation is appropriate in our context with three additional restrictions:

(A1) The unconditional variance of returns $r_\tau(t)$ is fixed.

(A2) The drift is not a function of the log-price process $p(t)$.

(A3) dW_q is statistically independent of dW_x .

The first restriction has already been discussed above—the very nature of prediction takes as *given* the object to be predicted, and meaningful comparisons of alternate prediction equations cannot be made if the “target” is allowed to change in any way. In particular, if the unconditional variance of $r_\tau(t)$ is not fixed, a reduction in the prediction error variance need not imply better predictability because it may be accompanied by a more-than-proportionate reduction in the unconditional variance to be predicted.

Restrictions (A2) and (A3) eliminate feedback relations between the conditional mean and the prediction error or residual, so that the discrete-time representation of the continuous-time system is a genuine prediction equation, i.e., the conditional expectation of the residual, conditioned on the drift, is zero.

Under these restrictions, it may be shown that an increase in predictability—as measured by R^2 —always decreases σ and therefore decreases option prices.¹⁸ The intuition for this relation is clear: holding fixed the unconditional variance of returns, an increase in the variability of the conditional mean must imply a decrease in the variability of the residual. More formally, the unconditional variance of returns may always be written as the following sum:

$$\text{Var}[r_\tau(t)] = \text{Var}\left[E[r_\tau(t)|\Omega_t] \right] + E\left[\text{Var}[r_\tau(t)|\Omega_t] \right] \quad (4.16)$$

¹⁸It is important to note that, here and throughout this paper, we consider changes in predictability or R^2 only through changes in the *specification* or functional form of the drift. Another way to change predictability, which we do not consider, is to change the information set under a given specification of the continuous-time process.

where Ω_t is the conditioning information set. Holding the left-hand side of (4.16) fixed, an increase in the first term of the right-hand side, i.e., an increase in predictability, must be accompanied by an equal decrease in the second term of the right-hand side. Furthermore, under restrictions (A1) – (A3), the variability of the residual can be shown to be monotonically related to the continuous-time parameter σ , hence increasing predictability implies decreasing option prices.

In particular, under the bivariate trending O-U process, increasing λ has the effect of increasing the variability of the conditional mean. Holding the unconditional variance of the returns $s^2(r_t)$ fixed, an increase in λ will therefore increase the predictability of returns, implying that the value of σ^2 must decrease since conditions (A1) – (A3) are satisfied by (4.12) – (4.13).¹⁹ As λ increases without bound so that progressively more variation in returns is attributable to the time-varying drift, returns become progressively more predictable, σ approaches 0, and the option's value approaches its lower bound of $e^{-r(T-t)} \text{Max} \{ P(T) - K, 0 \}$.²⁰ Only if predictability is defined in this narrow sense, and only under conditions (A1) – (A3), is there an unambiguous relation between predictability and option prices.

Under more general conditions, however, a simple relation between predictability and option prices is not available, and the very notion of predictability need not be well-defined. For example, condition (A2) is violated by the univariate trending O-U process of Section 3.1, and in that case, while increasing predictability does decrease the variance of the prediction error of $r_t(t)$, it also *increases* σ .

¹⁹Another way to see this is to consider the special case where the conditioning information set at t is the whole sample path of $X(\cdot)$ from t to $t+\tau$, i.e., $\Omega_t = [X_s : t \leq s \leq t+\tau]$. In this case,

$$\text{Var} \left[E[r_t(t) | \Omega_t] \right] = \left(\frac{\lambda}{\delta} \right)^2 \sigma_x^2 \left[\tau - \frac{1}{\delta} (1 - e^{-\delta\tau}) \right], \quad E \left[\text{Var}[r_t(t) | \Omega] \right] = \sigma_y^2 \tau.$$

Thus, the residual variance is simply $\sigma_y^2 \tau$. Increasing λ increases the predictability in returns and decreases σ^2 holding $\text{Var}[r_t(t)]$ fixed. We thank the referee for suggesting this special case.

²⁰Note that this particular limit is economically unrealizable because even though the stock price is still stochastic when σ vanishes (due to the drift), it is once-differentiable and therefore admits arbitrage [see Harrison, Pitbladdo, and Schaefer (1984)].

4.3 A Numerical Example

To illustrate the importance of predictability in determining the Black-Scholes input σ , we use historical daily returns on the CRSP value-weighted market index from 1962 to 1990 to calibrate the bivariate trending O-U process and evaluate σ explicitly. Since all second-order moments of continuously-compounded returns depend on the six underlying parameters of the bivariate process, γ , δ , λ , σ , σ_x and κ , we may choose any six moments and solve for the six underlying parameters. Moreover, if $\gamma \neq 0$, we can set $\lambda = \gamma$ without loss of generality, which reduces the total number of free parameters to five. To further simplify the calibration exercise, we set $\kappa = 0$. Thus, we require only four second-order moments to determine γ , δ , σ and σ_x .

For the four second-order moments, we use the sample variance of the returns $\text{Var}\{r_\tau(t)\}$, the first order autocorrelation coefficient $\rho_\tau(1)$ and two higher-order autocorrelation coefficients. If the bivariate trending O-U process is the true DGP and we possessed the actual population values of the moments, then of course choice of which two higher-order autocorrelation coefficients to fit is arbitrary, since they will arrive at the same parameter values. However, since we are using actual data to perform the calibration, and are not estimating the parameters of the system, some care is required in selecting the moments to match. In particular, since the autocorrelation function of the bivariate trending O-U process can change sign only once [from positive to negative], we must choose our moments to be consistent with this restriction. With this in mind, we select the following four moments for our calibration:

$$\begin{aligned} s(r_\tau) &= 0.0085 \\ \rho_\tau(1) &= 0.1838 \\ \rho_\tau(5) &= 0.0323 \\ \rho_\tau(25) &= -0.0092 \end{aligned} \tag{4.17}$$

where $\rho_\tau(k)$ denotes the k -th order autocorrelation of τ -period returns. Calibrating the

parameters to these moments yield the following values:²¹

$$\begin{aligned}\gamma &= 0.3748 \\ \delta &= 0.0106 \\ \sigma_x &= 0.0128 \\ \sigma &= 0.0074\end{aligned}\tag{4.18}$$

Observe that the value of the Black-Scholes input σ under the bivariate trending O-U specification, 0.0074, is approximately 13% smaller than the standard deviation of continuously-compounded returns 0.0085, which is the value of σ under an arithmetic Brownian motion specification.

The theoretical call option prices for a hypothetical \$40 stock in Table 2 show that such a difference can have potentially large effects, particularly for longer-maturity options just as in Tables 1a-c. However, in this case the naive Black-Scholes prices are *over-estimates* of the correct call price, since the σ that accounts for predictability is lower than the σ obtained under an iid assumption.

5 The Multivariate Trending O-U Process

Despite the flexibility of the bivariate trending O-U process, as a model of asset prices it has at least three unattractive features that are both related to the behavior of its increments as the differencing interval increases without bound: the variance of its increments approaches a finite limit, its first-order autocorrelation function approaches a limit of $-\frac{1}{2}$, and can change sign only once. Moreover, the bivariate process does not allow for additional economic or "state" variables that might affect the drift.

In this section we present a multivariate extension of the bivariate trending O-U process that addresses all of these concerns. By allowing the drift to depend linearly on additional state variables, resulting in the "multivariate trending O-U process", richer patterns of autocorrelation can be captured without sacrificing tractability. If the state variables are stationary, then log-prices are trend-stationary as in the bivariate case. If the state variables are random walks, then log-prices will contain stochastic trends, in which case the variance of

²¹Note that the solution for γ , δ and σ_x is not unique, however, the solution for σ is.

its increments can increase without bound and the first-order autocorrelation can approach 0 as the differencing interval increases.

In a straightforward generalization of the bivariate case, we let the detrended log-price process $q(t)$ fluctuate around a stochastic mean, now governed by a multivariate linear process. Specifically, let:

$$dq(t) = [-\gamma q(t) + \Lambda X(t)] dt + \sigma dW_q \quad (5.1)$$

$$dX(t) = -\Delta X(t) dt + B_x dW_x \quad (5.2)$$

with $q(0) = q_0$, $X(0) = X_0$, $t \in [0, \infty)$.

where $X(t)$ is an m -dimensional random process, $W_x(t)$ a k -dimensional standard Wiener process, γ and σ are scalar parameters, and Λ , Δ , B_x are $(1 \times m)$, $(m \times m)$, $(m \times k)$ matrix parameters, respectively. Without loss of generality we assume that Δ is diagonal, i.e., $\Delta = \text{diag}\{\delta_i\}$. The linear system $[q(t) \ X(t)']$ defined by (5.1) - (5.2) has the following explicit solution:

$$q(t) = e^{-\gamma t} q_0 + \Lambda(\gamma I - \Delta)^{-1} [e^{-\Delta t} - e^{-\gamma I t}] X_0 + \int_0^t e^{-\gamma(t-s)} \sigma dW_q(s) - \Lambda(\gamma I - \Delta)^{-1} \int_0^t [e^{-\gamma I(t-s)} - e^{-\Delta(t-s)}] B_x dW_x(s) \quad (5.3)$$

$$X(t) = e^{-\Delta t} X_0 + \int_0^t e^{-\Delta(t-s)} B_x dW_x(s) \quad (5.4)$$

where I is the $(m \times m)$ identity matrix.²² Since Δ is diagonal, $e^{-\Delta(t-s)} = \text{diag}\{e^{-\delta_i(t-s)}\}$. Given (5.3) - (5.4), we can readily derive the unconditional moments of $q(t)$ and $X(t)$ (if they exist), as well as those of returns over any finite holding period τ .²³

If the diagonal matrix Δ contains strictly positive diagonal entries δ_i , then the log-price process is trend-stationary as in the case of the bivariate trending O-U process. The unconditional moments of the detrended log-price process and returns follow analogously

²²We have implicitly assumed that $\gamma \neq \delta_i$, $i = 1, \dots, m$ so that the inverse of $\gamma I - \Delta$ exists. If not, we can derive the corresponding solution by taking the appropriate limit.

²³As in the bivariate case, when Δ is not of full rank, i.e., when $\delta_i = 0$ for some i , or when $\gamma = 0$, the unconditional moments of $q(t)$ no longer exist. However, the unconditional moments of returns do exist and they can be calculated by taking the appropriate limits; see the discussion in Section 5.

from (5.3) – (5.4) and some of these are reported in the Appendix.

Alternatively, if a subset of the state variables follow random walks or is “difference-stationary”, then the log-price process will also be difference-stationary and the variance of its increments will increase without bound as the differencing interval approaches infinity. For example, consider the following trivariate special case of (5.2). Let $\mathbf{X}(t) \equiv [X(t) \ Z(t)]'$, $\Delta = \text{diag}(\delta, 0)$, $\mathbf{B}_x = \text{diag}(\sigma_x, \sigma_x)$, and $d\mathbf{W}_x = [dW_x \ dW_x]'$. In this case, $X(t)$ follows an O-U process while $Z(t)$ follows a random walk. Assume that $\gamma > 0$. Without loss of generality, we can let $\Lambda = [\gamma \ \gamma]$. Then the explicit solution for the detrended price process $q(t)$ is:

$$q(t) = \tilde{q}(t) + Z(t) \quad (5.5)$$

$$\begin{aligned} \tilde{q}(t) = & e^{-\gamma t} \tilde{q}_0 + \frac{\gamma}{\gamma - \delta} (e^{-\delta t} - e^{-\gamma t}) X_0 + \int_0^t e^{-\gamma(t-s)} \sigma dW_q(s) - \\ & \int_0^t e^{-\gamma(t-s)} \sigma_x dW_x(s) + \frac{\gamma}{\gamma - \delta} \int_0^t [e^{-\delta(t-s)} - e^{-\gamma(t-s)}] \sigma_x dW_x(s) \end{aligned} \quad (5.6)$$

$$X(t) = e^{-\delta t} X_0 + \int_0^t e^{-\delta(t-s)} \sigma_x dW_x(s) \quad (5.7)$$

$$Z(t) = Z_0 + \int_0^t \sigma_z dW_z(s) \quad (5.8)$$

Observe that $q(t)$ can be decomposed into two components: a stationary component $\tilde{q}(t)$ and a random walk component $Z(t)$ where the stationary component $\tilde{q}(t)$ behaves like the detrended log-price in the stationary bivariate O-U case. However, in contrast to the trend-stationary case, the existence of a random walk component in the detrended log-price implies that the risk of holding the asset increases with the holding period.

Of course, when $q(t)$ is non-stationary, the unconditional moments of $q(t)$ are no longer well-defined. However, the unconditional moments of the *increments* of $q(t)$, which are simply the de-measured continuously compounded returns, are well-defined and may be obtained from the results for the stationary case by taking the limit that $\delta_x \rightarrow 0$.

6 Maximum Likelihood Estimation

The fact that the univariate and the bivariate trending O-U processes imply such different relations between autocorrelation and option values illustrates the complexity and importance of correctly identifying the data-generating process *before* implementing an option pricing formula. In the previous sections, we have shown that holding fixed the unconditional moments of the true data-generating process, a change in the specification of the drift can change the population value of the Black-Scholes input σ . As a result, a change in the specification of the drift can also change the empirical estimate of σ .

Perhaps the most direct approach to addressing these issues is to propose a reasonably flexible specification of the drift that can capture a wide variety of autocorrelation patterns, derive the exact discrete-time representation of the log-price process, estimate *all* the parameters of this discrete-time process simultaneously, and then solve for the parameters of the continuous-time process—which includes σ —as a function of the parameter estimates of the discretely-sampled data. Since all three of our specifications for the drift are linear, their discrete-time representations are readily available and are also linear processes, to which maximum likelihood estimation may be applied, as described in Lo (1986, 1988).

To this end, denote by t_k the sampling dates, where $k = 1, 2, \dots, n$, and let $t_k - t_{k-1} = \tau$ be a constant, hence $t_k = k\tau$.²⁴ Let $q_k \equiv q(t_k) = p(t_k) - \mu t_k$ and assume that q_k is observed. Of course, in practice the trend rate μ must be estimated, but as long as a consistent estimator of μ is available, replacing μ with $\hat{\mu}$ will have no effect upon the asymptotic properties of the parameter estimates.

6.1 The Univariate Trending O-U Process

From the explicit solution (3.4) of the univariate trending O-U process, it is easy to obtain a recursive representation of q_k which shows that its deviations from trend follow an AR(1):

$$q_k = e^{-\gamma\tau} q_{k-1} + \epsilon_k \quad , \quad \epsilon_k \equiv \sigma \int_{t_{k-1}}^{t_k} e^{\gamma(t_k-s)} dW(s) . \quad (6.1)$$

²⁴This last assumption is made purely for notational convenience—irregularly-sampled data may be just as easily accommodated but is notationally more cumbersome.

For this simple process, the maximum likelihood estimator of the discrete-time parameters is asymptotically equivalent to the ordinary least squares estimator applied to detrended prices. The continuous-time parameters μ , σ , and γ may then be obtained from the discrete-time parameter estimates.

6.2 The Bivariate Trending O-U Process

Let $X_k \equiv X(t_k)$. Then from (4.1) and (4.2), we have:

$$q_k = \alpha_q q_{k-1} + \phi X_{k-1} + \epsilon_{q,k} \quad (6.2)$$

$$X_k = \alpha_x X_{k-1} + \epsilon_{x,k} \quad (6.3)$$

where $\alpha_q \equiv e^{-\gamma\tau}$, $\alpha_x \equiv e^{-\delta\tau}$, $\phi \equiv \frac{\lambda}{\gamma-\delta}(\alpha_x - \alpha_q)$, and

$$\epsilon_{q,k} \equiv \int_{t_{k-1}}^{t_k} e^{-\gamma(t_k-s)} \sigma dW_q(s) + \frac{\lambda}{\gamma-\delta} \int_{t_{k-1}}^{t_k} [e^{-\delta(t_k-s)} - e^{-\gamma(t_k-s)}] \sigma_x dW_x(s)$$

$$\epsilon_{x,k} \equiv \int_{t_{k-1}}^{t_k} e^{-\delta(t_k-s)} \sigma_x dW_x(s).$$

Observe that $[\epsilon_{q,k} \ \epsilon_{x,k}]'$ is a bivariate normal vector that is temporally independently and identically distributed, with mean 0 and covariance matrix S_τ given in the Appendix. Re-writing (6.2) – (6.3) in vector form yields:

$$\begin{pmatrix} q_k \\ X_k \end{pmatrix} = \begin{pmatrix} \alpha_q & \phi \\ 0 & \alpha_x \end{pmatrix} \begin{pmatrix} q_{k-1} \\ X_{k-1} \end{pmatrix} + \begin{pmatrix} \epsilon_{q,k} \\ \epsilon_{x,k} \end{pmatrix}. \quad (6.4)$$

This is simply a bivariate AR(1) process, where the second component X_k may or may not be observed. The parameters of this discrete-time process may be estimated by maximum likelihood by casting (6.4) in state-space form and applying the Kalman filter [see, for example, Harvey (1989a) or Lütkepohl (1991)]. There are seven parameters to be estimated: μ , α_q , α_x , ϕ , and the elements of the symmetric (2x2) matrix S_τ . From the definition of these discrete-time parameters, we can uniquely determine the seven parameters of the underlying continuous-time process, μ , γ , δ , λ , σ , σ_x and κ [see (A.5), (A.6), (A.7), and (A.8) in the Appendix], hence the principle of invariance yields maximum likelihood estimators for these

as well.

6.3 The Multivariate Trending O-U Process

The discrete-time representation of (5.1) – (5.2) is a straightforward generalization of the bivariate case:

$$q_k = \alpha_q q_{k-1} + \Phi X_{k-1} + \epsilon_{q,k} \quad (6.5)$$

$$X_k = A_X X_{k-1} + \epsilon_{X,k} \quad (6.6)$$

where $q_k = q(t_k)$, $X(t_k) = X_k$, $\alpha_q \equiv e^{-\gamma\tau}$, $A_X \equiv e^{-\Delta\tau}$, $\Phi = \Lambda(\gamma I - \Delta)^{-1} (A_X - \alpha_q I)$, and

$$\begin{aligned} \epsilon_{q,k} &\equiv \int_{t_{k-1}}^{t_k} e^{-\gamma(t_k-s)} \sigma dW_q(s) - \\ &\quad \Lambda(\gamma I - \Delta)^{-1} \int_{t_{k-1}}^{t_k} [e^{-\gamma I(t_k-s)} - e^{-\Delta(t_k-s)}] B_X dW_X(s) \\ \epsilon_{X,k} &\equiv \int_{t_{k-1}}^{t_k} e^{-\Delta(t_k-s)} B_X dW_X(s). \end{aligned}$$

Observe that $e_k \equiv [\epsilon_{q,k} \ \epsilon'_{X,k}]'$ is an $(m+1)$ -dimensional normal random variable which is temporally independently and identically distributed. In vector form, we have:

$$\begin{pmatrix} q_k \\ X_k \end{pmatrix} = \begin{pmatrix} \alpha_q & \Phi \\ 0 & A_X \end{pmatrix} \begin{pmatrix} q_{k-1} \\ X_{k-1} \end{pmatrix} + e_k. \quad (6.7)$$

which is a VAR(1), and given observations $\{p_k\}$, or $\{p_k\}$ and some components of $\{X_k\}$, we can obtain maximum likelihood estimates of its parameters by applying the Kalman filter to the state-space representation as before.²⁵

In our trivariate example (5.5) of Section 5 which is difference-stationary, the discrete-time representation of (5.5) – (5.8) is:

$$(q_k - Z_k) = \alpha_p (q_{k-1} - Z_{k-1}) + \phi X_{k-1} + \epsilon_{\bar{q},k} \quad (6.8)$$

$$X_k = \alpha_x X_{k-1} + \epsilon_{x,k} \quad (6.9)$$

²⁵However, in the general multivariate case, identification is not guaranteed and is often difficult to verify. See Lütkepohl (1991, Chapter 13.4.2) for further discussion.

$$Z_k = Z_{k-1} + \epsilon_{z,k} \quad (6.10)$$

where $Z_k \equiv Z(t_k)$, and $\epsilon_{q,k}$, $\epsilon_{x,k}$ and $\epsilon_{z,k}$ are iid Gaussian shocks derived from the stochastic integrals in (5.6), (5.7), and (5.8). Since q_k is non-stationary here, prices cannot be used directly to estimate the parameters. Instead, de-meaned continuously compounded returns may be used since they are stationary under this current specification. Define $\tilde{r}_k \equiv q_k - q_{k-1}$, $\nu_k \equiv X_k - X_{k-1}$, and $e_k \equiv [\epsilon_{q,k} \ \epsilon_{x,k} \ \epsilon_{z,k}]'$. We then have:

$$\begin{pmatrix} \tilde{r}_k \\ \nu_k \end{pmatrix} = \begin{pmatrix} \alpha_p & \phi \\ 0 & \alpha_x \end{pmatrix} \begin{pmatrix} \tilde{r}_{k-1} \\ \nu_{k-1} \end{pmatrix} + \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix} e_k - \begin{pmatrix} 1 & 0 & -\alpha_p \\ 0 & 1 & 0 \end{pmatrix} e_{k-1} \quad (6.11)$$

which is simply a multivariate ARMA(1,1) process. Once again, given observations $\{r_k\}$, or $\{r_k, \nu_k\}$, maximum likelihood estimation of the discrete-time parameters may be readily performed as in the trend-stationary case via its state-space representation.

7 Extensions and Other Issues

There are several other aspects of the impact of predictability on option prices that deserve further discussion, such as extensions to option pricing models other than the Black-Scholes model, implications of the distinction between discrete and continuous time, the relation of our findings to those surrounding "estimation risk", and the interpretation of implied volatilities in the presence of predictability. We shall consider each of these issues in turn in the following sections.

7.1 Extensions to Other Option Pricing Models

Although we have confined our attention so far to the case where the diffusion coefficient σ is constant—the Black-Scholes case—predictability can affect other option and derivative pricing formulas in a similar fashion. Since analytical pricing formulas for options and other derivative assets are almost always obtained from no-arbitrage conditions, the drift plays no role in determining the formula but plays a critical role in determining both the population values and empirical estimates of the parameters that enter the formula as arguments. For example, although the drift does not enter into Merton's (1976a) jump-diffusion option

pricing formula, its specification will affect the values of σ [the volatility of the diffusion component], δ [the volatility of the logarithm of the jump magnitude], k [the expectation of the logarithm of the jump magnitude], and λ [the mean rate of occurrence of the Poisson jump]. Since all of our drift specifications in Sections 3, 4, and 5 are linear, they may be readily incorporated into more complex stochastic processes.

7.2 Discrete vs. Continuous Time

Clearly, the importance of the drift in implementing option pricing formulas comes from the fact that the data are sampled at discrete time intervals, while the theoretical models are formulated in continuous time. Now it is well known that the diffusion coefficient is a "sample-path property", so that any single realization of a *continuous* sample path over a finite interval is sufficient to reveal the true value of σ . However, continuous sample paths are practically unrealizable, therefore we are always confronted with some sampling error in our attempts to estimate σ . Assuming for the moment that the diffusion coefficient σ is indeed constant, the sampling error of any estimator of σ can be traced to two distinct sources: misspecification of the drift, and the discreteness of the sampling interval.

Of course, these two sources of sampling error are closely related. For example, the effects of misspecifying the drift diminishes as the sampling frequency increases, and in the limit there is no sampling error in estimating σ hence misspecification of the drift is irrelevant. In particular, consider the relation between the continuous-time parameter σ^2 and the finite holding-period return variance $s^2(r_\tau)$ for the univariate trending O-U case of Section 3.1. For any fixed value of σ^2 , (3.17) shows that as the return horizon τ decreases to 0 the ratio of σ^2 to $s^2(r_\tau)/\tau$ approaches 1, hence σ^2 may be recovered exactly in the limit of continuous sampling. This is a general property of diffusions (2.2) with a constant diffusion coefficient—the unconditional variance $s^2(r_\tau)$ approaches $[dq(t)]^2 = \sigma^2 dt$ as the holding period τ approaches zero.²⁶ Alternatively, $\sigma^2 dt$ may be viewed as the conditional variance of dq , conditional on the drift. But since all the infinitesimal variation in dq is attributable to the diffusion term $\sigma^2 dW$ [recall that the drift is of order dt and the diffusion term is of order \sqrt{dt}], the conditional and unconditional variance of the stochastic differential

²⁶In fact, even if the diffusion coefficient is time-varying, it may be estimated with arbitrary precision by sampling more frequently within a fixed time span. See Huang and Lo (1994) for further details.

dq are effectively the same [see Sims (1984) for further details].

This limiting result may lead some to advocate using the most finely-sampled data available to compute $s^2(r_\tau)/\tau$, so as to minimize the effects of the drift of the data-generating process. Of course, whether or not the most finely-sampled data available is fine enough to render $s^2(r_\tau)/\tau$ an adequate approximation to σ^2 is an empirical issue that depends critically on what the true data-generating process is, and on the types of market microstructure effects that may come into play.

It is also conceivable that the sampling error in $\hat{\sigma}$ induced by a misspecification of the drift is not nearly so great as the sampling error induced by discrete sampling. While specifying a "better" drift may yield a closer approximation to the continuous-time process, it may not improve the performance of $\hat{\sigma}$ for a given set of discretely-sampled data.

The potential importance of both sources of sampling error are, of course, empirical issues that must be resolved on an individual basis with a particular application and dataset at hand. For example, consider the univariate trending O-U process of Section 3.1, and recall from (3.19) - (3.21) that $\Delta(1, \tau, \rho_\tau(1))$ provides a convenient measure of the impact of serial correlation on the Black-Scholes input σ^2 as a function of the first-order autocorrelation coefficient $\rho_\tau(1)$ for τ -period returns.

For example, let $s^2(r_1)$ be defined for *daily* returns, and suppose that the first-order autocorrelation of *daily* returns is -30% . Table 3 shows that $\Delta(1, 1, -0.30) = 1.527$, hence the value of $s^2(r_1)$ must be increased by 52.7% to yield the correct value for the Black-Scholes input σ^2 . If, however, a -30% first-order autocorrelation is observed for *5-day* returns, this should yield a smaller autocorrelation for daily returns [recall that in the limit, the autocorrelation vanishes], which is confirmed by Table 3's entry of 1.094 for $\Delta(1, 5, -.30)$, i.e., σ^2 is only 9.4% larger than $s^2(r_1)$ in this case. Even in the extreme case of a -45% autocorrelation, if this autocorrelation is for 25-day returns, σ^2 is only 4.7% larger than $s^2(r_1)$, whereas the same autocorrelation for daily returns implies that σ^2 is 156% larger than $s^2(r_1)$. Contrary to conventional wisdom, the autocorrelation coefficient is *not* unitless, and has an important time element to it.

A second method of gauging the relative importance of a misspecification of the drift in the sampling error of $\hat{\sigma}$ in the case of the univariate trending O-U process is to perform a simple Monte Carlo simulation experiment. For a given sample size, say 250 observations,

consider simulating a sample path of daily returns under the univariate trending O-U specification, estimating σ with and without an adjustment for the drift, and repeating this 5,000 times to obtain the finite-sample distribution of the two estimators.

Table 4 reports the outcome of such an experiment for sample sizes ranging from 250 to 1,250 daily observations [roughly one to five years of daily data], and for first-order autocorrelations ρ_1 ranging from -10 to -45% , and holding the daily unconditional variance fixed at 2% . The first two columns report the parameter values of the simulation [note that σ changes with ρ_1 since we have fixed the daily unconditional variance at 2%], and columns 3 and 4 report the percentage bias of the naive estimator \hat{s} [which is simply the sample standard deviation of daily returns] and the estimator $\hat{\sigma}$ that adjusts for the mean-reverting drift [based on (3.18)], where the expected value of each of the two estimators is computed over the 5,000 replications of each simulation.

The first row of the first panel of Table 4 shows that \hat{s} is considerably more biased than $\hat{\sigma}$: -5.1% for \hat{s} versus 0.5% for $\hat{\sigma}$. For more extreme values of ρ_1 , the biases of both estimators worsen, but even in the worst case when ρ_1 is -45% , \hat{s} is more biased than $\hat{\sigma}$: -37.3% versus 12.1% . The relative performance of \hat{s} and $\hat{\sigma}$ is similar for larger sample sizes.

Table 4 also reports the theoretical value of Δ , which relates σ to s according to (3.19), and its expectation over each of the 5,000 replications. For most sample sizes and values of ρ_1 , Δ and $\hat{\Delta}$ are fairly close, which explains the superiority of $\hat{\sigma}$ to \hat{s} .

A somewhat more subtle issue surrounding the distinction between discrete and continuous time is the fact that while we have used the Black-Scholes formula to gauge the effects of asset return predictability on option prices, it may be argued that the Black-Scholes formula holds only if continuous trading is possible and costless. Indeed, to implement the replicating strategy literally requires observing the sample-path of prices continuously, which eliminates the need for *estimating* σ altogether. In this case, the relation between predictability and option prices still exists but is irrelevant since the true σ can always be recovered exactly.

However, the continuous-trading assumption underlying the pricing formulas does not invalidate our main conclusion: whenever option pricing formulas are implemented with discretely-sampled data, the drift matters.

Of course, ideally we should also incorporate the effect of discreteness into the pricing formulas to provide a complete and empirically relevant theory of option pricing. One

approach is to simply impose discrete trading, e.g., Black and Scholes (1972) and Boyle and Emanuel (1980), and another approach is to take into account directly the economic causes of discrete trading, such as transactions costs, in constructing replicating strategies, e.g., Leland (1985). These approaches will yield either approximate pricing formulas or bounds for option prices, and in both cases, the results will certainly depend on the numerical value of the diffusion coefficient σ which in turn will depend on the specification of the drift, *ceteris paribus*. Therefore, despite the fact that in the continuous-time limit σ becomes known, any empirical implementation must incorporate the effects of predictability on option prices.

7.3 Estimation Risk

It is important to note that the effects of predictability on option prices is closely related to, but not synonymous with the problem of "estimation risk" [see, for example, Barry et al. (1991)]. As we have just discussed in Section 7.2, the fact that σ^2 must be estimated from discretely-sampled data provides the primary motivation for our analysis. But the link between σ^2 and asset return predictability exists even when σ^2 is known without error. Of course, if σ^2 is known, then the degree of predictability in asset returns is irrelevant for purposes of pricing options even if the link is present. However, when σ^2 is unknown, the precise form of asset return predictability will affect both the estimate and the estimation risk of σ^2 .

7.4 Implied Volatilities

A consequence of the Black-Scholes model is that the parameter σ may be recovered from option prices directly by inverting (2.4). Therefore, why go through the trouble of relating asset return predictability to σ ? The response to this simple but perplexing question is quite straightforward.

The relevance of the implied volatility relies on the proper specification of the option pricing formula. If prices were truly Black-Scholes prices, then implied volatilities would be irrelevant since the Black-Scholes model requires that σ is known. But if the market price is not truly a Black-Scholes price, then an implied volatility obtained from the Black-Scholes formula is difficult to interpret and use.

Alternatively, for the Black-Scholes formula to yield the correct prices, *some* investors

must know the true value of σ [or equivalently, must observe a continuous sample path of the stock price over a finite time interval] and must be able to trade continuously. If these conditions held, other investors could obtain the true value of σ by simply inverting the Black-Scholes formula. If, however, all investors do not know the true value of σ but only observe [and trade at] prices at discrete time intervals, then they will have to infer the value of σ from the data instead of the Black-Scholes formula which no longer holds exactly. This is the situation we consider here. In this case, the market prices of options do not provide any information about the underlying stock price process other than what can be inferred from the data, such as the true value of σ .²⁷ In other words, market option prices imply the true value of σ only if some investors actually know σ . Market prices cannot reveal something that nobody knows.²⁸

Even in the case where some investors may have better information about σ , other investors can rely completely on the implied volatility from market prices only if the prices are informationally efficient [see footnote 28]. If there is noise in market option prices due to market imperfections, e.g., frictions, deviations from perfect competition, etc., less informed investors will have to rely on their own information as well in making inferences about σ , such as discretely-sampled stock prices [see Grossman and Stiglitz (1980)].

8 Conclusion

The fact that asset return predictability has nontrivial implications for option prices provides a link between two seemingly disparate strands of the asset pricing literature: linear multi-factor models of time-varying expected asset returns, and arbitrage-based models of

²⁷It should be emphasized that our analysis of the impact of predictability on option prices does not hinge on the particular pricing model we use. In particular, we focus on the Black-Scholes formula solely because of its simplicity. The fact that the Black-Scholes formula may not be correct in its exact form under the maintained assumptions of our analysis—that only discretely-sampled data is observed and investors may not be able to trade continuously—affects only the numerical values of our examples, and not the substance of our conclusions.

²⁸Of course, one can construct situations where prices can reveal information that no single investor has [see Grossman (1976)], but this is achievable only in the extreme case where prices can efficiently aggregate all the information in the market, i.e., full informational efficiency. For both theoretical and practical reasons, this case is of little interest in modeling financial markets [see Grossman and Stiglitz (1980)].

derivative asset prices.²⁹ Heuristically, when predictability is well-defined, i.e., when the asset return's conditional mean does not depend upon past prices or returns, and when the conditional expectation of the prediction error is zero, then increases in predictability generally decrease option prices when the unconditional variance of asset returns is fixed. In such cases, an increase in predictability is equivalent to a reduction in the asset's residual uncertainty or prediction-error variance, and since option prices are monotonically increasing in the volatility of this residual uncertainty in the Black-Scholes case where the diffusion coefficient σ is constant, option prices decline as predictability increases.

This has an interesting implication for the evolution of option premia through time: as we are better able to model the time-varying expected return of an asset, option premia on that asset should fall, *ceteris paribus*. Alternatively, the fact that option premia are positive may imply an upper bound on the predictability of the underlying asset's returns, which may partly address Roll's (1988) lament that the R^2 's in financial applications are disappointingly low. We hope to explore these implications in future research.

For alternatives to the Black-Scholes case, such as those with stochastic volatility or jump components, predictability also affects option prices nontrivially, but in considerably more complex ways. To capture such effects, each of our drift specifications can be paired with a particular specification for the diffusion coefficient. While closed-form adjustments for predictability may not always exist in these more general cases, maximum likelihood estimation is almost always feasible for our linear drift specifications.

Despite the fact that the drift of a diffusion process plays virtually no role in deriving theoretical pricing formulas for derivative assets, its importance cannot be overemphasized in the implementation of those formulas. The practical value of arbitrage-based models of derivative prices rests heavily on the existence of an empirically plausible and stable model of the true data-generating process for the underlying asset's price. Although changing specifications for the drift does not influence the derivative pricing formula, it does influence both the theoretical value and empirical estimate of the parameter(s) on which the formula depends.

Moreover, although our approach begins by first specifying the true distribution of the

²⁹At least three other papers have hinted at such a link: Dybvig and Ingersoll (1982), Grundy (1991), and Lo (1989).

data-generating process [which includes the modeling of predictability], and then deriving its implications for the risk-neutral distribution, it is also possible to do the reverse. Taking the properties of the risk-neutral distribution implicit in derivative asset prices as a starting point, it is possible to infer the properties of the true data-generating process [see Grundy (1991), for example]. The fact that the risk-neutral distribution and the true distribution of the data-generating process are linked makes both directions of inference possible and empirically relevant.³⁰

³⁰Which approach to take, ours or Grundy's, may depend on which market is considered more "efficient" in conveying information: the market for an underlying asset, or the market for derivatives of that asset. For example, some have argued that options markets are often more liquid and "informed" than their corresponding spot markets. In such cases, it may well be easier to use option prices to infer predictability in the underlying asset's returns, rather than to study the asset's returns directly. However, since the many recent studies of predictability focus exclusively on the behavior of the underlying asset's DGP, we take the true distribution of the DGP as our starting point.

Appendix

A The Bivariate Trending O-U Process

To derive (6.2) - (6.3), observe that from (4.1) - (4.2), we have:

$$q(t_k) = q(t_{k-1})e^{-\gamma\tau} + \frac{\lambda}{\gamma-\delta} (e^{-\delta\tau} - e^{-\gamma\tau}) X(t_k) + \int_{t_{k-1}}^{t_k} e^{-\gamma(t_k-s)} \sigma dW_q(s) + \frac{\lambda}{\gamma-\delta} \int_{t_{k-1}}^{t_k} [e^{-\delta(t_k-s)} - e^{-\gamma(t_k-s)}] \sigma_x dW_x(s) \quad (\text{A.1})$$

$$X(t_k) = e^{-\delta\tau} X(t_{k-1}) + \int_{t_{k-1}}^{t_k} e^{-\delta(t_k-s)} \sigma_x dW_x(s). \quad (\text{A.2})$$

Define $\alpha_q \equiv e^{-\gamma\tau}$, $\alpha_x \equiv e^{-\delta\tau}$, $\phi = \frac{\lambda}{\gamma-\delta} (\alpha_x - \alpha_q)$, and

$$\begin{aligned} \epsilon_{q,k} &= \int_{t_{k-1}}^{t_k} e^{-\gamma(t_k-s)} \sigma dW_q(s) + \frac{\lambda}{\gamma-\delta} \int_{t_{k-1}}^{t_k} [e^{-\delta(t_k-s)} - e^{-\gamma(t_k-s)}] \sigma_x dW_x(s) \\ \epsilon_{x,k} &= \int_{t_{k-1}}^{t_k} e^{-\delta(t_k-s)} \sigma_x dW_x(s). \end{aligned}$$

We then have:

$$q_k = \alpha_q q_{k-1} + \phi X_{k-1} + \epsilon_{q,k} \quad (\text{A.3})$$

$$X_k = \alpha_x X_{k-1} + \epsilon_{x,k}. \quad (\text{A.4})$$

Clearly, $\epsilon_{q,k}$ and $\epsilon_{x,k}$ are independently and identically distributed over time and jointly normally distributed. Furthermore,

$$\begin{aligned} s_{q,\tau}^2 \equiv \text{Var}[\epsilon_{q,k}] &= \frac{\sigma^2}{2\gamma} (1 - \alpha_q^2) + \frac{2\lambda\kappa\sigma_x\sigma}{\gamma-\delta} \left[\frac{1 - \alpha_q\alpha_x}{\gamma+\delta} - \frac{1 - \alpha_q^2}{2\gamma} \right] + \\ &\quad \frac{\lambda^2\sigma_x^2}{(\gamma-\delta)^2} \left[\frac{1 - \alpha_q^2}{2\gamma} + \frac{1 - \alpha_x^2}{2\delta} - \frac{2(1 - \alpha_q\alpha_x)}{\gamma+\delta} \right] \end{aligned} \quad (\text{A.5})$$

$$s_{x,\tau}^2 \equiv \text{Var}[\epsilon_{x,k}] = \frac{\sigma_x^2}{2\delta} (1 - \alpha_x^2) \quad (\text{A.6})$$

$$s_{qx,\tau} \equiv \text{Cov}[\epsilon_{q,k}, \epsilon_{x,k}] = \frac{\kappa\sigma_x\sigma}{\gamma+\delta} (1 - \alpha_q\alpha_x) + \frac{\lambda\sigma_x^2}{\gamma-\delta} \left[\frac{1 - \alpha_x^2}{2\delta} - \frac{1 - \alpha_q\alpha_x}{\gamma+\delta} \right]. \quad (\text{A.7})$$

There is a one-to-one mapping between the parameters of the discretely-sampled system, α_q , α_x , ϕ , $s_{q,\tau}^2$, $s_{x,\tau}^2$, $s_{qx,\tau}$, and the parameters of the underlying continuous-time process, γ , δ ,

$\lambda, \sigma, \sigma_x, \kappa$. Specifically, normalize the time units so that $\tau = 1$ and observe that:

$$\gamma = -\log \alpha_q, \quad \delta = -\log \alpha_x, \quad \lambda = \frac{\phi(\log \alpha_q - \log \alpha_x)}{\alpha_q - \alpha_x}. \quad (\text{A.8})$$

Substituting (A.8) into (A.5) – (A.7) then yields three equations which are linear in σ^2, σ_x^2 and $\kappa\sigma\sigma_x$, hence the remaining three continuous-time parameters may be easily recovered from these equations.

Since (p_k, X_k) follows a bivariate AR(1) process, a closed-form expression for the likelihood function of p_k may be obtained which can be used in the maximum likelihood estimation [see, for example, Jazwinski (1970)].

B The Multivariate Trending O-U Process

The multivariate trending O-U process $(q(t), X'(t))$ is defined by the following Itô integrals:

$$q(t) = e^{-\gamma t} q_0 + \Lambda(\gamma I - \Delta)^{-1} [e^{-\Delta t} - e^{-\gamma I t}] X_0 + \int_0^t e^{-\gamma(t-s)} \sigma dW_q(s) - \Lambda(\gamma I - \Delta)^{-1} \int_0^t [e^{-\gamma I(t-s)} - e^{-\Delta(t-s)}] B_x dW_x(s) \quad (\text{B.1})$$

$$X(t) = e^{-\Delta t} X_0 + \int_0^t e^{-\Delta(t-s)} B_x dW_x(s) \quad (\text{B.2})$$

where I is the $(m \times m)$ identity matrix. When γ and the real parts of all of the eigenvalues of Δ are strictly positive, $(q(t), X'(t))$ is stationary. The unconditional moments of $q(t)$ and $X(t)$ may be readily obtained from (B.1) and (B.2).

Since $W_x(t)$ is a k -dimensional standard Wiener process, $E[dW_x dW_x'] = I dt$ where I is the identity matrix of order k .³¹ Let $\sigma dW_q B_x dW_x = K dt$ where K is a $(k \times 1)$ vector. For notational convenience, define $\Sigma = \{\sigma_{ij}\} \equiv B_x B_x'$, $\Omega_\tau = \{\omega_{ij}(\tau)\}$ where $\omega_{ij}(\tau) \equiv \sigma_{ij} [1 - e^{-(\delta_i + \delta_j)\tau}] / (\delta_i + \delta_j)$ and $\Xi_\tau = \{\xi_{ij}(\tau)\}$ where $\xi_{ij}(\tau) \equiv \sigma_{ij} [1 - e^{-(\gamma + \delta_j)\tau}] / (\gamma + \delta_j)$. Then we have:

$$\begin{aligned} E[q(t)] &= E[X(t)] = 0 \\ \text{Var}[q(t)] &= \frac{\sigma^2}{2\gamma} + \Lambda(\gamma I - \Delta)^{-1} \left(\frac{1}{2\gamma} \Sigma - \Xi_\infty - \Xi'_\infty + \Omega_\infty \right) (\gamma I - \Delta)^{-1} \Lambda' + \end{aligned} \quad (\text{B.3})$$

³¹There is no loss of generality by assuming that $W_x(t)$ has independent components since components of $X(t)$ can have arbitrary covariance structure through B_x .

$$\Lambda(\gamma\mathbf{I}-\Delta)^{-1} \left[2(\gamma\mathbf{I}+\Delta)^{-1} - \frac{1}{\gamma}\mathbf{I} \right] \mathbf{K} \quad (\text{B.4})$$

$$\text{Var}[\mathbf{X}(t)] = \Omega_\infty \quad (\text{B.5})$$

$$\text{Cov}[q(t), \mathbf{X}(t)] = (\Xi_\infty - \Omega_\infty)(\gamma\mathbf{I}-\Delta)^{-1}\Lambda' + (\gamma\mathbf{I}+\Delta)^{-1}\mathbf{K} \quad (\text{B.6})$$

$$\begin{aligned} \text{Cov}[q(t), q(t+\tau)] &= e^{-\gamma\tau} \text{Var}[q(t)] + \\ &\quad \Lambda(\gamma\mathbf{I}-\Delta)^{-1} (e^{-\Delta\tau} - e^{-\gamma\mathbf{I}\tau}) \text{Cov}[q(t), \mathbf{X}(t)]. \end{aligned} \quad (\text{B.7})$$

From these expressions, the moments of $r_\tau(t)$ follow directly:

$$\mathbb{E}[r_\tau(t)] = \mu\tau \quad (\text{B.8})$$

$$\text{Var}[r_\tau(t)] = 2\text{Var}[q(t)] \left[(1 - e^{-\gamma\tau}) - \Lambda(\gamma\mathbf{I}-\Delta)^{-1} (e^{-\Delta\tau} - e^{-\gamma\mathbf{I}\tau}) \mathbf{b} \right] \quad (\text{B.9})$$

$$\begin{aligned} \text{Cov}[r_\tau(t+\tau), r_\tau(t)] &= -\text{Var}[q(t)] \left\{ (1 - e^{-\gamma\tau})^2 + \right. \\ &\quad \left. \Lambda(\gamma\mathbf{I}-\Delta)^{-1} \left[(\mathbf{I} - e^{-\Delta\tau})^2 - (\mathbf{I} - e^{-\gamma\mathbf{I}\tau})^2 \right] \mathbf{b} \right\}. \end{aligned} \quad (\text{B.10})$$

where $\mathbf{b} \equiv \text{Cov}[q(t), \mathbf{X}(t)]/\text{Var}[q(t)]$. The return autocorrelation function may then be obtained from these moments.

It is straightforward to derive the discrete-time representation of the system $(q(t), \mathbf{X}'(t))$:

$$q_k = \alpha_q q_{k-1} + \Phi \mathbf{X}_{k-1} + \epsilon_{q,k} \quad (\text{B.11})$$

$$\mathbf{X}_k = \mathbf{A}_x \mathbf{X}_{k-1} + \epsilon_{x,k} \quad (\text{B.12})$$

where $q_k = q(t_k)$, $\mathbf{X}_k = \mathbf{X}(t_k)$, $\tau \equiv t_k - t_{k-1}$, $\alpha_q \equiv e^{-\gamma\tau}$, $\mathbf{A}_x \equiv e^{-\Delta\tau}$, $\Phi \equiv \Lambda(\gamma\mathbf{I}-\Delta)^{-1}(\mathbf{A}_x - \alpha_q\mathbf{I})$ and

$$\epsilon_{q,k} \equiv \int_{t_{k-1}}^{t_k} e^{-\gamma(t_k-s)} \sigma dW_q(s) - \Lambda(\gamma\mathbf{I}-\Delta)^{-1} \int_{t_{k-1}}^{t_k} [e^{-\gamma\mathbf{I}(t_k-s)} - e^{-\Delta(t_k-s)}] \mathbf{B}_x dW_x(s)$$

$$\epsilon_{x,k} \equiv \int_{t_{k-1}}^{t_k} e^{-\Delta(t_k-s)} \mathbf{B}_x dW_x(s).$$

It is easy to show that

$$\begin{aligned} s_{q,\tau}^2 &\equiv \text{Var}[\epsilon_{q,k}] = \frac{1-\alpha_q^2}{2\gamma} \sigma^2 + \Lambda(\gamma\mathbf{I}-\Delta)^{-1} \left(\frac{1-\alpha_q^2}{2\gamma} \Sigma - \Xi_\tau - \Xi'_\tau + \Omega_\tau \right) (\gamma\mathbf{I}-\Delta)^{-1} \Lambda' + \\ &\quad \Lambda(\gamma\mathbf{I}-\Delta)^{-1} \left\{ 2(\gamma\mathbf{I}+\Delta)^{-1} [1 - e^{-(\gamma\mathbf{I}+\Delta)\tau}] - \frac{1-\alpha_q^2}{\gamma} \mathbf{I} \right\} \mathbf{K} \end{aligned}$$

$$s_{x,\tau}^2 \equiv \text{Var}[\epsilon_{x,k}] = \Omega_\tau$$

$$s_{qx,\tau} \equiv \text{Cov}[\epsilon_{q,k}, \epsilon_{x,k}] = (\Omega_\tau - \Xi_\tau)(\gamma\mathbf{I}-\Delta)^{-1} \Lambda' + (\gamma\mathbf{I}+\Delta)^{-1} [1 - e^{-(\gamma\mathbf{I}+\Delta)\tau}] \mathbf{K}.$$

Similar to the bivariate case, the mapping between the parameters of the discrete-time representation, $\alpha_q, \mathbf{A}_x, \Phi, s_{q,\tau}^2, s_{x,\tau}^2, s_{qx,\tau}$, and the parameters of the underlying continuous-time process, $\gamma, \Delta, \Lambda, \sigma^2, \Sigma$ and \mathbf{K} is one-to-one.³² Let $\tau = 1$, we have

$$\gamma = -\log \alpha_q, \quad \delta_i = -\log \alpha_i, \quad \Lambda = \Phi (\mathbf{A}_x - \alpha_q \mathbf{I})^{-1} (\gamma \mathbf{I} - \Delta). \quad (\text{B.13})$$

where $\alpha_j = \{\mathbf{A}_x\}_{jj}$ (Note that \mathbf{A}_x is diagonal). From $s_{x,\tau}^2$, we can solve for Ω_τ, Σ and Ξ_τ :

$$\Omega_\tau = s_{x,\tau}^2, \quad \sigma_{ij} = \frac{(\delta_i + \delta_j) \omega_{ij}(\tau)}{1 - \alpha_i \alpha_j}, \quad \xi_{ij} = \frac{\sigma_{ij} (1 - \alpha_i \alpha_j)}{\gamma + \delta_j}. \quad (\text{B.14})$$

We can then solve for \mathbf{K} given $s_{qx,\tau}$:

$$\mathbf{K} = \left[1 - e^{-(\gamma \mathbf{I} + \Delta)} \right]^{-1} (\gamma \mathbf{I} + \Delta) \left[s_{qx,\tau} - (\Omega_\tau - \Xi_\tau) (\gamma \mathbf{I} - \Delta)^{-1} \Lambda' \right]. \quad (\text{B.15})$$

From the definition of $s_{q,\tau}^2$, we can further solve for σ^2 .

³²Note that \mathbf{B}_x is simply the Cholesky decomposition of Σ .

References

- Abel, A., 1992, "Exact Solutions for Expected Rates of Return Under Markov Regime Switching: Implications for the Equity Premium Puzzle", NBER Working Paper No. 4110.
- Barry, C., French, D. and R. Rao, 1991, "Estimation Risk and Adaptive Behavior in the Pricing of Options", *Financial Review* 26, 15-30.
- Bekaert, G. and R. Hodrick, 1992, "Characterizing Predictable Components in Excess Returns on Equity and Foreign Exchange Markets", *Journal of Finance* 47, 467-509.
- Bessembinder, H. and K. Chan, 1992, "Time-Varying Risk Premia and Forecastable Returns in Futures Markets", to appear in *Journal of Financial Economics*.
- Black, F. and M. Scholes, 1972, "The Valuation of Option Contracts and a Test of Market Efficiency", *Journal of Finance* 27, 399-418.
- Black, F. and M. Scholes, 1973, "The Pricing of Options and Corporate Liabilities", *Journal of Political Economy* 81, 637-659.
- Boyle, P. and D. Emanuel, 1980, "Discretely Adjusted Option Hedges", *Journal of Financial Economics* 8, 259-82.
- Breen, W., Glosten, L. and R. Jagannathan, 1989, "Economic Significance of Predictable Variations in Stock Index Returns", *Journal of Finance* 44, 1177-1189.
- Campbell, J. and J. Ammer, 1993, "What Moves the Stock and Bond Markets? A Variance Decomposition for Long-Term Asset Returns", *Journal of Finance* 48, 3-37.
- Campbell, J. and Y. Hamao, 1992, "Predictable Stock Returns in the United States and Japan: A Study of Long-Term Capital Market Integration", *Journal of Finance* 47, 43-70.
- Campbell, J. and P. Perron, 1991, "Pitfalls and Opportunities: What Macroeconomists Should Know About Unit Roots", *NBER Macroeconomics Annual*, 141-201.
- Cecchetti, S., Lam, P. and N. Mark, 1991, "The Equity Premium and the Risk Free Rate: Matching the Moments", NBER Working Paper No. 3752.
- Chen, N., 1991, "Financial Investment Opportunities and the Macroeconomy", *Journal of Finance* 46, 529-554.
- Chen, N., Roll, R. and S. Ross, 1986, "Economic Forces and the Stock Market", *Journal of Business* 59, 383-403.
- Chopra, N., Lakonishok, J. and J. Ritter, 1992, "Measuring Abnormal Performance: Do Stocks Overreact?" *Journal of Financial Economics* 31, 235-268.
- Cox, J. and S. Ross, 1976, "The Valuation of Options for Alternative Stochastic Processes", *Journal of Financial Economics* 3, 145-166.
- DeBondt, W. and R. Thaler, 1985, "Does the Stock Market Overreact?" *Journal of Finance* 40, 793-805.

- Doob, J., 1953, *Stochastic Processes*. New York: John Wiley and Sons.
- Dybvig, P. and J. Ingersoll, 1982, "Mean-Variance Theory in Complete Markets", *Journal of Business* 55, 233-251.
- Engle, R., Lilien, D. and R. Robbins, 1987, "Estimating Time Varying Risk Premia in the Term Structure: the ARCH-M Model," *Econometrica* 55, 391-407.
- Fama, E. and K. French, 1988a, "Permanent and Temporary Components of Stock Prices", *Journal of Political Economy* 96, 246-273.
- Fama, E. and K. French, 1988b, "Dividend Yields and Expected Stock Returns", *Journal of Financial Economics* 22, 3-26.
- Fama, E. and K. French, 1990, "Business Conditions and Expected Returns on Stocks and Bonds", *Journal of Financial Economics* 25, 23-49.
- Ferson, W., 1989, "Changes in Expected Security Returns, Risk, and the Level of Interest Rates", *Journal of Finance* 44, 1191-1217.
- Ferson, W., 1990, "Are the Latent Variables in Time-Varying Expected Returns Compensation for Consumption Risk?" *Journal of Finance* 45, 397-430.
- Ferson, W., Foerster, S. and D. Keim, 1993, "General Tests of Latent Variable Models and Mean-Variance Spanning", *Journal of Finance* 48, 131-156.
- Ferson, W. and C. Harvey, 1991a, "Sources of Predictability in Portfolio Returns", *Financial Analysts Journal* May-June, 49-56.
- Ferson, W. and C. Harvey, 1991b, "The Variation of Economic Risk Premiums", *Journal of Political Economy* 99, 385-415.
- Ferson, W., Kandel, S. and R. Stambaugh, 1987, "Tests of Asset Pricing with Time-Varying Expected Risk Premiums and Market Betas", *Journal of Finance* 42, 201-220.
- Figlewski, S., 1989, "What Does an Option Pricing Model Tell Us About Option Prices?" *Financial Analysts Journal* September-October, 12-15.
- Gibbons, M. and W. Ferson, 1985, "Testing Asset Pricing Models with Changing Expectations and an Unobservable Market Portfolio", *Journal of Financial Economics* 14, 217-236.
- Grundy, B., 1991, "Option Prices and the Underlying Asset's Return Distribution", *Journal of Finance* 46, 1045-1069.
- Harrison, M., Pitbladdo, R. and S. Schaefer, 1984, Continuous Price Processes in Frictionless Markets Have Infinite Variation, *Journal of Business* 57, 353-365.
- Harvey, A., 1989a, *Forecasting, Structure Time Series Models and the Kalman Filter*. Cambridge, UK: Cambridge University Press.
- Harvey, C., 1989b, "Time-Varying Conditional Covariances in Tests of Asset Pricing Models", *Journal of Financial Economics* 24, 289-317.
- Heaton, J. and D. Lucas, 1992a, "The Effects of Incomplete Insurance Markets and Trading Costs in a Consumption-Based Asset Pricing Model", *Journal of Economic Dynamics and Control* 16, 601-620.

- Heaton, J. and D. Lucas, 1992b, "Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing", Working Paper No. 3491-92-EFA, Sloan School of Management, MIT.
- Huang, C. and A. Lo, 1994, "Estimating the Drift of a General Diffusion Process", unpublished working paper.
- Jagannathan, R., 1984, "Call Options and the Risk of Underlying Securities", *Journal of Financial Economics* 13, 425-434.
- Jazwinski, A., 1970, *Stochastic Processes and Filtering Theory*. New York: Academic Press.
- Jegadeesh, N., 1990, "Evidence of Predictable Behavior of Security Returns", *Journal of Finance* 45, 881-898.
- Kandel, S. and R. Stambaugh, 1988, "Modeling Expected Stock Returns for Long and Short Horizons", Rodney L. White Center Working Paper No. 42-88, Wharton School, University of Pennsylvania.
- Kandel S. and R. Stambaugh, 1990, "Expectations and Volatility of Consumption and Asset Returns", *Review of Financial Studies* 3, 207-232.
- Keim, D. and R. Stambaugh, 1986, "Predicting Returns in the Stock and Bond Markets", *Journal of Financial Economics* 17, 357-390.
- King, B., 1966, "Market and Industry Factors in Stock Price Behavior", *Journal of Business* 39, 139-190.
- Lehmann, B., 1990, "Fads, Martingales, and Market Efficiency," *Quarterly Journal of Economics* 105, 1-28.
- Leland, H., 1985, "Option Pricing and Replication with Transactions Costs", *Journal of Finance* 40, 1283-1301.
- Lo, A., 1986, "Statistical Tests of Contingent Claims Asset-Pricing Models: A New Methodology", *Journal of Financial Economics* 17, 143-173.
- Lo, A., 1988, "Maximum Likelihood Estimation of Generalized Itô Processes with Discretely Sampled Data", *Econometric Theory* 4, 231-247.
- Lo, A. and C. MacKinlay, 1988, "Stock Market Prices Do Not Follow Random Walks: Evidence From a Simple Specification Test", *Review of Financial Studies* 1, 41-66.
- Lo, A. and C. MacKinlay, 1990, "When Are Contrarian Profits Due to Stock Market Overreaction?" *Review of Financial Studies* 3, 175-206.
- Lo, A. and C. MacKinlay, 1992, "Maximizing Predictability in the Stock and Bond Markets", Working Paper No. 3450-92-EFA, Sloan School of Management, MIT.
- Lütkepohl, H., 1991, *Introduction to Multiple Time Series Analysis*. Berlin, Germany: Springer-Verlag.
- Mehra, R. and E. Prescott, 1985, "The Equity Premium: A Puzzle", *Journal of Monetary Economy* 15, 145-162.
- Merton, R., 1973, "Rational Theory of Option Pricing", *Bell Journal of Economics and Management Science* 4, 141-183.

- Merton, R., 1976a, "Option Pricing When Underlying Stock Returns are Discontinuous", *Journal of Financial Economics* 3, 125-144.
- Merton, R., 1976b, "The Impact on Option Pricing of Specification Error in the Underlying Stock Price Distribution", *Journal of Finance* 31, 333-350.
- Poterba, J. and L. Summers, 1988, "Mean Reversion in Stock Returns: Evidence and Implications", *Journal of Financial Economics* 22, 27-60.
- Roll, R., 1988, " R^2 ", *Journal of Finance* 43, 541-566.
- Sims, C., 1984, "Martingale-Like Behavior of Prices and Interest Rates", Discussion Paper No. 205, Center of Economic Research, Department of Economics, University of Minnesota.
- Sternberg, J. and J. Ying, 1992, "Option Pricing for Assets with Serially Correlated Returns", Working Paper, University of Arizona.
- Weil, P., 1989, "The Equity Premium Puzzle and the Risk-free Rate Puzzle", *Journal of Monetary Economy* 24, 401-421.

$\rho_{\tau}(1)$




Figure 1 - $\rho_{\tau}(1)$

$$\delta = 0.6, \sigma_p^2 = 0.5, \sigma_{\mu}^2 = 1.0, \kappa = 0$$

Table 1a

Comparison of theoretical call option prices on a hypothetical \$40 stock under an arithmetic Brownian motion versus a trending Ornstein-Uhlenbeck process for log-prices, assuming a standard deviation of 2 percent for daily continuously-compounded returns, and a daily continuously-compounded riskfree rate of $\log(1.05)/364$.

Strike Price	Black-Scholes Price	Trending O-U Price, with Daily $\rho_r(1) =$					
		-.05	-.10	-.20	-.30	-.40	-.45
Time-to-Maturity $T - t = 7$ Days							
30	10.028	10.028	10.028	10.028	10.028	10.028	10.028
35	5.036	5.037	5.038	5.042	5.051	5.074	5.108
40	0.863	0.885	0.910	0.973	1.062	1.216	1.368
45	0.011	0.013	0.016	0.024	0.041	0.082	0.137
50	0.000	0.000	0.000	0.000	0.000	0.001	0.005
Time-to-Maturity $T - t = 91$ Days							
30	10.526	10.547	10.572	10.640	10.754	10.989	11.262
35	6.331	6.387	6.451	6.614	6.855	7.289	7.735
40	3.270	3.350	3.439	3.661	3.978	4.526	5.068
45	1.459	1.532	1.615	1.822	2.124	2.657	3.195
50	0.574	0.623	0.680	0.829	1.058	1.491	1.954
Time-to-Maturity $T - t = 182$ Days							
30	11.285	11.336	11.394	11.548	11.786	12.238	12.725
35	7.558	7.646	7.746	7.998	8.365	9.014	9.668
40	4.740	4.851	4.976	5.286	5.728	6.491	7.244
45	2.810	2.922	3.048	3.361	3.812	4.595	5.375
50	1.592	1.687	1.797	2.073	2.482	3.214	3.963
Time-to-Maturity $T - t = 273$ Days							
30	12.040	12.113	12.198	12.415	12.745	13.352	13.989
35	8.587	8.698	8.824	9.139	9.596	10.396	11.199
40	5.905	6.039	6.191	6.565	7.099	8.019	8.925
45	3.943	4.082	4.239	4.627	5.185	6.147	7.098
50	2.573	2.702	2.849	3.217	3.753	4.695	5.642
Time-to-Maturity $T - t = 364$ Days							
30	12.753	12.845	12.950	13.218	13.620	14.349	15.102
35	9.493	9.622	9.769	10.133	10.661	11.582	12.501
40	6.908	7.061	7.234	7.660	8.269	9.315	10.343
45	4.941	5.102	5.283	5.732	6.374	7.478	8.566
50	3.489	3.645	3.821	4.261	4.896	6.003	7.106

Table 1b

Comparison of theoretical call option prices on a hypothetical \$40 stock under an arithmetic Brownian motion versus a trending Ornstein-Uhlenbeck process for log-prices, assuming a standard deviation of 2 percent for daily continuously-compounded returns, and a daily continuously-compounded riskfree rate of $\log(1.05)/364$.

Strike Price	Black-Scholes Price	Trending O-U Price, with Weekly $\rho_r(1) =$					
		-.05	-.10	-.20	-.30	-.40	-.45
Time-to-Maturity $T - t = 7$ Days							
30	10.028	10.028	10.028	10.028	10.028	10.028	10.028
35	5.036	5.036	5.037	5.037	5.037	5.038	5.040
40	0.863	0.866	0.870	0.878	0.891	0.912	0.933
45	0.011	0.011	0.011	0.012	0.013	0.016	0.019
50	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Time-to-Maturity $T - t = 91$ Days							
30	10.526	10.529	10.532	10.540	10.552	10.573	10.596
35	6.331	6.338	6.347	6.369	6.400	6.455	6.510
40	3.270	3.281	3.294	3.325	3.369	3.444	3.520
45	1.459	1.470	1.481	1.510	1.550	1.620	1.690
50	0.574	0.581	0.588	0.608	0.635	0.683	0.733
Time-to-Maturity $T - t = 182$ Days							
30	11.285	11.292	11.300	11.320	11.348	11.398	11.449
35	7.558	7.570	7.584	7.619	7.668	7.752	7.838
40	4.740	4.756	4.774	4.817	4.878	4.983	5.089
45	2.810	2.826	2.844	2.887	2.949	3.055	3.162
50	1.592	1.605	1.620	1.658	1.711	1.803	1.897
Time-to-Maturity $T - t = 273$ Days							
30	12.040	12.050	12.062	12.090	12.131	12.203	12.276
35	8.587	8.603	8.621	8.664	8.726	8.832	8.939
40	5.905	5.924	5.945	5.998	6.072	6.199	6.327
45	3.943	3.962	3.984	4.039	4.116	4.248	4.381
50	2.573	2.591	2.612	2.662	2.734	2.858	2.983
Time-to-Maturity $T - t = 364$ Days							
30	12.753	12.766	12.781	12.816	12.867	12.956	13.047
35	9.493	9.512	9.532	9.582	9.654	9.777	9.902
40	6.908	6.930	6.954	7.014	7.099	7.244	7.390
45	4.941	4.964	4.989	5.052	5.141	5.294	5.448
50	3.489	3.511	3.536	3.597	3.683	3.832	3.982

Table 1c

Comparison of theoretical call option prices on a hypothetical \$40 stock under an arithmetic Brownian motion versus a trending Ornstein-Uhlenbeck process for log-prices, assuming a standard deviation of 2 percent for daily continuously-compounded returns, and a daily continuously-compounded riskfree rate of $\log(1.05)/364$.

Strike Price	Black-Scholes Price	Trending O-U Price, with Monthly $\rho_r(1) =$					
		-.05	-.10	-.20	-.30	-.40	-.45
Time-to-Maturity $T - t = 7$ Days							
30	10.028	10.028	10.028	10.028	10.028	10.028	10.028
35	5.036	5.036	5.036	5.036	5.037	5.037	5.037
40	0.863	0.864	0.864	0.866	0.869	0.874	0.879
45	0.011	0.011	0.011	0.011	0.011	0.012	0.012
50	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Time-to-Maturity $T - t = 91$ Days							
30	10.526	10.527	10.528	10.529	10.532	10.537	10.541
35	6.331	6.332	6.334	6.339	6.346	6.359	6.371
40	3.270	3.273	3.276	3.283	3.293	3.310	3.327
45	1.459	1.462	1.464	1.471	1.480	1.496	1.512
50	0.574	0.575	0.577	0.581	0.588	0.598	0.609
Time-to-Maturity $T - t = 182$ Days							
30	11.285	11.287	11.289	11.293	11.300	11.310	11.321
35	7.558	7.561	7.564	7.572	7.583	7.602	7.621
40	4.740	4.744	4.748	4.758	4.772	4.796	4.820
45	2.810	2.814	2.818	2.828	2.842	2.866	2.890
50	1.592	1.595	1.598	1.607	1.619	1.640	1.660
Time-to-Maturity $T - t = 273$ Days							
30	12.040	12.042	12.045	12.051	12.061	12.076	12.092
35	8.587	8.591	8.595	8.605	8.619	8.643	8.667
40	5.905	5.909	5.914	5.926	5.943	5.972	6.002
45	3.943	3.947	3.952	3.965	3.982	4.012	4.043
50	2.573	2.577	2.582	2.594	2.610	2.638	2.666
Time-to-Maturity $T - t = 364$ Days							
30	12.753	12.756	12.759	12.768	12.779	12.799	12.819
35	9.493	9.497	9.502	9.514	9.530	9.558	9.586
40	6.908	6.913	6.919	6.933	6.952	6.985	7.018
45	4.941	4.946	4.952	4.966	4.987	5.022	5.057
50	3.489	3.494	3.500	3.514	3.534	3.567	3.601

Table 2

Comparison of theoretical call option prices on a hypothetical \$40 stock under an arithmetic Brownian motion versus a bivariate trending Ornstein-Uhlenbeck process for log-prices, both calibrated to match the daily CRSP value-weighted returns index from 1962 to 1990. The time-to-maturity is given by $T - t$, entries under the 'B-S' heading are call prices calculated under the Black-Scholes assumption of arithmetic Brownian motion [for which $\sigma = 0.0085$], and entries under the 'O-U' heading are call prices calculated under the bivariate trending Ornstein-Uhlenbeck [for which $\sigma = 0.0074$]. A daily continuously-compounded risk-free rate of $\log(1.05)/364$ is assumed.

Strike Price	$T - t = 7$		$T - t = 91$		$T - t = 182$		$T - t = 273$		$T - t = 364$	
	B-S	O-U	B-S	O-U	B-S	O-U	B-S	O-U	B-S	O-U
30	10.028	10.028	10.363	10.363	10.725	10.722	11.092	11.080	11.463	11.438
35	5.033	5.033	5.467	5.442	6.005	5.931	6.537	6.422	7.049	6.901
40	0.378	0.331	1.542	1.378	2.328	2.100	2.990	2.716	3.586	3.276
45	0.000	0.000	0.155	0.088	0.562	0.393	1.020	0.771	1.485	1.174
50	0.000	0.000	0.005	0.001	0.085	0.038	0.263	0.148	0.508	0.320

Table 3

Ratio of σ^2 to $s^2(r_1)$ for various values of the first-order autocorrelation $\rho_\tau(1)$ and holding period τ , where τ is measured in units of the holding period used to construct $s^2(r_1)$.

τ	$\rho_\tau(1)$						
	0.00	-0.05	-0.10	-0.20	-0.30	-0.40	-0.45
1	1.000	1.054	1.116	1.277	1.527	2.012	2.558
2	1.000	1.027	1.057	1.133	1.247	1.456	1.684
3	1.000	1.018	1.038	1.088	1.160	1.292	1.432
4	1.000	1.013	1.028	1.065	1.119	1.215	1.315
5	1.000	1.011	1.022	1.052	1.094	1.170	1.248
6	1.000	1.009	1.019	1.043	1.078	1.140	1.204
7	1.000	1.008	1.016	1.037	1.067	1.119	1.173
8	1.000	1.007	1.014	1.032	1.058	1.104	1.151
9	1.000	1.006	1.012	1.029	1.052	1.092	1.133
10	1.000	1.005	1.011	1.026	1.047	1.083	1.120
11	1.000	1.005	1.010	1.023	1.042	1.075	1.108
12	1.000	1.004	1.009	1.021	1.039	1.069	1.099
13	1.000	1.004	1.009	1.020	1.036	1.063	1.091
14	1.000	1.004	1.008	1.018	1.033	1.059	1.084
15	1.000	1.004	1.007	1.017	1.031	1.055	1.079
16	1.000	1.003	1.007	1.016	1.029	1.051	1.074
17	1.000	1.003	1.007	1.015	1.027	1.048	1.069
18	1.000	1.003	1.006	1.014	1.026	1.045	1.065
19	1.000	1.003	1.006	1.014	1.024	1.043	1.062
20	1.000	1.003	1.006	1.013	1.023	1.041	1.059
21	1.000	1.003	1.005	1.012	1.022	1.039	1.056
22	1.000	1.002	1.005	1.012	1.021	1.037	1.053
23	1.000	1.002	1.005	1.011	1.020	1.035	1.051
24	1.000	1.002	1.005	1.011	1.019	1.034	1.049
25	1.000	1.002	1.004	1.010	1.018	1.033	1.047

Table 4

Monte Carlo simulation of naive estimator \hat{s} and autocorrelation-adjusted estimator $\hat{\sigma}$ of Black-Scholes volatility input σ for various sample sizes and first-order autocorrelation coefficients ρ_1 . Each row corresponds to an independent simulation of 5,000 sample paths, where each path is generated according to a univariate trending Ornstein-Uhlenbeck process calibrated to *daily* returns with a standard deviation of 2% per day. Note that σ changes with ρ_1 because the unconditional standard deviation is fixed at 2% per day for all the simulations.

ρ_1	$\sigma \times 100$	%Bias(\hat{s}) (SE)	%Bias($\hat{\sigma}$) (SE)	$\sqrt{\Delta}$	$\sqrt{E[\Delta]}$ (SE)
Sample Size = 250					
-0.10	2.1126	-5.1488 (0.0614)	0.5048 (0.0852)	1.0563	1.0377 (0.0597)
-0.20	2.2601	-11.4339 (0.0598)	0.4501 (0.1000)	1.1301	1.1349 (0.0406)
-0.30	2.4716	-19.0601 (0.0569)	0.7178 (0.1263)	1.2358	1.2458 (0.0528)
-0.40	2.8368	-29.4781 (0.0515)	3.6719 (0.2851)	1.4184	1.4378 (0.0906)
-0.45	3.1990	-37.3341 (0.0469)	12.1053 (0.4813)	1.5995	1.5053 (0.1279)
Sample Size = 500					
-0.10	2.1126	-5.3581 (0.0439)	0.0625 (0.0610)	1.0563	1.0531 (0.0415)
-0.20	2.2601	-11.3826 (0.0414)	0.3298 (0.0679)	1.1301	1.1326 (0.0336)
-0.30	2.4716	-19.1139 (0.0395)	0.2136 (0.0844)	1.2358	1.2396 (0.0427)
-0.40	2.8368	-29.5296 (0.0364)	1.1622 (0.1465)	1.4184	1.4348 (0.0689)
-0.45	3.1990	-37.4316 (0.0333)	6.4146 (0.3412)	1.5995	1.5886 (0.1148)

Table 4 (Continued)

ρ_1	$\sigma \times 100$	%Bias($\hat{\theta}$) (SE)	%Bias($\hat{\phi}$) (SE)	$\sqrt{\Delta}$	$\sqrt{E[\Delta]}$ (SE)
Sample Size = 750					
-0.10	2.1128	-5.3334 (0.0354)	0.0766 (0.0492)	1.0563	1.0566 (0.0303)
-0.20	2.2601	-11.4443 (0.0333)	0.1975 (0.0580)	1.1301	1.1317 (0.0301)
-0.30	2.4716	-19.0865 (0.0328)	0.1785 (0.0683)	1.2358	1.2385 (0.0380)
-0.40	2.8368	-29.4002 (0.0297)	0.8319 (0.1083)	1.4184	1.4308 (0.0593)
-0.45	3.1990	-37.5008 (0.0273)	4.3151 (0.2665)	1.5995	1.6121 (0.1037)
Sample Size = 1,000					
-0.10	2.1128	-5.3239 (0.0312)	0.0578 (0.0433)	1.0563	1.0569 (0.0249)
-0.20	2.2601	-11.5035 (0.0293)	0.1033 (0.0483)	1.1301	1.1313 (0.0280)
-0.30	2.4716	-19.0636 (0.0275)	0.1800 (0.0572)	1.2358	1.2381 (0.0348)
-0.40	2.8368	-29.4680 (0.0258)	0.5782 (0.0885)	1.4184	1.4271 (0.0535)
-0.45	3.1990	-37.4445 (0.0238)	3.0434 (0.2170)	1.5995	1.6191 (0.0963)
Sample Size = 1,250					
-0.10	2.1128	-5.3159 (0.0276)	0.0682 (0.0381)	1.0563	1.0569 (0.0236)
-0.20	2.2601	-11.5280 (0.0267)	0.0517 (0.0434)	1.1301	1.1310 (0.0264)
-0.30	2.4716	-19.0891 (0.0256)	0.0846 (0.0527)	1.2358	1.2369 (0.0331)
-0.40	2.8368	-29.5306 (0.0229)	0.3844 (0.0796)	1.4184	1.4254 (0.0506)
-0.45	3.1990	-37.4603 (0.0211)	2.1408 (0.1700)	1.5995	1.6206 (0.0854)