RISK-SMOOTHING ACROSS TIME AND THE DEMAND FOR INVENTORIES:

A MEAN-VARIANCE APPROACH

Richard D. Farmer
Congressional Budget Office

INTRODUCTION

Businesses face uncertainty every day in the form of incomplete information needed for making optimal choices. A common assumption in modeling economic decisions involving time and uncertainty is that, from the decision maker's perspective, all such uncertainty is in the future. Conversely, economic agents are assumed to act as if current market parameters and technological factors are known with certainty. That view of uncertainty is inappropriate for a growing and important set of storable commodities. The presence of uncertainty surrounding current prices, for example, is well recognized for widely traded commodities such as oil, natural gas, precious metals, and grains. For a large percentage of transactions in those raw materials, traders contract today to purchase or sell goods for delivery some weeks or months in the future. A final price is determined only at the time of delivery, often depending on an initially agreed-upon reference price and some index of change.

It is also the case that the market shocks affecting those commodities can be more complex than economists commonly assume, with expectations of rising prices over time accompanying diminishing uncertainty, and vice versa. Changes in uncertainty about prices can be independent of changes in expectations. The impact of events that alter uncertainty about today's prices in particular may be quite different than current models of the demand for commodity inventories would predict or be capable of predicting.[2] Given that manufacturer stocks of raw materials are among the most volatile components of aggregate inventory investment, understanding the determinants of that volatility may be especially important for explaining total stock changes.[3] With the well-known contribution of aggregate inventory investment to business cycles, some appreciation of the complexity of market shocks and the role of current uncertainty also may contribute to a better understanding of macroeconomic performance.

To that end, this paper describes how inventory investment (or the rate of accumulation) by a competitive supplier of a continuously-produced good responds to new information that alters the firm's view of present and future expectations and uncertainty. The model indicates that, all else constant, an increase in current uncertainty relative to later periods would cause the rate of inventory accumulation

---

Richard D. Farmer: Congressional Budget Office, Washington, DC 20515. E-mail: richf500@rcn.com.

Eastern Economic Journal, Vol. 32, No. 4, Fall 2006

to rise. One example of how a focus on such changes in current uncertainty can be helpful relates to the relationship between certain government actions (that have an unintended effect of creating market uncertainty) and business decisions. In the context of price stabilization programs such as the Strategic Petroleum Reserve, official policy may be to sell government stocks in response to a supply disruption, with the goal being to moderate price increases. But if the market becomes more uncertain about the degree of government intervention, businesses may hold on to more of their own stocks than otherwise, with the effect being to push current prices upward—not downward.[4] That perverse response followed each announcement of plans to release SPR oil—in 1990 and 2003.

Past contributions to inventory theory under uncertainty provide a very limited capability for analyzing such information shocks: changes in expectations and uncertainty are not separable or, if they are, changes in current uncertainty relative to later periods are not separable. Much of that literature derives from the expected utility hypothesis of von Neumann and Morgenstern [1944], where decisions are based on the expected value of the utility of some outcome. That includes three well-known representations of inventory demand: the production smoothing model (and important variations such as the production-cost smoothing models of Blinder [1986] and Eichenbaum [1989]), the buffer-stock model (including the stock-out avoidance model of Kahn [1987]), and the optimal lot size (or S-s) model of Blinder [1981].[5] (Extensive reviews of the research on the demand for inventories appear in Blinder and Maccini [1991] and Ramey and West [1998]. Tests of models that combine the production-smoothing, cost-smoothing, and stockout hypotheses appear in Pindyck [1994] and Durlauf and Maccini [1995].)

The approach taken in this paper derives from the mean-variance theorem of Markowitz [1959] and Tobin [1958], in which individuals are assumed to maximize the value of a two-parameter function reflecting the aggregate mean and variance of a portfolio of choices.[6] In contrast to the expected utility approach, mean-variance decisions reflect the combined utility of the expected outcome and the variance of that outcome. The best-known applications of mean-variance utility have come in macroeconomics and finance, with Tobin's theory of liquidity preference describing the efficient substitution between risky assets and cash. Applications to microeconomics and the theory of the firm, as in this paper, are rare—although important exceptions appear in the analysis of storable commodity futures [McKinnon, 1967; Turnovsky, 1983].

Specifically, this paper applies the mean-variance utility function of Markowitz and Tobin to maximize utility across time in a deterministic control theory framework. That approach makes it possible to analyze the effects of market shocks that affect expectations separately from uncertainty or that alter expectations and uncertainly differently across time. The model side-steps the narrow focus of many optimal control problems on the properties of the optimal time paths for decision variables (e.g., the Hotelling effect). Using comparative statics, it instead becomes possible to investigate changes in the optimal rates of production and inventory investment at distinct points in time—most importantly the present.

Answering questions about the effect of time-varying changes in uncertainty is possible with that approach because uncertainty has an explicit role in the individual's utility function and because both the planning horizon (or terminal date) and the ending stock level are free. The model is consistent with the well-known production-smoothing motivation for holding inventories—with rising marginal production costs—plus a new, intertemporal risk-smoothing motivation.

That approach also has implications for an important area of research in inventory demand—how to account for the evidence that output levels are more volatile than sales. The paper reconciles the finding of excess production volatility with the production-smoothing motive for holding stocks by showing that output should respond to the same changes in expectations and uncertainty as does inventory accumulation, and in the same direction. The volatility of output relative to sales will increase as the costs of changing inventory levels (as distinct from the costs of holding inventories) increase and as businesses give more weight in their decisions to current changes in expectations and uncertainties (relative to future changes).

THE MODEL: OPTIMIZING MEAN AND VARIANCE IN CONTINUOUS TIME

This section develops a model of inventory decisions under price uncertainty that is based on mean-variance utility and solved using deterministic optimal control theory. The model represents two intertemporal goals of maintaining stocks: smoothing production cycles over time to minimize total production and inventory costs (and maximize total expected profits), and smoothing risks over time to minimizing total exposure to uncertainty. The model omits a third, static goal of maintaining stocks: facilitating transactions or production. In particular, to simplify the analysis, we ignore the possibility of sharply higher production or transaction costs from running low on stocks.[7] Following Turnovsky [1983], utility in that representation is linear in expected profits and the conditional variance of those profits. The utility from current sales and the terminal value of stocks each increase with expected price ($\alpha > 0$) and decrease with the conditional variance of price ($\beta < 0$).[8]

In the terms of control theory, the firm's objective is to find values for the control variables (production and inventory accumulation), the state variable (inventory level), and a costate variable λ (the net present value of one unit added to inventories in t) that will maximize the present- value sum of the utility from net cash flow over time and the terminal value of inventories. Decisions are subject to an initial level of inventories, but other than a requirement that inventory levels cannot be negative, no restraints exist on the planning horizon or on the terminal inventory level. In mathematical terms (see Table 2), the objective is to maximize

$$\int_0^T U(q,i,N,t)e^{-rt}dt + \varphi(N,T)e^{-rT}$$

$$\text{(1)} \qquad \text{where} \qquad U = \alpha[E(p)(q - i) - C(q) - H(N,i)] + \beta[\mathrm{VAR}(p)(q - i)^2]$$

$$\varphi = \alpha E[p(T)]N(T) + \beta\ \mathrm{VAR}[p(T)]N(T)^2$$

The necessary conditions for an optimal solution are well known, but we state them here because each will be referred to in the process of deriving a solution for inventory change:

(2) $U_q e^{-rt} = 0$ (optimality condition for q)

$$(\text{or } \{\alpha[E(p) - C_q] + 2\beta\,\text{VAR}(p)(q - i)\}e^{-rt} = 0)$$

(3) $U_i e^{-rt} + \lambda = 0$ (optimality condition for i)

$$(\text{ or } \{-\alpha[E(p) + H_i] - 2\beta\,\text{VAR}(p)(q - i)\}\, e^{-rt} + \lambda = 0)$$

(4) $\dot{N} = i$ (state equation, subject to the terminal condition $N(0) = N_0$)

(5) $\dot{\lambda} = -U_N e^{-rt} (= \alpha H_N e^{-rt})$ (costate equation)

(6) $\lambda(T) = \varphi_N(T) e^{-rT}$ (transversality condition for free N(T))

$$(= \{\alpha E[p(T)] + 2\beta\,\text{VAR}[p(T)]N(T)\}e^{-rT})$$

(7) $U e^{-rt} + \lambda i + \varphi e^{-rt} \left(\text{at } t = T\right) = 0$ (transversality condition for free T)

TABLE 1
Definitions

Functions, Variables, Parameters	**Definition**
$U(q, i, N, t)$	Utility of net cash flow, a function of the mean and conditional variance of net cash flow in t
$\varphi(N, T)$	Utility of terminal inventories, a function of the mean and conditional variance of the value of inventories at the end of the planning horizon
$C(q)$	Production costs in t, a function of the production rate
$H(N, i)$	Inventory holding costs in t, a function of inventories and inventory change
Q	Production in t, a control variable
I	Rate of addition to inventories in t, a control variable
$Q - I$	Sales in t
N	Inventories in t, the state variable, with the terminal level $N(T)$ free
λ	Costate variable, or the net present value of one unit added to inventories in t
T	Free terminal time for the planning horizon, a dependent variable
P	Sales price
R	Discount rate
$E(p) = E(p \mid \Omega)$	Price expected in t, given the available information set
$\text{VAR}(p) = \text{VAR}(p \mid \Omega)$	Conditional variance of the price in t, given the available information set
Ω	The current ($t = 0$) information set relevant to future expectations and uncertainty
α, β	Parameters of the mean-variance utility function

Note in that representation that the solutions for production and inventory accumulation in $t = 0$ are actual changes, while changes in $t > 0$ represent current plans. Revenues come from planned sales, or $q - i$. As a competitive supplier, able to sell as much it finds profitable at any price, the firm plans how much to sell and from what combinations of production and stock draw. The rate of inventory accumulation may take a positive or negative value, but production is assumed to take place in a range that guarantees that sales will be positive.[9] Implicit in the requirement for positive sales and stocks are assumptions that the marginal utilities of sales (U_{q-i}) and stocks (U_N) and the marginal salvage value of stocks (φ_N) are all positive. Similarly, production cost is a convex function of q, such that $C_q > 0$ and $C_{qq} > 0$. And inventory holding cost is a linear function of the current inventory level ($H_N > 0$ and $H_{NN} = 0$) and an increasing function of the absolute value of i ($H_i > 0$ and $H_{ii} > 0$), where the effects of i and N are separable ($H_{Ni} = 0$).[10]

The future path of the commodity price—or, more correctly, the path of p relative to other prices—is uncertain, while the production and storage cost relations and the discount rate are known with certainty.[11] For simplicity, covariances between price and all costs are zero. All information relevant to forming expectations about prices and assessing uncertainty about prices is contained in the vector Ω (where $\Omega \in [E(p), VAR(p)]$ for all t). Expectations of prices in t reflect the set of information currently available (in $t = 0$), so that $E(p) = E(p \mid \Omega)$. Similarly, uncertainty about prices in t reflects the subjective variance of prices, conditional on currently available information, so that $\mathrm{VAR}(p) = \mathrm{VAR}(p \mid \Omega)$.[12] That conditional variance of prices in each period is assumed to equate with the nondiversifiable risk of the decision maker at that time; the diversifiable risk is implicit in the constant discount rate. The existence of a solution for this model relies on assumptions that price expectations and price uncertainty (and the paths for market shocks to those variables) are well-behaved, continuous functions of time. Also, the model assumes that no irreversible costs result from current production or inventory investment decisions.[13]

That approach contrasts with previous research in several ways, in addition to the mean-variance utility function and the optimal control framework. First, it allows for two control variables, not one. The standard inventory model generally represents a competitive industry that can only sell the market demand, so sales are exogenous. In that case, only production or stock draw can be a control variable, with the other calculated as a residual. Second, the assumption that H is linear in the inventory level and a nonlinear function of inventory investment is a marked departure from the traditional assumption of holding costs as a quadratic function of inventory levels alone. However, the new assumption of H increasing with absolute i is consistent with the goal of such models to provide a cost check on changing stock levels.

A COMPARATIVE STATICS ANALYSIS OF THE GENERAL SOLUTION

How would the optimal rate of inventory accumulation at a distinct point in time (most importantly, the present) change in response to a current information shock ($d\Omega$) that alter the time paths for price expectations or for the conditional variances of expected prices? An analytical solution describing the impact of new information comes from the total differential of the system of equations presented by the necessary conditions for maximum utility over the planning horizon.[14] As intermediate steps, the total differential for that system reduces to three equations in terms of the three unknowns, dT^*, dq^*, and di^*, and the four independent variables, $dE(p)$, $dE[p(T)]$, $d\mathrm{VAR}(p)$, and $d\mathrm{VAR}[p(T)]$.[15] (The asterisks indicate optimal values.) That is, the relevant information set for answering that question is fully characterized by the changes in expectations and uncertainties in just two periods, the present and the terminal date. Partial solutions for dT^*, dq^*, and di^* are described in the Appendix. The general solution for the optimal change in inventory accumulation in t is

$$di^* = \tag{8}$$

$$\frac{\left(\begin{array}{l} U_{i\Omega} + \varphi_{N\Omega} e^{-r(T-t)} - U_{iq} \dfrac{U_{q\Omega}}{U_{qq}} - \\ \left[U_N(T) - r\varphi_N(T) + \dot{\varphi}_N(T)\right] \dfrac{U_\Omega(T) + i(T)\varphi_{N\Omega}(T) + \varphi_\Omega(T)}{\dot{U}(T) \ + \ i(T)\dot{\varphi}_N(T) \ + \ \dot{\varphi}(T)} e^{-r(T-t)} \end{array} \right)}{-[U_{ii} - U_{iq} \dfrac{U_{qi}}{U_{qq}}]} d\,\Omega$$

That complex expression can be simplified to demonstrate that information shocks can affect inventory decisions in three ways, through cost- and risk-arbitraging incentives, current production costs, and terminal holding costs. First note that the denominator of equation (8) can be shown to be strictly positive, since

$$-[U_{ii} - U_{iq} \frac{U_{qi}}{U_{qq}}] = -\frac{U_{qq}U_{ii} - U_{iq}U_{qi}}{U_{qq}}$$

$$= -\frac{[-\alpha C_{qq} + 2\beta\,\mathrm{VAR}(p)][-\alpha H_{ii} + 2\beta\,\mathrm{VAR}(p)] - [-2\beta\,\mathrm{VAR}(p)]^2}{-\alpha C_{qq} + 2\beta\,\mathrm{VAR}(p)} > 0$$

Thus, the key to understanding the direction of di^* is the numerator of equation (8). Expanding the model's term for market shocks ($d\Omega$) into its four components ($dE(p)$, $dE[p(T)]$, $d\mathrm{VAR}(p)$, and $d\mathrm{VAR}[p(T)]$), the numerator can be broken down to identify three sources of influence on the rate of inventory accumulation:

$$\begin{aligned} &(U_{i\Omega} + \varphi_{N\Omega} e^{-r(T-t)}) d\Omega = \\ &\quad = \alpha\left\{dE(p) - dE\left[p(T)\right]e^{-r(T-t)}\right\} - \\ &\quad\quad 2\beta\left\{(q-i)d\,\mathrm{VAR}(p) - N(T)d\,\mathrm{VAR}\left[p(T)\right]e^{-r(T-t)}\right\} \end{aligned} \tag{9.1}$$

$$-U_{iq}\frac{U_{q\Omega}}{U_{qq}} d\Omega = 2\beta\,\mathrm{VAR}(p)\frac{\alpha dE(p) + 2\beta(q-i)d\,\mathrm{VAR}(p)}{-\alpha C_{qq} + 2\beta\,\mathrm{VAR}(p)} \tag{9.2}$$

and

$$(9.3)\quad -[U_N(T) - r\varphi_N(T) + \dot{\varphi}_N(T)]\frac{U_\Omega(T) + i(T)\,\varphi_{N\Omega}(T) + \varphi_\Omega(T)}{\dot{U}(T) + i(T)\dot{\varphi}_N(T) + \dot{\varphi}(T)}e^{-r(T-t)}d\Omega$$

where $\dot{U}(T) + i(T)\,\dot{\varphi}_N(T) + \dot{\varphi}(T) < 0$ (the condition for existence of a finite T)

$$-[U_N(T) - r\varphi_N(T) + \dot{\varphi}_N(T)] > 0$$

and, rearranging terms,

$$\begin{aligned}&[U_\Omega(T) + i(T)\varphi_{N\Omega}(T) + \varphi_\Omega(T)]d\Omega\\ &\quad = \alpha\left\{q(T) + N(T)\right\}dE[p(T)] +\\ &\qquad\quad \beta\left\{[q(T) - i(T)]^2 + N(T)^2 + 2i(T)N(T)\right\}d\,\mathrm{VAR}[p(T)]\end{aligned}$$

Equation (9.1) represents the combined production-smoothing and risk-smoothing incentives for holding stocks. As a basic finding of the paper to be tested empirically, it indicates

Proposition 1. *Market shocks that cause expected profitability to rise (fall) in the current period relative to later periods or cause current uncertainty to fall (rise) relative to later periods will, all else constant, cause the rate of inventory accumulation to drop (rise).*

Two important parts of that "all else" derive from the independent effects of an information shock on current production and the terminal value of stocks. Equation (9.2) represents the effect on *di** of changes in current production costs. Specifically, higher current expectations or lower current uncertainty will cause the current production and, as a result, production costs to rise. That will reduce the marginal utility of inventory investment, diminishing the drop in accumulation rates that would otherwise result from those shocks to current market conditions.

Equation (9.3) represents the effect on *di** of changes in the utility associated with terminal stocks. Specifically, higher expectations or lower uncertainty in later periods may increase the incentive to build stocks for later sale, but those same shocks will also add to the costs of continuing to hold stocks at that time (due to storage costs, the time value of capital, and continuing changes in conditional uncertainty with time). Those added costs diminish the rise in accumulation that would otherwise result from shocks to terminal market conditions that increase expected price or reduce the conditional uncertainty in *T*.

Comparison of the coefficients to the market-shock terms suggests that the production-smoothing/risk-smoothing effect (from equation (9.1)) on change in inventory accumulation would dominate the current production cost effect (from equation (9.2)). The dominance of production-smoothing/risk-smoothing over the terminal-value effect (from equation (9.3)) is not as clear, especially in later periods; but because that terminal effect is discounted from T, its importance for change in the current rate of accumulation would probably be secondary as well. That points to the special case

Proposition 2. *If shocks to expectations or uncertainty are relatively uniform over time, the combined production- and risk-smoothing effect on inventories would be zero, and the effect of changes in current production costs or the value of terminal stocks must dominate.*

"Uniform" means, for example, $dE(p)$ equals $dE[p(T)]e^{-r(T-t)}$. Which effect, equation (9.2) or equation (9.3), would dominate is an empirical issue. Circumstances may well exist that cause the change in inventory accumulation to be counter to what a direct comparison of current and terminal marginal utilities (from equation (9.1)) might indicate on its own.

FITTING THE MODEL TO NATIONAL DATA—CONCEPTS AND DATA SOURCES

The mean-variance model of inventory change is inherently a microeconomic tool. However, estimation of the model at the firm or industry levels presents many challenges, not the least of which are the practical difficulties of guessing now what businesses may have been thinking at the time about uncertain events that are now in the past. That said, it is still possible to demonstrate the utility of the approach as well as some of the problems that analysts would need to confront in such industry-level estimations by fitting the model to aggregate national data on inventory changes and information that can represent data on price expectations and price uncertainty.

To simplify the task, the estimation addresses only the combined production- and risk-smoothing incentives for holding stocks, as described in Proposition 1. That assumes that all uncertainty is in the price and omits the effects of current production costs and terminal stock values, as described in Proposition 2. Accordingly, the relationship to be tested is:

$$\begin{aligned}\Delta i(t) = \gamma_1\left\{\Delta E_t[p(t)] - \Delta E_t[p(t+1)]e^{-rt}\right\} + \\ \gamma_2 sales(t)\left\{\Delta V_t(p) - \Delta V_t[p(t+1)]e^{-rt}\right\}\end{aligned} \tag{10}$$

(See Table 2 for a description of terms and data sources.) From Proposition 1, we expect that the coefficient on the term for change in price expectations (γ_1) will be negative (since $-\alpha < 0$), and the coefficient for the change in uncertainty (γ_2) will be positive (since $-2\beta > 0$). The reasons for focusing only on price uncertainty (rather than, say, cost uncertainty) were of expedience: the only forecast information relevant to expectations and uncertainty that is available in a time series relates to forecasts of inflation. It also was expected that any proxy for such cost information at the aggregate national level (for example, from the producer price index) would be highly correlated with final prices and, hence, would not add to the explanation. The reason for omitting information on the effects of current production costs and the effects of terminal values relates to the identity problem that arises because those terms are based on expressions for expectations and uncertainty that appear in the production- and risk-smoothing term. A multi-stage estimation that first approximates the parameters of those terms may help to yield useful time series, but that work is beyond the scope of this paper.

TABLE 2
Data Sources for Regression Variables

Regression Term	**Definition, Period, and Source**
Dependent Variable:	
$\Delta i(t) = i(t) - i(t\text{-}1)$	Change in the real change in private inventories ($2000) over the past quarter, 1982:1 - 2001:3 (Bureau of Economic Analysis)
Independent Variables:	
$\Delta E[p(t)] = \text{PGDP}(t+1) - \text{PGDP}(t)$	Change in the actual GDP price deflator over the forthcoming quarter, 1982:2 - 2001:4 (Bureau of Economic Analysis)
$\Delta E[p(T)] = E[\text{PGDP}_{t+1}(t+2)] - E[\text{PGDP}_t(t+1)]$	Change in the mean of the next-quarter forecast of PGDP over the forthcoming quarter, 1982:2 - 2001:4 (Survey of Professional Forecasters)
$\Delta V[p(t)] = \text{var}[\text{PGDP}_t(t+1)] - \text{var}[\text{PGDP}_{t-1}(t)]$	assumed to equal zero
$\Delta V[p(T)] = \text{var}[\text{PGDP}_{t+1}(t+2)] - \text{var}[\text{PGDP}_t(t+1)]$	Change in the calculated variance of forecaster disagreement for PGDP over the forthcoming quarter, 1982:2 - 2001:4 (Lahiri and Liu (2003))
$\text{sales}(t) = q(t) - i(t)$ and $= N(T)$*	Final sales ($2000), 1982:1 - 2001:3 (Bureau of Economic Analysis)
$r = r(T\text{-}t)$	Federal funds rate, 1982:2 - 2001:4

NOTE: PGDP is the implicit price deflator for gross domestic product. $\text{PGDP}_t(t+1)$ refers to the expectation (E) or the conditional variance (V) of deflator forecasts for t+1 that are formulated in t.

* The use of current sales as a proxy for terminal stocks ($N(T)$) was a convenience, based on the assumption that the factors that affect current sales would affect the optimal inventory level at the end of the planning horizon in the same way. However, terminal stocks would likely be some multiple of sales, and the estimated value of γ_2 is likely overstated by that same multiple.

An initial challenge for the estimation was deciding how to measure market expectations and uncertainties. The difficulty of defining those concepts encompasses the issue of informational efficiency, or how much of the information that is relevant to assessing expectations and risks is actually utilized by the market. In contrast to the nomenclature of mean values and conditional variances, actual expectations and uncertainty are perceptions that need not derive from any empirical data or, when they do, need not fully reflect the relevant data. The literature on the predictability of inflation points to two general approaches: one based on statistical observations of past prices and the other on observations of different forecasts from different models.[16] Within the forecast-based approach, common measures of uncertainty are those based on forecast disagreement, the variance of the average individual forecaster errors, and the conditional variance of forecast error.

This study adopts the second approach, using an average forecast of the implicit GDP deflator to measure future price expectations and a measure of forecast disagreement developed by Lahiri and Liu [2003] to measure the future conditional variance. The choice of uncertainty measures was largely one of convenience. Engle finds a close correspondence between his ARCH variances and various measures of dispersion calculated from a survey of inflation forecasts, indicating that it may ultimately matter only little which approach—past prices or forecast surveys—one adopts. Similarly, Lahiri and Liu find that there is little difference among three measures of forecaster uncertainty—forecast disagreement, variance of the average individual forecaster errors, and the conditional variance of forecast error—except in periods of particularly unstable or high inflation.

Specifically, to represent changes in price expectations, this analysis uses the mean of the next-quarter forecast of the GDP deflator (PGDP) for terminal expectations and the historical price deflator for current expectations. The time series of mean forecasts is from the Survey of Professional Forecasters (SPF), conducted by the Federal Reserve Bank of Philadelphia. (Each quarter, survey participants submit their forecasts of the implicit GDP price deflator for six quarters and annual forecasts for the current and following year. They also submit information that describes their uncertainty about their inflation forecasts for the current and following year in the form of probabilities that they attach to a number of discrete intervals of possible change in inflation.)

That approach for expectations was supported by the fact that the forecast and the historical price changes are relatively independent. In contrast, the changes in individual quarterly and annual forecasts were highly correlated with one another, so that little insight could be gained by relying, for example, on a next-quarter forecast for current expectations and an annual forecast for terminal expectations. (The use of next-quarter forecasts for current expectations also should minimize any problem of reverse causality—current inventory change is only one small component of the aggregate demand that would influence aggregate prices.) Moreover, for terminal expectations, the next-quarter forecasts were expected to be more useful in a regression than the SPF forecasts with a longer-term outlook because there is more quarter-to-quarter variability in the near-term series.[17] (The use of that forecast series does not imply the planning horizon for inventories is one quarter.)

To represent changes in price uncertainty, the analysis uses a measure of forecast disagreement to represent terminal-period uncertainty and assumes current-period uncertainty is constant.[18] The measure of forecast disagreement was developed by

Lahiri and Liu, based on the SPF individual probability distributions for forecast uncertainty. It presents the successive quarterly variances, adjusted to take out systematic bias across forecasters and leave a series that represents disagreement attributable to different information sets (not different modeling approaches or forecasting goals). The reason for assuming no change in current-period uncertainty—despite a belief that the current period is critical to understanding inventory changes—is the absence of information on current uncertainties that is both independent of the measure for future uncertainty and calculated on a comparable scale.[19]

A further challenge was to decide how to apply the discrete-period data on inventory change, expectations, and uncertainty to approximate the instantaneous changes that are implicit in the comparative statics of model. The appropriate regression appears to relate the change in inventory accumulation over the past quarter to the changes in expectations and uncertainties over the forthcoming quarter. The logic behind that advanced-lag structure is that the information on which forecasters base their SPF submissions was likely to be available to businesses months earlier, when those businesses were making the decisions that lead to the observed inventory change.

REGRESSION RESULTS FOR A NATIONAL MODEL

Regression results from estimating equation (10) using ordinary least squares (with no adjustments for serial correlation) are

Regression Results

		γ_1	γ_2
Coefficient (*t* Statistic):		-6.640	0.003
		(-2.43)	(1.64)
Value of rho		0.018	0.106
R-squared	0.105		
F Statistic	4.53 (Significance of F: 0.01)		
Sign Test[20]	2.94 (Significance of Sign Test: 0.01)		
Number of observations	78		

Note: Results derived using the regression function available in the Microsoft Excel data analysis tools.

Several observations are in order. The overall fit is significant as evidenced by the F statistic, showing that the hypothesis of no-relationship can be rejected with nearly 99 percent confidence. The estimated coefficients have the anticipated signs: negative for γ_1 and positive for γ_2. The explanatory power of the coefficient on the term for relative change in current and future prices (γ_1)—the production smoothing term—is especially strong. The hypothesis that that coefficient equals zero can be rejected with more than 98 percent confidence. The coefficient for the relative change in current and future uncertainty (γ_2)—the risk smoothing term—is also significant, but less so (at nearly 90 percent confidence). Sensitivity regressions (not reported here) indicate that the introduction of seasonal variables would not enhance the significance of the relationship tested.

It is important to note that the ability of the model to explain all variation in the change in inventories is limited, as evidenced by the low R-squared (10.5 percent).

That is to be expected for several reasons, related to omitted-variable and error-in-variable effects. Omitted variables for the effects of current production costs and terminal stock values may introduce a bias in the regression, since there is no way to effectively hold those effects constant, but it is difficult to know what that bias may be. Proposition 2 summarizes conditions under which their omission could be significant. A further omitted variable is the change in uncertainty in the current period (assumed to equal zero).

Error-in-variable problems are likely to derive from the use of aggregate data and from the imprecisions of the statistical approximations to the concepts of expectations and uncertainty, including any imprecisions that result from the use of quarterly data and the lag structure for representing when expectations are formed. It is possible, if not likely, that independent changes in inventories for different industries will offset one another and, hence, dampen aggregate changes. The question of bias resulting from measures of expectations and uncertainty is an open one, with significant research, described earlier, concluding that few systematic differences among measures exist but with no one claiming that any measures fully capture those concepts.

Yet within that partial explanation, the model itself performs well, as evidenced by its ability to explain the turning points in historical inventory investment. The change in real inventories is a highly volatile series, yet it is the changes in direction of that flow—that is, the turning points—in combination with the accelerator principle that are critical to understanding and forecasting cyclical changes in domestic output. One statistic that supports the usefulness of the mean-variance approach for predicting such turning points is a sign test that records how well the model explains changes in direction of the dependent variable. In 52 of 78 observations, the values of di that the estimated model predicts are changing in the same direction as the actual values of di. The normal statistic z, calculated to test whether the signs of the two series match, is 2.94. That indicates the hypothesis that the turning points are random can be rejected with over 99 percent confidence.

Figure 1 presents a comparison of the observed and predicted values. It shows how much of the total variation remains to be explained (consistent with the low R-squared) as well as how many of the turning points the model captures (consistent with the sign test). Note that the ability to explain the turning points in such a volatile data series as change in inventory investment is accomplished without the introduction of distributed-lag dependent variables. Moreover, both the changes in expected prices and the changes in uncertainty, which are only weakly correlated (the correlation coefficient is -0.12), appear to have important roles in predicting those turning points.

EXPLAINING EXCESS PRODUCTION VOLATILITY IN INDUSTRY DATA

Further support for the mean-variance model of inventory demand comes from its ability to help resolve a long-standing issue in the literature on inventory demand. That issue is excess production volatility and the apparent conflict between a prediction of the basic production-smoothing model (that inventories, not output, should vary to accommodate demand shocks) and observations that the output of most industries is more variable than their sales. The benchmark for that observation is data from Blinder [1981; 1986] that show that the variance of production exceeds the variance of sales in all major retail and manufacturing industries save two (where they are virtually equal). Specific attributes of the mean-variance model that would lead to that result are the separate representations of production and inventory investment as decision variables (with holding costs dependent on the rate of inventory accumulation) and the separate representations of current and future uncertainty. This section develops two propositions that may help guide future research into differences in excess volatility across industries.

FIGURE 1
Observed and Predicted Quarterly Changes in Real Inventory Investment ($2000)

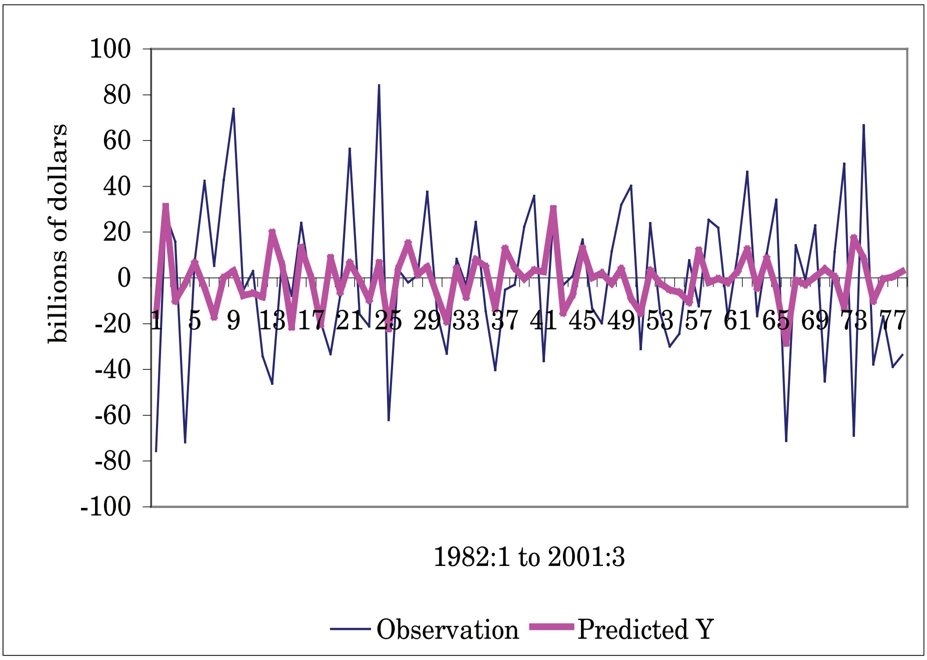

Efforts to resolve the excess-volatility conflict have motivated important alternative explanations of why businesses hold inventories. They include models of production-cost smoothing—with cost shocks in the form of input prices [Blinder, 1986] or technology [Eichenbaum, 1989]—and models of stock-out avoidance [Kahn, 1987]. The limited explanatory power of those theories in isolation and their more general success in combination [Pindyck, 1994; Durlauf and Maccini, 1995] suggests that observations of excess volatility may require different explanations for different industries or periods. The prospect of special theories for special cases is unappealing to many.

Starting with the accounting identity that the change in utility-maximizing sales (*dsales**) equals the difference between dq^* and di^*, an observation of excess production volatility ($dsales^*/dq^* < 1$) would require that dq^*/di^* be positive and greater than one.

But that would be inconsistent with the basic prediction of the production-smoothing model. We ask what circumstances would give that same result in the present model? From the partial solution for dq^* (see Appendix, equations (A5)) and the first term of equation (8)—that is, focusing on the combined production- and risk-smoothing motives for holding stocks—excess production volatility would be consistent with

$$\frac{dq^*}{di^*} = -\frac{\dfrac{U_{q\Omega}}{U_{qq}}d\Omega}{\left\{\begin{array}{l} [U_{i\Omega} + \varphi_{N\Omega}e^{-r(T-t)}]d\Omega \\ -[U_{ii} - U_{iq}U_{qi}/U_{qq}] \end{array}\right\}} - \frac{U_{qi}}{U_{qq}} > 1 \qquad \text{(11.1)}$$

(defined for $d\Omega$ .ne. 0)

Or, equivalently,

$$\frac{U_{q\Omega}d\Omega}{[\ U_{i\Omega} + \varphi_{N\Omega}e^{-r(T-t)}\ ]\,d\,\Omega} - \frac{U_{qi} + U_{qq}}{[\ U_{ii} - U_{iq}\ U_{qi}/U_{qq}\ ]} < 0 \qquad \text{(11.2)}$$

where

$$[U_{ii} - U_{iq}U_{qi}/U_{qq}] < 0$$

(from equation (8))

and

$$U_{qi} + U_{qq}(= -\alpha C_{qq}) < 0.$$

That result suggests there are two basic relationships underlying the relative change in output and inventory accumulation that could yield excess production volatility. The clearest relationship can be seen in the second term of the inequality in equation (11.2). It indicates that the relative change should reflect changes in the marginal utilities of production and inventory accumulation that are directly attributable to the respective changes in production and accumulation. That second term reduces to

$$-\frac{[-\alpha C_{qq}][-\alpha C_{qq} + 2\beta\,\mathrm{VAR}(p)]}{[-\alpha C_{qq} + 2\beta\,\mathrm{VAR}(p)][-\alpha H_{ii} + 2\beta\,\mathrm{VAR}(p)] - [-2\beta\,\mathrm{VAR}(p)]^2} \qquad \text{(11.3)}$$

It is always negative. In a simple case with current uncertainty VAR(p) (or risk averseness β) approaching zero, it becomes $-C_{qq} / H_{ii}$. In general, as a basic finding of the paper that may guide future research, it indicates

Proposition 3. *The greater the convexity of inventory holding costs (with respect to inventory investment) relative to production costs, the greater the volatility of output relative to sales.*

The standard production-smoothing model, which represents holding costs as a quadratic function of the total stock level alone (and not the rate of accumulation) and commonly represents inventory demand as a residual, misses that dampening effect on inventory investment and, hence, sales.

The second relationship that would contribute to excess volatility comes from the first term of equation (11.2). It indicates that the relative change in output and inventory accumulation should reflect the changes in marginal utilities that are directly attributable to changes in expectations and uncertainty. That first term becomes

$$\frac{\alpha dE(p) + 2\beta(q-i)d\,\mathrm{VAR}(p)}{-\alpha\{dE(p) - dE[p(T)]e^{-r(T-t)}\} - 2\beta\{(q-i)dVAR(p) - N(T)d\,\mathrm{VAR}\,[p(T)]e^{-r(T-t)}\}} \tag{11.4}$$

It may take either sign. But because the second term of inequality (11.2) is strictly negative, sufficient conditions for excess production volatility arise whenever market shocks only affect current expectations and/or uncertainty—so that expression (11.4) would be strictly negative—or only affect future expectations and uncertainty—so that expression (11.4) would equal zero. Of course, a premise of this paper is the likelihood that many shocks will affect both current and future expectations and uncertainties. Given that, expression (11.4) and, hence, the total value of inequality (11.2) would be negative in two further circumstances. One is whenever current and future expectations (or uncertainty) change in the same direction. The other is whenever the absolute value of the current changes exceeds that for the present value of future changes. If such shocks are more commonplace than ones where current and future values have opposite signs, or if a greater weight on current changes is likely because of discounting and the general lack of information about the future, we can state

Proposition 4. *The more that businesses myopically concern themselves with changes in current prices and risks, the greater the volatility of output relative to sales.*

The standard inventory model, by implicitly assuming that all uncertainty is in the future, would miss that effect of changes in current expectations and uncertainty.

That said, market shocks that would lead to excess sales volatility are theoretically possible. If future expectations consistently increase relative to current expectations—for example, there is a sustained outlook for rising prices—then inequality (11.4) could be positive. If that value then exceeds the absolute value of the second term of the expression (11.3), excess sales volatility would prevail. The Blinder data on sales and output for a broad range of industries provide strong evidence that such shocks are not commonplace.

THE ADVANTAGES OF MEAN-VARIANCE UTILITY—A CASE FOR EXPANDED USE

The key to obtaining those results is the assumption that utility should be modeled as a separable function of expected returns and uncertainty. Yet despite the apparent explanatory power of a mean-variance utility function, that approach has seen little application in representing decision making under uncertainty in such basic stock-flow problems as inventory demand, exhaustible resource production, or capital investment. The reason may come from theoretical work that initially emphasized the advantages of the expected utility approach or possibly from common terminology that has clouded the distinction between mean-variance and expected utility.

Early criticisms of the mean-variance approach focused on the lack of generality of the two-parameter utility function and the concern whether markets could indeed rank-order choices based on those parameters (for example, see Hirshleifer [1965] and Feldstein [1969]). However, later work that has gone largely unnoticed seemed to satisfy those criticisms by demonstrating conditions under which expected utility and mean-variance utility yield equivalent results without imposing unreasonable restrictions on consumer preferences or the nature of risks. In particular, Levy and Markowitz [1979] have shown that both approaches generate the same ranking of choices if the random parameters of the model are normally distributed and both utility functions are consistent with risk averse behavior. (With expected utility, a quadratic utility function with declining marginal utility ($U'' < 0$) satisfies that requirement; with mean-variance utility, a negative coefficient on the conditional variance term β is sufficient.) More generally, Sinn [1983], Meyer [1987], and others have shown that the mean-variance utility generates a consistent ranking of choices so long as the outcomes of the choices are all linearly related to one another (or, in Meyer's terms, differ by location and scale factors).[21] That condition is met for the most common types of economic problems, where a single-outcome variable (for example, profits) depends on choice variables (production, inventory investment, etc.) and linearly depends on one random parameter (for example, price). In that circumstance, the necessity of assuming normalcy of the random parameter in order to gain consistency goes away, too.

Despite the withdrawal of some early theoretical concerns, the mean-variance approach of Markowitz and Tobin has fallen into general disuse for microeconomic applications. That fact may itself have gone unnoticed in part because the literature on capital asset pricing models (CAPM) has largely co-opted the terms "portfolio theory" and "mean-variance" theory (for example, see Rubinstein [1973]). CAPM, however, derives from expected utility under the specific assumptions of declining marginal utility and a normal distribution for the market portfolio's return—not from mean-variance utility. Also, the search for consistency may in some ways have colored perceptions of mean-variance utility as a tool that, at its best, can only replicate what expected utility offers.

To the contrary, mean-variance utility conveys several distinct advantages over expected utility as it has variously been applied to inventory demand, especially if the goal is to model the effects of time-varying changes in expectations or uncertainty. Those advantages all derive from the difficulties of mathematically representing choices based on expected utility. As Meyer points out, a great value of representing some problems in a simple mean-variance framework is that that approach is much more conducive to the comparative statics analysis of problems in economics and can generate simple analytical solutions. Those single equations can serve as objects of econometric analyses to identify otherwise unknown parameter values—as demonstrated in this paper.

Relative advantages from a mean-variance approach also derive from the restrictive assumptions that analysts must make to apply the expected utility hypothesis. One such assumption is that of risk neutrality ($U'' = 0$). A rationale for that restriction comes from the view that businesses represent the interests of shareholders, who in turn hold diversified portfolios.[22] The focus of the production-smoothing, stock-out avoidance, and *S-s* models of inventory demand is then on the different outcomes that a stochastic process yields and on the cumulative effects of those outcomes on decisions. But in the production-smoothing models, where certainty-equivalent values are substituted for expected values in the mathematical formulation, the capability for analyzing the impact of changes in uncertainty on production and inventory decisions is diminished. Changes in expectations cannot be addressed independently of changes in uncertainty in any period.[23] In the stock-out avoidance and *S-s* models, variance of the stochastic variable remains as an element of the utility function and, hence, still affects decisions despite risk neutrality. But that means that the variance of only the future values enters the utility function or that the variance is represented as constant over time. Either way, change in the variance can only affect production decisions in subsequent periods. Those models are unable to address the effects of changes in current uncertainty relative to later periods.

In contrast, with optimization based on mean-variance utility and risk averse behavior, the effects of expectations and uncertainty can be independent in all periods. Support for assuming risk averse behavior in that utility comes from the literature on agency theory and corporate governance, which supports the view that the personal interests of management or of particular trading blocks may run counter to those of well-diversified stockholders (for example, see Morck et al. [1988]). An assumption of risk averse behavior is critical for the ability to analyze the effects of time-varying uncertainty.

Risk averseness is also a key to the mathematical solution of the problem—not an obstacle to solution. (See discussion in Appendix.) Inventory models based on expected utility are often formulated as problems in dynamic programming, with a solution generally requiring the analyst to assume an infinite planning horizon and a fixed ending level of stocks.[24] The logic supporting those assumptions is straightforward: the expected utility from future sales will be positive so long as stocks are positive, so there is no logical point to cut off planning. The advantages of the mean-variance approach, however, are best realized by formulating the model as a problem in deterministic optimal control [Pontryagin et al., 1962] because of the flexibility that control theory offers for representing both the planning horizon and ending stock level as free variables.[25] Indeed, this paper demonstrates that the existence of a unique solution with risk averseness may depend on those free terminal conditions and the assumption that individuals are increasingly uncertain about (and hence associate less utility with) more distant events. In one sense, the costs of considering decisions beyond some date become too great. The ability to cut off consideration of how current actions affect increasingly distant periods may be a necessary part of making decisions. The notion of a finite planning horizon that changes with market conditions is certainly consistent with observations of diminishing but variable open interest in commodity futures contracts of increasing maturity.

FINAL OBSERVATIONS

Support for the usefulness of a mean-variance model of inventory demand comes from its ability to explain cyclical changes in aggregate inventory investment over the past two decades. However, the model is inherently a microeconomic tool. Its value may lie more in the perspective it brings to the general understanding of how current market events or government policies can affect current inventory decisions than in its ability to accurately predict inventory changes over time. Economists will always struggle with the definitions and measurements of expectations and uncertainty that would be necessary for such forecasts. In contrast, analysts who closely follow particular industries are generally able to sense changes in market expectations and uncertainties as they occur, and will know whether they are growing or diminishing over their planning horizon. Those analysts also are likely to recognize the role that current market uncertainty and producer incentives for risk smoothing play in determining inventory investment.

For researchers who want to better understand inventory change, this paper suggests new variables to consider that can help explain the relative volatility of output and sales in different industries. It indicates the value of new research to help distinguish current and forward measures of uncertainty, as opposed to questioning what single concept best represents uncertainty in general. (Better measures would yield better explanations of inventory investment.) And it points to a need for new information on the effects of current production costs and the terminal value of stocks on storage decisions, so that those omitted variables can be brought into the explanation.

APPENDIX

The first step in the exercise to derive an equation for the optimal change in inventory investment is to understand how information shocks affect the optimal planning horizon dT*. That relationship comes from equations (6) and (7), which yield the terminal condition for total utility

$$U(T)+\varphi_N i(T)+\varphi(T)=0 \tag{A1}$$

and from equation (2) and the combination of equations (3) and (6), which together yield the terminal conditions for marginal utility

$$U_q(T)=0 \tag{A2}$$

and

$$U_i(T)+\varphi_N(T)=0 \tag{A3}$$

Those conditions require that sales cease (for planning purposes) when the total and marginal utility of profits (including the salvage value) are both zero. A finite T will exist if the total and marginal utilities of profits are both falling near *T*. Several conditions would guarantee the existence of a finite planning horizon. One is expectations of either rising unit costs or declining prices with time.[26] A less restrictive, behavioral assumption would require only that businesses be risk averse and that the conditional variance of prices be rising at *T*. In particular, from the perspective of the present (t = 0), firms expect the conditional variance of prices in successive periods to rise solely as a function of time ($d\text{VAR}(p)/dt > 0$ at $t = T$), as a cumulative result of independent stochastic events [Samuelson, 1965] or perhaps of increasingly costly information about the future [Grossman and Stiglitz, 1980]. In the context of those models, no logic exists for expecting prices to change solely with time; hence, firms behave as if $dE(p)/dt = 0$. By the same logic, the marginal and total salvage value of stocks would fall with time, too. Note that a positive partial derivative on conditional variance with respect to time is not inconsistent with new information that make future prices more or less uncertain than today's prices. The importance of those behavioral assumptions cannot be overemphasized: *without risk averseness and the prospect that more distant events are necessarily more uncertain, all else constant, problems with inventories are underdefined mathematically and must rely on arbitrary and restrictive conditions to yield even partial solutions*.

Taking the total differential of equation (A1) and using the terminal conditions for marginal utility equation ((A2) and equation (A3)) to eliminate terms with $dq(T)$ and $di(T)$ yields

$$dT^* = -\frac{U_{\Omega}(T) + i(T)\varphi_{N\Omega}(T) + \varphi_{\Omega}(T)}{\dot{U}(T) + i(T)\dot{\varphi}_N(T) + \dot{\varphi}(T)}\,d\Omega \tag{A4}$$

where the denominator (the condition for finite *T*) is negative

That result supports several common-sense notions about how planning horizons change. It indicates that T^* (perhaps as represented by the open interest of out-month futures contracts) should increase with future prices and decrease with future uncertainty. The second step derives a partial solution for the change in optimal rate of production dq^* from the total differential of equation (2)

$$dq^* = -\frac{U_{q\Omega}}{U_{qq}}\,d\Omega - \frac{U_{qi}}{U_{qq}}\,di \tag{A5}$$

That is, in contrast to the basic representation of the standard model of inventory demand, current production will be affected by uncertainty (current and future). The third and final step derives a partial solution for change in the optimal rate of inventory accumulation di^*. The starting point for this step is the Euler equation, combining equation (3) and the particular solution for λ from equations (5) and (6), or

$$U_i e^{-rt} + \lambda = U_i e^{-rt} + \int_{s=t}^{T} U_N(s) e^{-rs} ds + \varphi_N(T) e^{-rT} = 0 \tag{A6}$$

Taking the total differential of that expression, recalling the simplifying assumptions that H_{iN} and H_{NN} both equal zero (or, equivalently, $U_{Ni} = U_{NN} = 0$), yields[27]

$$\begin{aligned} di^* = &-\frac{(U_{i\Omega} + \varphi_{N\Omega} e^{-r(T-t)})}{U_{ii}} d\Omega - \frac{U_{iq}}{U_{ii}} dq \\ &- \frac{[U_N(T) - r\varphi_N(T) + \dot{\varphi}_N(T)] e^{-r(T-t)}}{U_{ii}} dT \end{aligned} \tag{A7}$$

(The significance of assuming that storage costs are linear in N shows up in equation (A7): the initial stock level N_0 cannot affect the change in inventory accumulation rates. In some cases, however, one might want to see a bigger initial stock contributing to a higher marginal cost of storage and, hence, a smaller rise in accumulation rates than might otherwise occur.)

Combining equations (A4), (A5) and (A7) yields a general solution for the optimal change in inventory accumulation in t, as

$$di^* = \frac{\begin{pmatrix} U_{i\Omega} + \varphi_{N\Omega} e^{-r(T-t)} - U_{iq} \dfrac{U_{q\Omega}}{U_{qq}} - \\ [U_N(T) - r\varphi_N(T) + \dot{\varphi}_N(T)] \dfrac{U_{\Omega}(T) + i(T)\varphi_{N\Omega}(T) + \varphi_{\Omega}(T)}{\dot{U}(T) + i(T)\dot{\varphi}_N(T) + \dot{\varphi}(T)} e^{-r(T-t)} \end{pmatrix}}{-[U_{ii} - U_{iq} \dfrac{U_{qi}}{U_{qq}}]} d\Omega \tag{A8}$$

NOTES

1. The paper has benefited from comments by Moheb Ghali, Brad Humphreys, Robert Pindyck, James Tobin, anonymous referees, and colleagues at CBO. Special thanks to Kajal Lahiri and Fushang Liu for sharing with me the results of their research on inflation uncertainty and to Dean Croushore for his assistance on using data from the Survey of Professional Forecasters.
   The views expressed in this paper are those of the author and do not represent any policy positions of the Congressional Budget Office.
2. Emphasis on the relative importance of current and future changes may be familiar to financial analysts who study changes in the slopes of bond yield curves (for example, see Mishkin [1990]). Those studies note the importance of distinguishing whether slope changes result from changes in current or future uncertainty or whether the yield changes are uniform across maturities.
3. For example, see literature reviews in Blinder and Maccini [1991] and Humphreys, Maccini, and Schuh [2001].
4. The literature on irreversibility and capital investment offers a related explanation for a similar phenomenon (e.g., see Ingersoll and Ross [1992]). But inventory decisions, unlike capital investments, are largely reversible so long as costs associated with running low on stocks are not considered.

5. The Blinder model is in the tradition of expected-utility maximization, because the firm's selection of the optimal lot size (*S* minus *s*) is assumed to reflect the conditional probability distribution of sales. All three approaches may be classified as games against nature. Games against people have found only limited application to inventory theory—primarily to strategic stockpiling decisions in the face of a cartel (e.g., Nichols and Zeckhauser [1977]).
6. Non-expected utility models of preference, where utility does not depend linearly on probability weights, exist as well. Those models address certain observed phenomena of choice under uncertainty that appear inconsistent with expected utility maximization—as reviewed by Machina [1987]. Such research has seen little application to stock-flow problems like inventories and, with the exception of Epstein and Zin [1991], is ignored here.
7. Stock-out costs, which would depend on current stock levels (N) and either the current rates of production (e.g., for manufacturing industries needing to fill the pipeline) or the current sales (e.g., for retail industries needing to fill shelves), should rise whenever stock levels approach some minimum operating level. The *S-s* models address that cost, as does the production-smoothing model of Eichenbaum [1989]. In this study, production does not affect holding costs at all, and sales (represented as the difference between production (q) and inventory accumulation (i)) do not appear as a separate variable. The study could incorporate stock-out costs by including a term for production or sales in the holding cost function or, equivalently, a term for stocks in the production cost function.
8. The assumption of linear utility is for ease of presentation only; general expressions for first-order and partial derivatives of utility could be substituted for α and β throughout with no loss in generality, so long as the effects of changes in expectations and uncertainty are separable.
9. Explicitly constraining q-i to be nonnegative may change this problem, but not the sign of relationships between changes in expectations or uncertainty and changes in production and inventory accumulation—see Kamien and Schwartz [1981], footnote page 172. Understanding those relationships is the goal of the comparative-statics analysis that this paper presents.
10. A marginal cost representation such as $H_i = a^*i^2$ ($a > 0$) would work. The assumption of symmetry in H_i for $i > 0$ and $i < 0$ derives from the shared resources needed to build and draw stocks (pumps, forklifts, sales staff, etc.) (If that symmetry does not exist, it would be necessary to reflect builds and draws as separate variables with different costs.) The assumption of separability in N and i reflects a view that those resources differ from the requirements for holding stocks (tanks, distribution lines, shelves, etc.).
11. The focus of this paper is on changes in expectations and uncertainty for prices. However, the analysis of time-varying changes for marginal production costs (dC_q) or holding costs (dH_N or dH_i) is possible. Similarly, the recently identified relationship between internal costs of financing and inventory investment [Carpenter et al., 1994] could be analyzed on the basis of changes in the discount rate (dr). (Most simply, defining p as constant net revenues per unit of sales would enable the model to encompass changes relevant to any constant unit costs, $C_q = c$. An increase in $E(c)$ would diminish the effect of an increase in expected prices, and an increase in VAR(c) would enforce the effect of increased price uncertainty.)
12. The strength of that information set, per Fama [1970], is not relevant here.
13. If irreversibilities are present and the actual and expected paths of shocks diverge, the optimal path of inventories over time may diverge from the planning solution derived here—although solutions for the current period t=0 should be the same.
14. Those steps may be contrasted with the approach for solving the basic production-smoothing model in discrete time. From Eichenbaum [1989], for example, sales are exogenous (and arbitrary) and production is represented as the residual of expected sales and the optimal stock change. Hence, the solution can proceed from the first order condition for the single control variable, the optimal stock level in the next period.
15. A rigorous representation would acknowledge the multiple dimensions of the information set, identifying $d\Omega$ as a vector of information relevant to current and future expectations and uncertainty. The equations defining the problem should appear as summations of five matrix rows, with each row collecting all terms with constants for one of the four independent values.
16. Studies that follow the statistical approach include work by Engle [1983] and others, who use an ARCH structure to identify conditional means based on extrapolations from past inflation and unemployment and conditional variances based on deviations between extrapolated and actual values. Early studies that investigate various measures of dispersion of independent forecasts include Evans and Wachtel [1993], Huizinga [1983], and Zarnowitz and Lambros [1987].

17. Diminishing variability with the forecast horizon may reflect an emphasis of many forecasters on predicting the growth of potential GDP. Forecaster views about long-term trends may be less sensitive to new information than are views about the current state of the economy. In any case, that diminishing variability conflicts with a key assumption underlying the mathematical solution of the model: that uncertainty increases with the forecast horizon, all else constant.
18. See Giordani and Soderlind [2003] for a similar analysis, also using SPF data.
19. Serial correlation in the forecast-disagreement measure of uncertainty is very low, which suggested the possibility that lagged changes in the Lahiri-Liu index could proxy for changes in current uncertainty. However, trials using past changes (lagged one quarter) did not improve the overall regression results, and the significance of the term for uncertainty weakened. Alternative regressions that introduced current-period uncertainties based on differences between actual prices and naive prices forecasts—specifically, GARCH forecasts based on past prices and past unemployment—were not very different from the results presented here.
20. The normal approximation to the binomial distribution is $z = \frac{|52 - 0.5 * 78|}{\sqrt{78 * 0.5 * (1 - 0.5)}}$, where 0.5 is the probability that the two series move in the same direction (or in opposite directions) if they are not related, 52 is the number of matched signs, and 78 is the total number of observations of the change in *di*. For a discussion, see Larsen and Marx [1986] or other standard texts.
21. In particular, see comments on Meyer [1987] by Levy [1989] (extending Meyer's work) and Sinn [1989] (listing earlier studies on the compatibility of expected utility and mean-variance utility). Empirical evidence appears in Meyer and Rasche [1992].
22. Besides making those models easier to solve, the assumption of risk neutrality enables the analyst to avoid the unpleasant implications of negative marginal utility and satiation at some level of expected returns or wealth.
23. Epstein and Zin [1991], noting the importance of separating risk averseness from other factors affecting consumption-investment substitution, look at the optimization of time-varying (recursive), non-expected utility functions as one solution to the problem. Their approach, however, assumes all uncertainty is in the future, which limits its usefulness for addressing shocks that also affect current uncertainty.
24. More commonly, those models include an accelerator relationship that describes the adjustment of inventories towards some target level. But the effect is still to dictate the ending stock level.
25. In contrast to the applied literature, the theoretical literature on the stochastic methods utilized in the modeling of inventory demand does address the derivation of optimal stopping rules with free terminal time. For example, see Malliaris and Brock [1982].
26. One well-known source of declining utility with time is the increase in production costs with resource depletion. Additional circumstances for nondepletable resource industries may arise with the accumulating effects of environmental pollution and regulatory mandates to clean up or with expected changes in the regulations themselves (for example, see Farmer [1997]).
27. Extending this problem to include transaction costs (*CT*) that arise from stock outages poses no complications if *CT* depends on N/q, since CT_{Ni} would still equal zero. For *CT* as a function of $N/(q\text{-}i)$, CT_{Ni} is not equal to zero, and an analytical solution may not be possible.

REFERENCES

Blinder, A. Retail Inventory Behavior and Business Fluctuations. *Brookings Papers on Economic Activity*, 1981, 443-505.

____________. Can the Production Smoothing Model of Inventory Behavior Be Saved? *Quarterly Journal of Economics*, August 1986, 431-453.

Blinder, A. and Maccini, L. Taking Stock: A Critical Assessment of Recent Research on Inventories. *Journal of Economic Perspectives*, Winter 1991, 73-96.

Carpenter, R. Fazzari, S. and Petersen, B. Inventory Investment, Internal-Finance Fluctuations, and the Business Cycle. *Brookings Papers on Economic Activity*, 1994, 75-138.

Durlauf, S. and Maccini, L. Measuring Noise in Inventory Models. *Journal of Monetary Economics*, August 1995, 65-89.

Eichenbaum, M. Some Empirical Evidence on the Production Level and Production Cost Smoothing Models of Inventory Investment. *American Economic Review*, September 1989, 853-864.

Engle, R. Estimates of the Variance of U.S. Inflation Based upon the ARCH Model. *Journal of Money*, Credit, and Banking, August 1983, 286-301.

Epstein, L. and Zin, S. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns. *Journal of Political Economy*, April 1991, 263-286.

Evans, M. and Wachtel, P. Inflation Regimes and the Sources of Inflation Uncertainty. *Journal of Money, Credit, and Banking*, August 1993 (Part 2), 475-520.

Fama, E. Efficient Capital Markets: A Review of Theory and Empirical Work. *Journal of Finance,* May 1970, 383-417.

Farmer, R. Intertemporal Effects of Environmental Mandates. *Environmental and Resource Economics*, April 1997, 365-381.

Federal Reserve Bank of Philadelphia, Survey of Professional Forecasters. Documentation and survey results are available at http:/www.phil.frb.org/files/spf.

Feldstein, M. Mean-Variance Analysis in the Theory of Liquidity Preference and Portfolio Selection. *Review of Economic Studies*, January 1969, 5-12.

Giordani, P. and Soderlind, P. Inflation Forecast Uncertainty. *European Economic Review*, December 2003, 1037-1059.

Grossman, S. and Stiglitz, J. On the Impossibility of Informationally Efficient Markets. *American Economic Review*, June 1980, 393-408.

Hirshleifer, J. Investment Decision Under Uncertainty: Choice-Theoretic Approaches. *Quarterly Journal of Economics*, November 1965, 509-536.

____________. Investment Decision Under Uncertainty: Applications of the State-Preference Approach. *Quarterly Journal of Economics*, May 1966, 252-277.

Huizinga, J. Inflation Uncertainty, Relative Price Uncertainty, and Investment in U.S. Manufacturing. *Journal of Money, Credit, and Banking*, August 1983, 521-549.

Humphreys, B., Maccini, L. and Schuh, S. Input and Output Inventories. *Journal of Monetary Economics*, April 2001, 347-375.

____________. Input and Output Inventories: Erratum. *Journal of Monetary Economics*, March 2002, 455.

Ingersoll Jr., J. and Ross, S. Waiting to Invest: Investment and Uncertainty. *Journal of Business*, January 1992, 1-29.

Kahn, J. Inventories and the Volatility of Production. *American Economic Review*, September 1987, 667-679.

Kamien, M. and Schwartz, N. *Dynamic Optimization: The Calculus of Variations and Optimal Control in Economics and Management*. New York: North-Holland Publishing, 1981.

Lahiri, K. and Liu, F. Measuring Inflation Forecast Uncertainty with a Panel on Probability Forecasts. Working paper, Department of Economics, University at Albany SUNY, October 2003.

Larsen, R. and Marx, M. *An Introduction to Mathematical Statistics and Its Applications*. Englewood Cliffs: Prentice-Hall, 1986.

Levy, H. Two-Moment Decision Models and Expected Utility Maximization: Comment. *American Economic Review*, June 1989, 597-600.

Levy, H. and Markowitz, H. Approximating Expected Utility by a Function of Mean and Variance. *American Economic Review*, June 1979, 308-317.

Machina, M. Choice Under Uncertainty: Problems Solved and Unsolved. *Journal of Economic Perspectives*, Summer 1987, 121-154.

Malliaris, A. and Brock, W. *Stochastic Methods in Economics and Finance*. New York: North-Holland Publishing, 1982.

Markowitz, H. *Portfolio Selection: Efficient Diversification of Investments*. New York: John Wiley & Sons, 1959 (reprinted, New Haven: Yale University Press, 1970).

McKinnon, R. Futures Markets, Buffer Stocks, and Income Stability for Primary Producers. *Journal of Political Economy*, December 1967, 844-861.

Meyer, J. Two-Moment Decision Models and Expected Utility Maximization. *American Economic Review*, June 1987, 421-430.

Meyer, J. and Rasche, R. Sufficient Conditions for Expected Utility to Imply Mean-Standard Deviation Rankings: Empirical Evidence Concerning the Location and Scale Condition. *Economic Journal*, January 1992, 91-106.

Meyer, J. Two-Moment Decision Models and Expected Utility Maximization: Reply, *American Economic Review,* June 1989, 603.

Mishkin, F. What Does the Term Structure Tell Us About Future Inflation? *Journal of Monetary Economics*, January 1990, 77-95.

Morck, R., Schleifer, A. and Vishny, R. Management Ownership and Market Valuation: an Empirical Analysis. *Journal of Financial Economics*, 1988, 293-315.

von Neumann, J. and Morgenstern, O. *Theory of Games and Economic Behavior*. Princeton: Princeton University Press. 1944.

Nichols, A. and Zeckhauser, R. Stockpiling Strategies and Cartel Prices, *The Bell Journal of Economics*, Spring 1977, 66-96.

Pindyck, R. Inventories and the Short-Run Dynamics of Commodity Prices, *Rand Journal of Economics*, Spring 1994, 141-159.

Pontryagin, L., Boltyanskii, V., Gamkrelidze, R., and Mishchenko, E. *The Mathematical Theory of Optimal Processes*. New York: John Wiley & Sons, 1962.

Ramey, V. and West, K. *Handbook of Macroeconomics*, editied by J. Taylor and M. Woodford. New York: North-Holland Publishing, 1998.

Rubinstein, M. A Mean-Variance Synthesis of Corporate Financial Theory. *Journal of Finance*, March 1973, 167-181.

Samuelson, P. Proof that Properly Anticipated Prices Fluctuate Randomly. *Industrial Management Review*, Spring 1965, 41-49.

____________. Lifetime Portfolio Selection by Dynamic Stochastic Programming. *Review of Economics and Statistics*, August 1969, 239-246.

Sinn, H. *Economic Decisions Under Uncertainty*. NewYork: North-Holland, 1983.

____________. Two-Moment Decision Models and Expected Utility Maximization: Comment. *American Economic Review*, June 1989, 601-602.

Tobin, J. Liquidity Preference as Behavior Towards Risk. *Review of Economic Studies,* February 1958, 65-86.

Turnovsky, S. The Determination of Spot and Futures Prices with Storable Commodities. *Econometrica*, September 1983, 1363-1387.

Zangwill, W. A Deterministic Multiproduct, Multifacility Production and Inventory Model. *Operation Research,* May-June 1966, 486-507.

Zarnowitz V. and Lambros, L. Consensus and Uncertainty in Economic Prediction. *Journal of Political Economy*, June 1987, 591-621.