

Equilibrium Forward Curves for Commodities

BRYAN R. ROUTLEDGE, DUANE J. SEPPI,
and CHESTER S. SPATT*

ABSTRACT

We develop an equilibrium model of the term structure of forward prices for storable commodities. As a consequence of a nonnegativity constraint on inventory, the spot commodity has an embedded timing option that is absent in forward contracts. This option's value changes over time due to both endogenous inventory and exogenous transitory shocks to supply and demand. Our model makes predictions about volatilities of forward prices at different horizons and shows how conditional violations of the "Samuelson effect" occur. We extend the model to incorporate a permanent second factor and calibrate the model to crude oil futures data.

COMMODITY MARKETS IN RECENT YEARS have experienced dramatic growth in trading volume, the variety of contracts, and the range of underlying commodities. Market participants are also increasingly sophisticated about recognizing and exercising operational contingencies embedded in delivery contracts.¹ For all of these reasons, there is a widespread interest in models for pricing and hedging commodity-linked contingent claims. In this paper we present an equilibrium model of commodity spot and forward prices. By explicitly incorporating the microeconomics of supply, demand, and storage, our model captures some fundamental differences between commodities and financial assets.

Empirically, commodities are strikingly different from stocks, bonds and other conventional financial assets. Among these differences are:

* Graduate School of Industrial Administration, Carnegie Mellon University. We thank René Stulz and an anonymous referee for their helpful advice and David Backus, Jonathan Berk, Frank de Jong, Eric Ghysels, Rick Green, Rob Heinkel, Glen Kentwell, John Lehoczky, Bart Lipman, Kristian Miltersen, R. (Nic) Nicolaides, Christine Parlour, Pedro Santa Clara, Eduardo Schwartz, Steven Shreve, Sridhar Tayur, and Harold Zhang for comments. We also thank seminar participants at Carnegie Mellon, Northwestern, Penn State, Toulouse, UC Berkeley, UCLA, the University of Iowa, the University of Texas at Austin, the University of Vienna, the Vienna Institute for Advanced Studies, and Yale as well as at the Seventh Annual Derivative Securities Conference at Queen's University, the 1997 NBER Summer Asset Pricing Workshop, the 1997 European Finance Association, the 1998 American Finance Association, and the 1998 Q-Group meetings. The Decision, Risk, and Management Science Program of the National Science Foundation and the Institute for Quantitative Research in Finance provided financial support for this research. The Q-Group also awarded the paper its 1998 Roger Murray Prize.

¹ For example, energy supply contracts often include so-called "swing" options which give industrial consumers the flexibility to increase their "take" above a baseload at a fixed price for a pre-agreed number of extra days each month. See Jaillet, Ronn, and Tompaidis (1997) and Pilipovic and Wengler (1998).

- Commodity futures prices are often “backwardated” in that they decline with time-to-delivery. For example, Litzenger and Rabinowitz (1995) document that nine-month futures prices are below the one-month prices 77 percent of the time for crude oil.
- Spot and futures prices are mean reverting for many commodities.
- Commodity prices are strongly heteroskedastic (see Duffie and Gray (1995)) and price volatility is positively correlated with the degree of backwardation (see Ng and Pirrong (1994) and Litzenger and Rabinowitz (1995)).
- The term structure of commodity forward price volatility typically declines with contract horizon. This is known as the “Samuelson (1965) effect.” However, violations of this pattern occur when inventory is high (see Fama and French (1988)).
- Unlike financial assets, many commodities have pronounced seasonalities in both price levels and volatilities.

In equilibrium, backwardation implies that immediate ownership of the physical commodity entails some benefit or convenience which deferred ownership (via a long forward position) does not. This benefit, expressed as a rate, is termed the “convenience yield” (e.g., see Hull (1997)). A convenience yield is natural for goods, like art or land, that offer exogenous rental or service flows over time. However, substantial convenience yields are also observed in other commodities, such as agricultural products, industrial metals, and energy, which are consumed at a single point in time.

The “theory of storage” of Kaldor (1939), Working (1948, 1949), and Telser (1958) explains convenience yields in terms of an embedded timing option. In particular, the holder of a storable commodity (e.g., oil, natural gas, copper) can decide when to consume it. If it is optimal to store a commodity for future consumption, then it is priced like an asset, but if it is optimal to consume it immediately, then the commodity is priced as a *consumption good*. Thus, a commodity’s spot price is the *maximum* of its current consumption and asset values. In contrast, forward prices derive solely from the asset value of the *deferred* right to consume after delivery. Inventory decisions are important for commodities because—by influencing the relative current and future scarcity of the good—they link its current (consumption) and expected future (asset) values. This is unlike equities and bonds where outstanding quantities are fixed. This link is imperfect, however, because inventory is physically constrained to be *nonnegative*. Inventory can always be *added* to keep current spot prices from being too low relative to expected future spot prices. However, once the aggregate *discretionary* inventory² of a

² Measured inventory is never literally driven to zero in practice since some stocks are held as committed inputs in production (e.g., gas or oil in transit). By discretionary inventory we mean commodity stocks in excess of those committed to the production process (e.g., exchange warehouse holdings or inventory held otherwise by traders). For example, Brennan (1991) documents high copper prices when the aggregate inventory/sales ratio falls below a few weeks. Nondiscretionary inventories have their own convenience value by reducing production disruptions, minimizing delivery costs, etc.

commodity is driven to zero, its spot price is tied solely to the good's (high) "immediate use" consumption value. Thus, "stockouts" break the link between the current consumption and expected future asset values of a good. The result is backwardation and positive convenience yields.

In this paper we follow Deaton and Laroque (1992, 1996), Williams and Wright (1991), and Chambers and Bailey (1996) and use a competitive rational expectations model of storage to study the impact of the embedded timing option on commodity spot and futures/forward pricing.³ We assume that the "immediate-use" consumption value is driven by a mean-reverting Markov process and solve for the equilibrium inventory of competitive, risk-neutral agents. The shock process and inventory rule then jointly determine the spot and forward price processes. Our main results are the following:

- The equilibrium term structure of spot and forward prices is decreasing in inventory and—under a natural sufficient condition—increasing in the current Markov shock.
- Endogenous binomial price trees are constructed for pricing and hedging commodity-linked futures and (by extension) options and other derivatives.
- Conditional violations of the "Samuelson effect" occur when inventory is sufficiently high in the model. In particular, forward price volatilities can initially *increase* with contract horizon.
- Hedge ratios for long-dated forward positions using short-dated forwards are not constant, but are conditional on the current demand shock and the endogenous inventory level.
- A one-factor version of the model cannot match both the high unconditional volatility of long-horizon Nymex crude oil futures prices and the conditional volatilities given contango and backwardation. However, a tractable two-factor augmentation of the basic model is more successful.

An alternative to modeling forward and spot prices explicitly from economic primitives is to treat the convenience yield as an exogenous "dividend" process. For example, Brennan (1991), Gibson and Schwartz (1990), Amin, Ng, and Pirrong (1995), and Schwartz (1997) all model spot prices and convenience yields as separate stochastic processes with a constant correlation.⁴ Although these models are powerful tools for derivative pricing and hedging,

³ Routledge, Seppi, and Spatt (1999) extend this analysis to multiple commodities and, specifically, to electricity. While electricity itself is not directly storable, potential marginal fuels (e.g., natural gas, coal) are storable. Heinkel, Howe, and Hughes (1990) analyze a three-period economy in which optionality induces convenience yields and Bresnahan and Spiller (1986) discuss the relationship between inventory and the slope of the forward curve. Litzenberger and Rabinowitz (1995) model the impact of timing options on the optimal extraction path for a depletable resource and equilibrium prices. Pirrong (1998) studies commodity option pricing with storage.

⁴ Schwartz and Smith (2000) present a model with two factors, a "long-run" price component and a transitory disturbance, which they show is equivalent to the Gibson and Schwartz (1990) convenience yield model. Miltersen and Schwartz (1998) model a stochastic term structure of convenience yields based on Heath, Jarrow, and Morton (1992).

our approach also has several attractions. First, explicitly modeling the joint evolution of inventory and spot prices ensures the consistency of the spot price and convenience yield (i.e., forward price) dynamics.⁵ Second, our model predicts that the correlation between spot prices and convenience yields (or their innovations) is unlikely to be constant due to its dependence on inventory. Third, inventory acts as a second state variable summarizing past shocks, which allows our model to capture a stochastic convenience yield with only one exogenous factor.

The paper is organized as follows. Section I describes the basic one-factor model, demonstrates existence of equilibrium, and derives properties of the equilibrium inventory and spot price processes. Section II investigates properties of the forward curve and endogenous implied convenience yields. We also study the forward price volatility term structure and hedge ratios. Section III carries out a numerical calibration exercise and presents a tractable two-factor version of the basic model. Section IV concludes. Proofs are collected in the Appendix.

I. Equilibrium Model of Commodity Prices

The goal of this paper is to characterize spot and forward commodity prices in an equilibrium model of inventory with nonnegative storage. Our analysis builds on and extends Deaton and Laroque (1992, 1996), Chambers and Bailey (1996), Williams and Wright (1991), and Wright and Williams (1989). We start with general functional forms and develop tractable numerical implementations for derivative security valuation. We also investigate properties of the term structure of forward price volatility.

A. Model Structure

Consider a discrete-time, infinite horizon model in which a single homogeneous commodity is traded in a competitive market at dates $t = 1, 2, \dots$. Current production and consumption demand for “immediate use” are modeled as stochastic, reduced-form functions, g_t and c_t , of the spot price, P_t . The commodity can be stored by a group of competitive risk-neutral inventory traders who have access to a costly storage technology.

Storage is costly due to a constant proportional depreciation or wastage factor, $\delta \in (0, 1]$. Storage of q units of the commodity at $t - 1$ yields $(1 - \delta)q$ at t . One can interpret δ as spoilage (for agricultural commodities) or as a volumetric cost (for metals and energy). For technical reasons δ is strictly positive, but we abstract, for simplicity, from any other fixed or marginal costs.

⁵ Analogously, in some term structure settings assuming exogenous dynamics for interest rates of bonds of different maturities can be incompatible with equilibrium (i.e., admit arbitrage opportunities). Examples along these lines are noted in Cox, Ingersoll, and Ross (1985). Heath et al. (1992) describe the restrictions on forward price processes that are implied by the absence of arbitrage. Arbitrary joint dynamics for stock prices and dividends may also be inconsistent with equilibrium.

The spot price, P_t , is determined by market clearing. At each date t the two sources of the commodity, current production and incoming inventory $g_t + (1 - \delta)Q_{t-1}$, must equal the two types of demand, immediate consumption and outgoing inventory $c_t + Q_t$. This can be rearranged to get

$$c_t(P_t) - g_t(P_t) = -\Delta Q_t, \tag{1}$$

where $\Delta Q_t = Q_t - (1 - \delta)Q_{t-1}$. When ΔQ_t is positive (i.e., when inventory is increased), less of the good is left for immediate consumption. If the “immediate use” net-demand $c_t(P_t) - g_t(P_t)$ is monotone decreasing in the spot price P_t , the spot market can be summarized with an inverse net demand function

$$P_t = f(a_t, \Delta Q_t). \tag{2}$$

We initially abstract from permanent shocks and model the net demand shocks $a_t \in \Omega$ as realizations of a finite-dimensional, irreducible, m -state Markov process ($m \geq 2$), with transition probabilities $\pi(a|a_t)$ in a matrix Π . The shocks, a_t , represent the transitory effects of weather and/or production disruptions on the “immediate use” net demand, $c_t - g_t$. They are unaffected by current and/or past inventory. Where noted, we sometimes make an additional assumption that $\pi(a|a_t) > 0$ for all $a_t, a \in \Omega$ (denoted $\Pi \gg 0$), which implies that the demand shocks have a limiting distribution that is independent of the current state.⁶

Inventory (or the change in inventory ΔQ) is the key endogenous variable. We assume that the net demand function, f , has the following properties for all realizations $a \in \Omega$:

- (A) $f(a, \Delta Q) > 0$
- (B) $f(a, \Delta Q)$ is increasing in ΔQ
- (C) $f(a, \Delta Q)$ is continuous and unbounded from above in ΔQ .

The first assumption states that prices are positive. The second says that increased storage raises the good’s marginal valuation since it reduces the amount available for immediate use. The third assumption, that f is continuous and unbounded, is sufficient to ensure that a market-clearing spot price exists.⁷ The marginal valuation, f , need not be linear. Nonlinearities in f can arise from either the demand, c_t , or supply, g_t .

⁶ An irreducible Markov process has the property that Ω is the unique ergodic set. This means that all states are reachable with positive probability in some finite number of transitions from any initial state. For example, there are no absorbing states. The additional assumption that Π is strictly positive implies that there are no cyclically moving subsets. Alternatively, everywhere we assume $\Pi \gg 0$, we could instead assume some finite n exists such that $\Pi^n \gg 0$ (see Karlin and Taylor (1975)).

⁷ Our analysis does not require that ΔQ is unbounded, but rather only that $f(a, \Delta Q)$ is unbounded. The special case in which the maximum feasible current production output is bounded by some quantity \bar{g} (e.g., as with a harvest) has $f(a, \Delta Q) \rightarrow \infty$ as $\Delta Q \rightarrow \bar{g}$.

The one-period risk-free interest rate is assumed here to be constant, $r \geq 0$. Nonstochastic interest rates simplify the analysis and let us abstract from the difference between forward and futures prices.

B. Equilibrium

Inventory decisions of the risk-neutral traders are easily characterized since the only motive for holding inventory is trading profit.⁸ If spot prices were expected to rise by more than the “carrying costs” (i.e., wastage and interest), additional inventory would be purchased. This would then increase current (and lower future) spot prices. Conversely, if prices were expected to fall (or rise by less than the carrying costs), then inventory would be sold. However, inventory—or, more precisely, traders’ *discretionary* holdings in excess of stocks irreversibly committed to production—can only be reduced to zero. Optimal trading, given this physical nonnegativity constraint, implies that equilibrium spot prices and aggregate inventory must jointly satisfy

$$P_t = \theta E_t[P_{t+1}] \quad \text{if } Q_t > 0 \quad (3a)$$

$$P_t \geq \theta E_t[P_{t+1}] \quad \text{if } Q_t = 0, \quad (3b)$$

where $\theta = (1 - \delta)/(1 + r) < 1$ and $E_t[\cdot]$ denotes (rational) expectations conditional on the information, a_t and Q_{t-1} , available at time t . Positive equilibrium inventory, $Q_t > 0$, implies traders are indifferent to marginal changes in their inventory. However, if inventory is “stocked out,” $Q_t = 0$, traders may *want* to sell additional units but are physically constrained from doing so. One cannot consume goods that do not yet exist. It is precisely this asymmetry that makes the embedded timing option in the spot commodity valuable and that, in equilibrium, leads to option-like behavior in commodity prices.

We use $P(a, q)$ and $J(a, q)$ to denote the equilibrium spot price and aggregate inventory functions given a current demand state, a , and *previous* aggregate inventory, q . Without loss of generality, we assume $Q_0 = 0$.

Definition of Equilibrium: $\{Q_t, P_t\}$ is a *stationary rational expectations equilibrium* inventory and price process if inventory and spot-price functions $Q_0 = 0$, $Q_t = J(a_t, Q_{t-1})$ and $P_t = P(a_t, Q_{t-1}) = f(a_t, Q_t - (1 - \delta)Q_{t-1})$ exist that satisfy condition (3) for all t .

⁸ Jagannathan (1985) and Richard and Sundaresan (1981) derive equilibrium commodity prices under risk aversion. These papers do not consider the short-sale constraint while our model abstracts from risk aversion. With risk-averse traders the expectations in condition (3) should be taken using the risk-neutral martingale measure.

Proposition 1 establishes existence. The corollaries that follow identify some useful properties of the equilibrium inventory process. These are standard (e.g., see Deaton and Laroque (1992)) and are not the focus of our paper.

PROPOSITION 1 (Equilibrium): *A stationary rational expectations equilibrium exists and has the following properties:*

- (a) *the equilibrium inventory $J(a, q)$ is continuous in q and, for all $a \in \Omega$, $q \geq 0$, and perturbations $\epsilon > 0$, satisfies*

$$0 \leq J(a, q + \epsilon) - J(a, q) < (1 - \delta)\epsilon \tag{4}$$

- (b) *a unique finite upper bound, $Q_{max} \geq 0$, exists such that $J(a, q) \leq Q_{max}$ for all $a \in \Omega$ and $q \in [0, Q_{max}]$ and $J(a, Q_{max}) = Q_{max}$ for some $a \in \Omega$, and*
- (c) *the equilibrium spot price $P(a, q)$ is continuous and decreasing in q and is bounded such that $0 < P(a, q) < \infty$ for all $a \in \Omega$ and $q \in [0, Q_{max}]$.*

Costly storage and the assumptions on f ensure that the slope of the inventory function J is less than $1 - \delta$ and, hence, that equilibrium inventory is bounded. Another implication of the slope of J being less than $1 - \delta$ is that equilibrium spot prices P_t are smaller when previous storage Q_{t-1} is higher. Prices are positive due to the proportional storage cost and finite inventory (and inventory changes) are bounded. If carrying costs, θ , are large enough, equilibrium inventory is always zero (i.e., $Q_t = Q_{max} = 0$). However, equilibrium storage can be positive if $f(a_j, 0) < \theta E[f(a, 0)|a_j]$ in some demand states $a_j \in \Omega$. In equilibria with positive storage, certain $a \in \Omega$ can be identified as “sell states” in which inventory is never accumulated. In particular, if the prior inventory is sufficiently low, then all incoming inventory is sold for consumption and a “stockout” occurs.

COROLLARY 1.1 (Properties of J): *In a rational expectations equilibrium with $Q_{max} > 0$, there exists a nonempty set Ω_s of “sell states” such that for all $a_s \in \Omega_s$*

- (a) *inventory is reduced so that $J(a_s, q) \leq (1 - \delta)q$ for all q and*
- (b) *a critical inventory level $q_s \geq 0$ exists for each $a_s \in \Omega_s$ such that a stockout, $J(a_s, q) = 0$, occurs whenever previous inventory is sufficiently low, $q < q_s$.*

By construction, the equilibrium inventory and demand shock processes, (a_t, Q_t) , are jointly Markovian. However, the sequence of demand shocks, $\{a_\tau\}_{\tau=0}^t$, affects the inventory level, Q_t , because of the kink in the equilibrium condition (3). For example, a “sell state” followed by a “buy state” (in which inventory is accumulated) does not necessarily lead to the same ending inventory as vice versa. However, any time inventory “stocks out” (i.e., is

driven to zero) the inventory process *regenerates* or renews.⁹ This leads to the useful property that the distribution of long-run inventory (and, hence, of long-run prices) does not depend on current inventory.

COROLLARY 1.2 (Regeneration): *In the rational expectations equilibrium,*

- (a) *stockouts occur with positive probability; in other words, the probability of perpetually positive inventory, $\text{Prob}(Q_t > 0, Q_{t+1} > 0, Q_{t+2} > 0, \dots | a_t, Q_{t-1})$, is zero for any a_t and Q_{t-1} ,*
- (b) *the long-run probability distributions of future inventory and prices, Q_T and P_T , are independent of current inventory, Q_t , as $T \rightarrow \infty$*
- (c) *if the transition probabilities are all strictly positive, $\Pi \gg 0$, then*
 - (i) *limiting inventory and price distributions exist and do not depend on a_t or Q_{t-1} in that $\lim_{T \rightarrow \infty} \text{Prob}(Q_T \leq q | a_t, Q_{t-1}) = \phi_Q(q)$ and $\lim_{T \rightarrow \infty} \text{Prob}(P_T \leq p | a_t, Q_{t-1}) = \phi_P(p)$, and*
 - (ii) *the limiting probability of a stockout is strictly positive, $\phi_Q(0) \in (0, 1]$.*

Inventory cannot *always* be positive with positive carrying costs. Hence, stockouts occur and current inventory Q_t only has a temporary effect on future ($T > t$) inventory, Q_T , and prices, P_T . If, moreover, the Markov shock process has an invariant limiting distribution, then the current net demand state a_t also has only a temporary influence.

C. Numerical Example

To implement our model we must specify a Markov process for the net-demand shocks a_t , the inverse demand function f , the storage cost rate δ , and interest rate r . Our baseline parameterization is a linear two-state special case:

$$\begin{aligned}
 a_H &= 1 & a_L &= 0 \\
 \Pi &= \begin{pmatrix} \pi(a_H|a_H) & \pi(a_L|a_H) \\ \pi(a_H|a_L) & \pi(a_L|a_L) \end{pmatrix} = \begin{pmatrix} 0.75 & 0.25 \\ 0.25 & 0.75 \end{pmatrix} & (5) \\
 P_t &= f(a, \Delta Q) = a + \Delta Q,
 \end{aligned}$$

with a storage cost $\delta = 0.1$ and interest rate $r = 0$. Several features of this example deserve comment. First, seasonality is easily added by introducing cyclically moving subsets into the Markov structure.¹⁰ Second, future high (low) demand states are more likely if the current state is high (low). Some

⁹ As is common in the study of inventory systems, the inventory process can be used to define a *renewal process* over the stockout event. See Chapter 5 of Karlin and Taylor (1975).

¹⁰ For example, we could partition Ω into four seasons Ω_k where $k \in \{A, W, S, M\}$ with positive transition probabilities $\pi(a|a_t) > 0$ only if $a \in \Omega_{k(t+1)}$ where $k(t+1)$ denotes the season of date $t+1$.

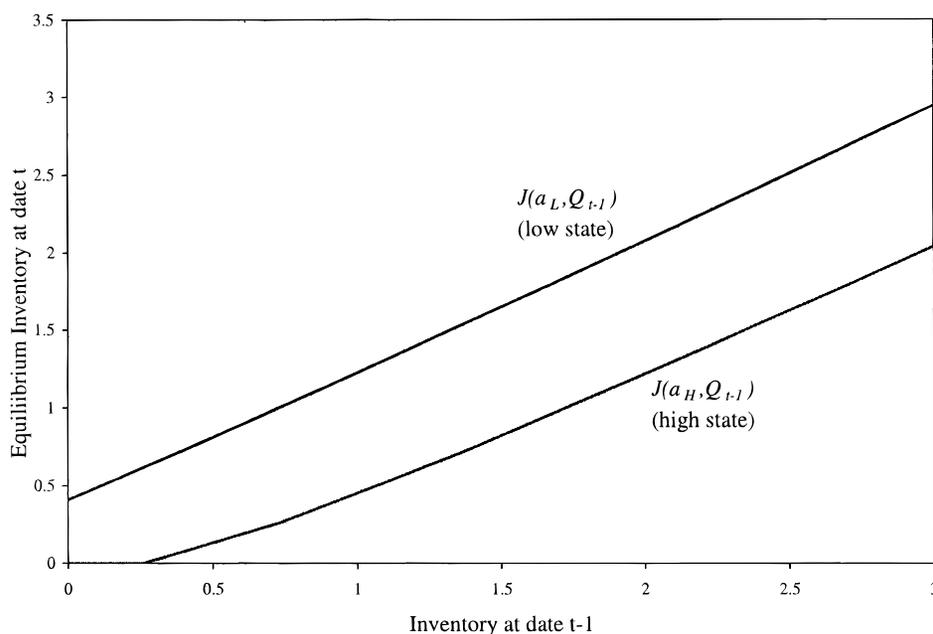


Figure 1. Numerical example of equilibrium inventory. The equilibrium outgoing inventory, Q_t , is plotted as a function of the previous inventory, Q_{t-1} , in the high and low net demand states given the baseline parameters of the numerical example: $a_H = 1$, $a_L = 0$, $\pi(a_H|a_H) = \pi(a_L|a_L) = 0.75$, $r = 0$, and $\delta = 0.1$.

of our theoretical results on the volatility of forward prices assume that the Markov process has this second property. Despite its simplicity, this two-state, linear specification is able to generate endogenous binomial trees of spot (and forward) prices with many interesting features.

Figure 1 shows the equilibrium inventory function $Q_t = J(a, Q_{t-1})$ for this numerical example.¹¹ Given the previous inventory Q_{t-1} , outgoing inventory Q_t is greater in the low state a_L than in the high state a_H . In particular, $J(a_H, Q_{t-1}) \leq (1 - \delta)Q_{t-1} \leq J(a_L, Q_{t-1})$ so that, consistent with Corollary 1.2 (with just $m = 2$ states), inventory is accumulated in the low state and consumed in the high state.

¹¹ We calculate $J(a, q)$ using a piecewise linear approximation for each $a \in \Omega$ on an equally spaced grid, G , with 1000 points. From an initial conjecture J_0 , successive approximations J_n are calculated in a fixed-point contraction mapping algorithm. For each $a \in \Omega$ and $q \in G$, $J_{n+1}(a, q)$ is the maximum of the solution to the equilibrium condition with equality (equation (3a) and zero. This is very similar to the construction in the proof of Proposition 1. The convergence rate is roughly $1 - \delta$, requiring typically 20 to 50 iterations to be within tolerance of the fixed point, $J(a, q)$. See Deaton and Laroque (1992) for additional details. In Section III we use a third-order polynomial approximation of J (in place of the discrete grid) to increase the speed of the algorithm.

II. Properties of Forward and Spot Prices

Define $F_{t,t+n}(a_t, Q_{t-1})$ as the *forward price* agreed to at date t for one unit of the commodity to be paid for and delivered at a future date $t + n$ given the information, a_t and Q_{t-1} , available at t . Since traders are risk neutral and the interest rate is nonstochastic, market clearing requires

$$F_{t,t+n}(a_t, Q_{t-1}) = E[P_{t+n} | a_t, Q_{t-1}]. \quad (6)$$

Since forward contracts only involve payment at delivery $t + n$, the carrying costs, δ and r , enter forward prices only indirectly through their effect on the inventory and spot price processes.

For concreteness, Figure 2 illustrates the range of forward-price term structures possible in our numerical example.¹² Each forward curve is for a different combination of net demand shock, a_L or a_H , and previous inventory, Q_{t-1} . The forward curves are upward sloping in the low demand state a_L and slope downward in the high demand state a_H with a low (or zero) prior inventory. When demand is high and the incoming inventory is at a moderate level, then the forward curve can be “hump shaped.” In particular, forward prices initially rise (i.e., $P_t = \theta E_t[P_{t+1}] < E_t[P_{t+1}] = F_{t,t+1}$ since outgoing inventory, Q_t , is positive), but eventually decline to a state-independent (constant), long-term forward price, F_∞ (see Proposition 3 below). One artifact of this two-state example is that the forward curves are identical in a stockout since outgoing inventory is 0 (by definition) and the state is always a_H . The stockout spot prices differ, however, because of differences in the level of incoming inventory. The result is the “fan” at $F_{t,t+1}$ in the stockout states.

A. Inventory and Forward Prices

Inventory plays a crucial role as an endogenous state variable summarizing the cumulative impact of past shocks, a_{t-1}, a_{t-2}, \dots in our model. To understand its impact on prices, consider two equilibrium inventory and price sequences that differ only by an (exogenous) perturbation to the initial

¹² Calculating exact forward prices requires a tree of all possible future spot prices. Since the inventory process is not recombining, each iteration to determine forward prices to horizon n potentially requires, given m demand states, calculations on the order of m^n . With a discrete grid of 1000 levels of inventory, this produces a large (but sparse) Markov transition matrix that can be used to calculate forward prices of any horizon as well as the limiting (unconditional) distribution. The impact of approximation errors on prices from rounding the inventory level are limited by Proposition 2. One can improve the accuracy (or reduce the number of discrete inventory levels) by choosing the inventory grid based on the equilibrium inventory function. The renewal property implies that one only needs to track paths until inventory reaches a stockout. Since stockouts happen with non-zero probability, most of the time inventory is within only a few steps (i.e., demand realizations) from a stockout. For example, Figure 2 was produced with just 100 points (carefully chosen) on the inventory grid. The 60 most likely forward curves are shown.

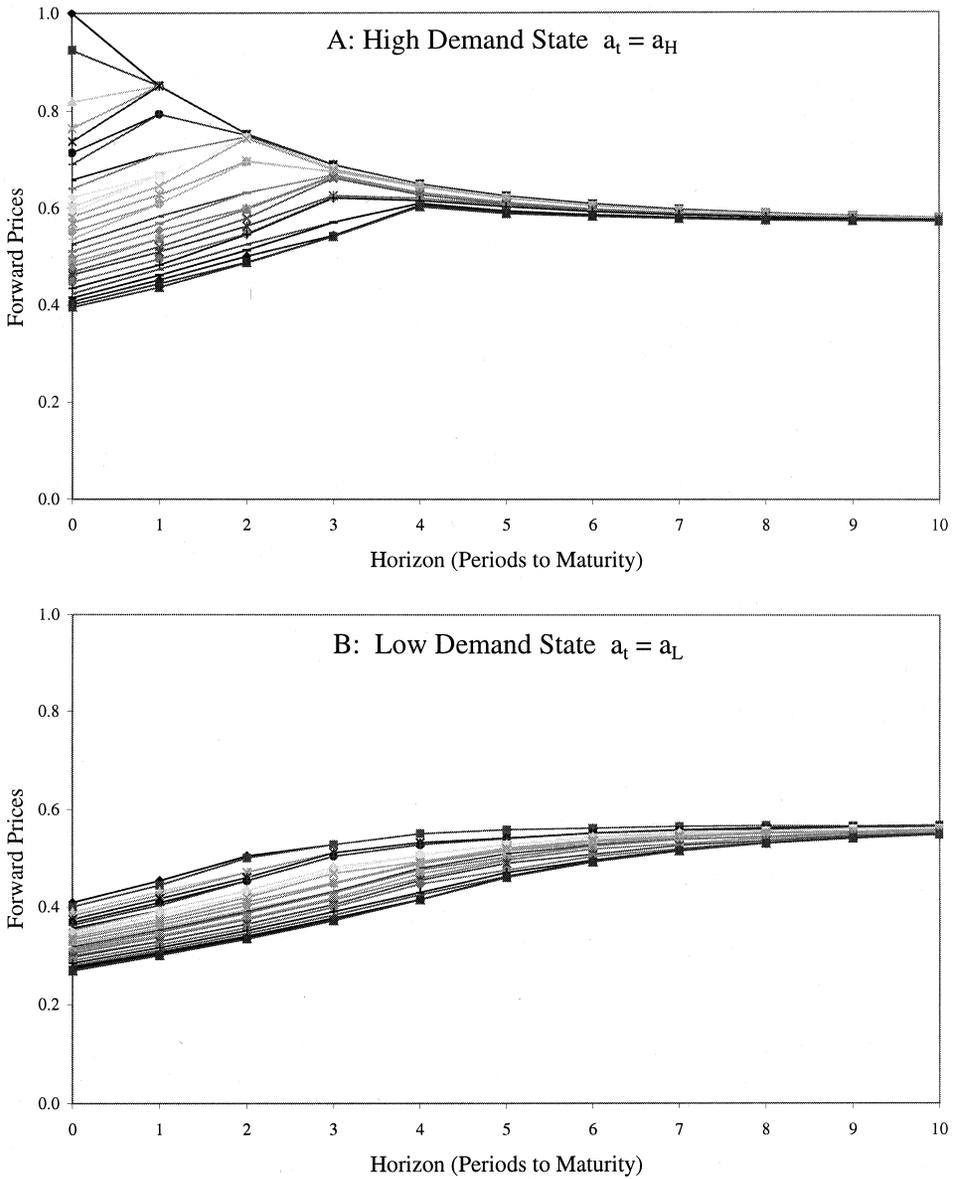


Figure 2. Forward curves. The 60 most common forward curves are depicted in the (A) high, a_H , and (B) low, a_L , demand states in the numerical example with baseline parameters: $a_H = 1$, $a_L = 0$, $\pi(a_H|a_H) = \pi(a_L|a_L) = 0.75$, $r = 0$, and $\delta = 0.1$.

inventory. In particular, the path of demand realizations, $\{a_t\}$, is identical for the two sequences. The following proposition shows that spot and forward prices are decreasing in initial inventory, but that the effect is only temporary.

PROPOSITION 2 (Increase Inventory): Consider two inventory processes $\{Q_t\}$ and a perturbed process $\{Q_t^x\}$ where $Q_t = J(a_t, Q_{t-1})$ and $Q_t^x = J(a_t, Q_{t-1}^x)$. If $Q_0 = 0$ and $Q_0^x = x > 0$, then for every realization of $\{a_t\}_{t=1}^\infty$

- (a) the perturbed inventory is (weakly) larger, but the difference shrinks with time, $0 \leq Q_t^x - Q_t < (1 - \delta)^t x$,
- (b) the perturbed spot and forward prices are weakly lower, $P_t^x \leq P_t$ and $F_{t,t+n}^x \leq F_{t,t+n}$ for all $n \geq 0$,
- (c) $\lim_{t \rightarrow \infty} (Q_t^x - Q_t) = 0$, $\lim_{t \rightarrow \infty} (P_t^x - P_t) = 0$, $\lim_{t \rightarrow \infty} (F_{t,t+n}^x - F_{t,t+n}) = 0$, and
- (d) for all $\epsilon \in (0, 1)$, there exists a date τ such that $\text{Prob}(Q_t^x \neq Q_t) < \epsilon$ for all $t > \tau$.

Since the slope of J is less than $1 - \delta$, any positive perturbation x reduces the current and future net inventory changes ΔQ_t . Thus spot and forward prices are lower. However, the perturbation is strictly temporary and dies out at a rate $(1 - \delta)^t$. This is why statements (a) to (c) hold *path-by-path* for every sequence of demand shocks. Furthermore, after the first stockout in the Q process, both inventory processes regenerate and are identical thereafter. Since inventory does not affect the probabilities of future demand shocks, changes in inventory simply shift the probability distribution of future spot prices. Although our focus is primarily on forward prices, Proposition 2 has an obvious but important implication for option prices, as stated in Corollary 2.1.

COROLLARY 2.1 (Derivative Prices): The prices of European calls (and all other derivatives whose payoffs are increasing in future spot or forward prices) are decreasing in inventory in the stationary rational expectations equilibrium.

Since current inventory has no long-term impact, the equilibrium prices preserve many of the features of the Markov structure of the underlying demand shocks. For example, if all demand shocks have strictly positive probability ($\Pi \gg 0$), then forward prices, $F_{t,t+n}$, are bounded in a range that tightens with the contract horizon, n , and a limiting forward price, F_∞ , exists that is invariant to the current demand state and inventory level.¹³

PROPOSITION 3 (Forward Prices): In the stationary rational expectations equilibrium of an economy with $\Pi \gg 0$, two monotonic sequences, $\{F_{min}^n\}_{n=0}^\infty$ (which is increasing in n) and $\{F_{max}^n\}_{n=0}^\infty$ (which is decreasing in n), bound forward prices $F_{t,t+n} \in [F_{min}^n, F_{max}^n]$ for all t and n with $\lim_{n \rightarrow \infty} F_{min}^n = \lim_{n \rightarrow \infty} F_{max}^n = F_\infty$.

The limiting forward price, F_∞ , is the unconditional mean spot price. In general, it is hard to determine the effect of storage on the mean spot price. Since storage is costly, less of the good is available for consumption, which increases its price. However, storage also smooths shocks that may, de-

¹³ With seasonal cyclical subsets in Ω , there is not a global limiting price distribution, but rather one for each season.

pending on the curvature of f , reduce the mean price. However, even if the net-demand function f is linear, the effect of moving from efficient storage ($\delta = 0$) to no-storage ($\delta = 1$) on F_∞ is not monotonic. Increasing δ , holding inventory constant, raises spot prices since more of the good is lost to the storage cost. However, equilibrium storage also changes with δ . At low values of δ (close to 0), an increase in storage costs raises the mean spot price since the mean equilibrium inventory levels are large. However, for high values of δ , and, hence, lower equilibrium inventory levels, the mean spot price decreases since less is lost to storage costs. In the limit as $\delta \rightarrow 1$, nothing is stored and nothing is lost to storage cost, the limiting forward price, F_∞ , is $E[f(a_t, 0)]$.

B. Net Demand Shocks and Forward Prices

The a_t realizations influence spot and forward prices directly, through current net demand, and indirectly, via the probability distribution over future shocks. Due to these two influences, two further assumptions (in addition to (A), (B), and (C) in Section I.A) are used to give an unambiguous order to the m demand states of Ω :

- (D) For all ΔQ , the inverse net demands are ordered with $f(a_1, \Delta Q) < \dots < f(a_m, \Delta Q)$.
- (E) First-Order Stochastic Dominance: for all $a_\ell, a_h \in \Omega$ where $f(a_\ell, \Delta Q) < f(a_h, \Delta Q)$,

$$\sum_{j=1}^k \pi(a_j | a_\ell) \geq \sum_{j=1}^k \pi(a_j | a_h) \tag{7}$$

for all $k = 1, \dots, m$ (with strict inequality for at least one k).

The first assumption, (D), says that the spot prices corresponding to the net demand states have the same order regardless of the change in inventory. The second, (E), states that the transition probabilities from high states, a_h , first-order stochastically dominate those from lower states, a_ℓ . In our two-state example, this implies that shocks are *persistent* in that $\pi(a_H | a_H), \pi(a_L | a_L) > 0.5$. Under these two (sufficient) assumptions forward prices have a particularly simple structure.

PROPOSITION 4 (Increase Demand): *In the stationary rational expectations equilibrium of an economy satisfying (D) and (E),*

- (a) *spot prices are ordered:* $P(a_1, q) < \dots < P(a_m, q)$,
- (b) *forward prices are ordered:* $F_{t, t+n}(a_1, q) < \dots < F_{t, t+n}(a_m, q)$ for all horizons $t + n \geq t$.

Part (a) establishes that if (D) and (E) hold, then inventory just smooths spot prices in that the demand-state order of shocks in (D) is preserved in the equilibrium spot prices. However, (D) alone is not sufficient for this

spot-price ordering when (the number of states) $m > 2$. Without an assumption on the conditional transition probabilities such as (E), a “low” demand state (according to (D)) could be associated with a large probability of high future demand. In such a case the resulting large equilibrium inventory demand might induce a high spot price despite the low current demand. The stochastic dominance assumption in (E), however, is sufficient to ensure that low states lead to low equilibrium spot prices.

Part (b) of the proposition establishes that, given (D) and (E), the forward prices also preserve the demand-state order. In other words, given any fixed level of incoming storage, the forward curve in any state, a_ℓ , lies everywhere below (i.e., does not cross) the forward curve in each higher-demand state, $a_{\ell+1}, \dots, a_m$. A key implication of Proposition 4(b) is that forward prices all move in the *same* direction and, thus, that hedge ratios (using one maturity forward contract to hedge another) are all positive.

A perturbation of the demand state from a low state a_ℓ to a higher state a_h has two effects on the distribution of future spot prices. First, the first-order stochastic dominance assumption of (E) improves the distribution of future demand states. Second, the perturbation affects current equilibrium inventory. In the case where inventory is reduced, $J(a_h, Q_{t-1}) < J(a_\ell, Q_{t-1})$, the effect on forward prices is reinforced. However, if inventory is increased, $J(a_h, Q_{t-1}) > J(a_\ell, Q_{t-1})$, this partially offsets the improved demand state distribution.¹⁴ To see what can happen without an assumption like (E), consider our simple two-state example. With just two states, $\{a_L, a_H\}$, spot prices at each date t are ordered, $P(a_H, Q_{t-1}) > P(a_L, Q_{t-1})$, even without (E). The Markov process violates (E) if $\pi(a_H|a_H)$ and $\pi(a_L|a_L)$ are less than 0.5. In this case a realization of a_L at time t implies a higher conditional probability of a_H at $t + 1$ than does a_H at t . Thus, perturbing the current demand state from a_L to a_H can lead to *lower* forward prices at odd horizons and *higher* prices at even horizons. This will occur if the difference in ΔQ_t is small (e.g., if δ is large). Therefore without assumption (E), forward prices can move in opposite directions.

C. Numerical Example (Continued)

Our numerical example has only two net-demand states, a_L and a_H , but the dynamics of inventory induce an endogenous binomial tree of forward curves over time. Figure 3 illustrates part of the tree. Four features deserve comment. First, since the starting node here at t is a stockout with zero incoming inventory, this node and its two successors are repeated at date $t + 1$ if the next shock is another high realization a_H . This is the renewal feature of inventory and prices (see Corollary 1.2). Second, the degree of

¹⁴ In economies with more than two demand states ($m \geq 3$), it is possible that higher demand states may lead to higher current inventory levels. For example, if the lowest state has a very high conditional probability of recurring, then equilibrium storage in that state can be zero while storage in an intermediate state is strictly positive. Despite this, spot prices and forward prices increase with the demand state.

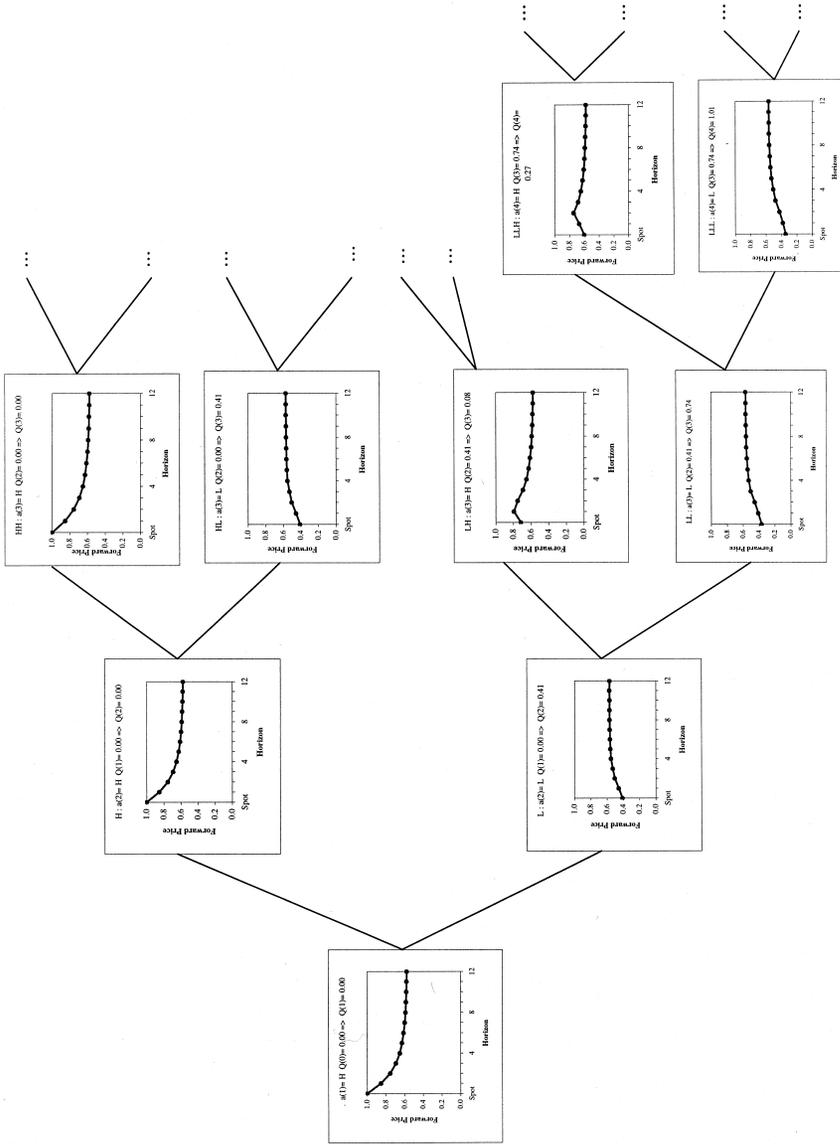


Figure 3. Binomial tree of forward curves. Part of the binomial tree of forward curves is shown, starting in the high state, $a(0) = a_H$, following a stockout, $Q(0) = 0$, in the previous period. The baseline parameters are used here: $a_H = 1$, $a_L = 0$, $\pi(a_H|a_H) = \pi(a_L|a_L) = 0.75$, $r = 0$, and $\delta = 0.1$.

backwardation or contango of the forward curve depends on how much inventory has been accumulated. Thus, an a_L at $t + 1$ (making inventory cheap) can flip the curve from backwardation to contango. Third, the model can also produce hump-shaped curves. Moreover, the hump moves further back along the forward curve if a run of low shocks leads to a buildup of inventory and thus pushes the possibility of a stockout further into the future. Fourth, the tree is nonrecombining. The forward curve given a shock sequence $\{a_H, a_L\}$ is very different from the forward curve given the reverse sequence $\{a_L, a_H\}$ due to the equilibrium implication of the nonnegative inventory constraint.

D. Spanning

Although spot prices are binomial here, the inability to short the commodity in a stockout limits the payoffs that can be spanned with the spot good.¹⁵ However, our example is still dynamically complete using a bond and a one-period forward contract. Therefore, our framework can determine a market price of convenience yield risk in a more complex model with risk-averse agents. This property can be extended to economies with binomial (rather than just two-state) shock processes.¹⁶

PROPOSITION 5 (Spanning): *The market is dynamically complete using a one-period forward contract and a riskless bond if at each date t the Markov process is binomial and satisfies (D) and (E).*

Unfortunately, extending this result to general Markov processes is not straightforward. The path dependence of equilibrium inventory prevents the use of standard arguments of generic completeness. The path dependence implies that there are uncountably many paths. Therefore, in order to check the dynamic completeness, one needs sufficient conditions for the payoff matrix using forward prices and a riskless bond to be invertible for all incoming inventory and demand state pairs. Fortunately, the two-state and binomial settings are quite flexible.

E. Forward Curve Slopes

The relationship between contemporaneous spot and forward prices is usually described either in terms of the slope of the forward curve or with implied convenience yields. We consider slopes first and then restate these results for convenience yields in Section II.F.

¹⁵ In addition, replication with a *long* spot position is a dominated strategy in a stockout (from equilibrium condition (3)). Ross (1976) and Breeden and Litzenberger (1978) discuss using derivative securities to complete the market and Raab and Schwager (1993) consider replication with short-sale restrictions on some assets.

¹⁶ By binomial we mean that, given any shock a_t at any time t , only two shocks $a_j \in \Omega$ have positive probability, $\pi(a_j|a_t) > 0$. In contrast to our ongoing example, these need not be the same two shocks for each a_t .

Since forward prices are positive, we can define the one-period (scaled) slope with respect to delivery horizon, n , as

$$F'_{t,t+n} = \frac{F_{t,t+n+1} - F_{t,t+n}}{F_{t,t+n}}. \quad (8)$$

A forward curve is in *contango* if it is upward sloping ($F' > 0$) and in *backwardation* if it is downward sloping ($F' < 0$).¹⁷ Storage and interest carrying costs make forward curves slope upward; negative slopes are possible due to current or future potential stockouts.

PROPOSITION 6 (Forward Slopes): *The scaled slopes of forward prices in the stationary rational expectations equilibrium are bounded by $F'_{t,t+n} \leq (r + \delta)/(1 - \delta)$ for all t and horizons $n \geq 1$. Additionally, if $\Pi \gg 0$, then*

- (a) *if $F'_{t,t+n} < (r + \delta)/(1 - \delta)$ for a delivery horizon n , then $F'_{t,t+h} < (r + \delta)/(1 - \delta)$ for all later delivery horizons $h > n$,*
- (b) *for every date t there is a future date $\tau^* \geq t$ such that the probability of future backwardation, $\text{Prob}(F'_{\tau,\tau+1} < 0 | a_t, Q_{t-1})$, is strictly positive for all dates $\tau \geq \tau^*$, and*
- (c) *the forward curve at long horizons is “flat” with $\lim_{n \rightarrow \infty} F'_{t,t+n} = 0$.*

Part (a) follows from the equilibrium condition (3) and iterated expectations. If outgoing inventory is positive, then condition (3a) holds and the forward slope is positive and at its maximum, $(r + \delta)/(1 - \delta)$, between P_t and $F_{t,t+1}$. If, instead, current inventory is zero and condition (3b) holds strictly, then the slope is less than the maximum. Using iterated expectations, the argument extends to longer horizons. Parts (a) through (c) all use the assumption that $\pi(a|a_t) > 0$. If the probability of a stockout at a future date $\tau^* > t$ is positive, then it is also positive at all longer horizons since “sell” states have a positive probability of recurring. Strictly positive transition probabilities also imply that the combination of the highest possible demand state and zero incoming inventory has a positive probability. Since this produces the highest possible spot price (P_{max}), the forward curve must have a negative slope in this situation. Finally, a limiting forward slope of zero is an immediate consequence of Proposition 3.

¹⁷ Some authors (e.g., Litzenger and Rabinowitz (1995)) distinguish between strong and weak backwardation. Strong backwardation occurs when the forward price is lower than the spot price (i.e., $F_{t,t+n} < P_t$). Weak backwardation is when the forward price is lower than the spot price adjusted for time value of money (i.e., $F_{t,t+n} < P_t(1 + r)^n$). Other authors (e.g., Keynes (1930)) use the term *backwardation* to refer to the (average) difference between forward prices and the future spot price. Given risk neutrality here, we restrict attention to the contemporaneous spot-forward relationship.

F. Convenience Yields

Another way to describe forward curves is in terms of convenience yields, which express the interim benefits accruing to the physical owners of a commodity as a rate. In our model, this benefit is the timing option to consume the good in high demand states (or to sell it for consumption by someone else) and then buy it back at a lower expected price in the future. Given any forward curve in our model, we can calculate the corresponding implied endogenous convenience yields using the cost-of-carry relation. Thus, convenience yields are an *output* of our model rather than an *input* as in Brennan (1991), Schwartz (1997), and Amin et al. (1995).

The one-period *convenience yield*, $y_{t,\tau}(a_t, Q_{t-1})$, over the interval τ to $\tau + 1$ given the spot and forward prices at date t is defined implicitly by

$$F_{t,T}(a_t, Q_{t-1}) = P(a_t, Q_{t-1}) \left(\frac{1+r}{1-\delta} \right)^{T-t} \prod_{\tau=t}^{T-1} (1 - y_{t,\tau}(a_t, Q_{t-1})). \tag{9}$$

Since interest rates and storage costs are constant here, this can be written as¹⁸

$$y_{t,\tau} = 1 - \frac{(1-\delta) F_{t,\tau+1}}{(1+r) F_{t,\tau}}. \tag{10}$$

Using this definition, the results in Proposition 6 can be rewritten to show that the implied convenience yield in our model is always well defined and nonnegative, and that $y_{t,\tau} < 1$.

PROPOSITION 7 (Convenience Yield): *In the stationary rational expectations equilibrium,*

- (a) *implied convenience yields are nonnegative, $y_{t,\tau} \geq 0$,*
- (b) *the convenience yield is positive, $y_{t,\tau} > 0$, if and only if the stockout probability is positive, $\text{Prob}(Q_\tau = 0) > 0$, and*
- (c) *in an economy with $\Pi \gg 0$,*
 - (i) *if $y_{t,\tau} > 0$, then $y_{t,\tau'} > 0$ for all later delivery dates, $\tau' > \tau$, and*
 - (ii) *the limiting convenience yield is constant with $\lim_{\tau \rightarrow \infty} y_{t,\tau} = (\delta + r)/(1 + r)$.*

Convenience yields inherit their option-like quality from the embedded timing (i.e., storage) option in spot prices. Convenience yields are strictly positive only if there is a positive probability of a stockout (which is like an option exercise). As in Proposition 6, if “sell” states repeat with positive prob-

¹⁸ This is the discrete-time analogue to the instantaneous convenience yield in continuous-time models such as Amin et al. (1995) or Schwartz (1997). Convenience yields are typically expressed net of storage. However, since storage costs are constant here, this distinction is not important.

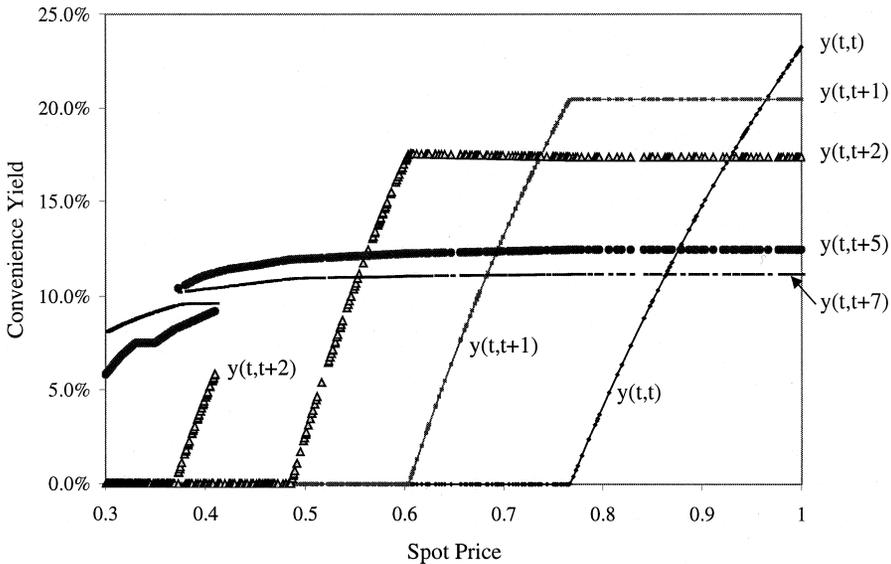


Figure 4. Convenience yields versus spot price. The convenience yield is plotted at different horizons against the corresponding spot price given the baseline parameters: $a_H = 1$, $a_L = 0$, $\pi(a_H|a_H) = \pi(a_L|a_L) = 0.75$, $r = 0$, and $\delta = 0.1$.

ability (e.g., if $\Pi \gg 0$), then a positive convenience yield at one horizon implies positive convenience yields at all longer horizons. Finally, the limiting convenience yield, like the limiting forward price, does not depend on the current demand state or inventory level.

Comparing the implied (equilibrium) convenience yields in our model with the usual (exogenous) convenience yield processes reveals some differences. First, our model does not reduce to a one-factor Brennan (1991) model where convenience yield is exogenously specified as a function of the spot price. Here, convenience yields are functions of both the exogenous demand state and the endogenous inventory level. Second, Schwartz (1997) assumes that innovations in spot prices and instantaneous spot convenience yields have a constant correlation. This in turn implies that the correlation between spot prices and convenience yields at longer horizons, $y_{t,\tau}$ is also constant. In our model, however, the correlation between the spot price and the one-period spot convenience yield, $y_{t,t}$, depends on the endogenously determined inventory level, Q_{t-1} (Proposition 7(b)). In particular, $y_{t,t}$ is nonzero only when outgoing inventory is zero. At longer horizons the convenience yield $y_{t,\tau}$ depends (from Proposition 7(b)) on both the probability and likely severity of a stockout at date τ and, therefore, on the current inventory Q_t .

Figure 4 illustrates these points by plotting the implied convenience yields versus the corresponding spot price in our numerical example. When the current spot price is high at t , the corresponding one-period convenience yield is also high. High spot prices are the result of both high current de-

mand, a_H , and low incoming storage. This situation leads to a stockout. P_t is decreasing in incoming inventory since the inventory buffers the high demand shock. However, $F_{t,t+1}$ is unaffected in a stockout since outgoing inventory is zero. This is the renewal feature seen in Figures 2 and 3. The result is that $y_{t,t}$ is decreasing in incoming inventory and is maximized when $Q_{t-1} = 0$. At low spot prices, outgoing inventory Q_t is positive and, therefore, $y_{t,t} = 0$.

At longer horizons, $n > 1$, the spot-price/convenience-yield relation is also state-dependent, but the nonlinearity is much less dramatic. Compare, for example, the convenience yield at the seven-period horizon, $y_{t,t+7}$, with $y_{t,t}$ in Figure 4. Since our limiting convenience yield is constant, the spot-price/convenience-yield correlation falls to zero (i.e., a constant) as the horizon lengthens. The convenience yield $y_{t,t+n}$ depends on the *probability* of a future stockout at date $t + n$ and its likely severity. The probability and severity depend on the *current* state and inventory. However, due to the renewal feature, the dependence on the current state at very long horizons is small.¹⁹

G. Conditional and Unconditional Volatilities

The volatility of forward prices at different horizons is important for both derivative security pricing and dynamic hedging. Empirically, commodities typically exhibit a pattern of forward price volatility which is declining with contract horizon known as the Samuelson (1965) effect. This is usually attributed to the smoothing of expectations over a mean-reverting process. However, in our model the Samuelson effect need not hold *conditionally* in all states at all horizons.

PROPOSITION 8 (Volatility): *In the stationary rational expectations equilibrium of an economy with $\Pi \gg 0$, the volatility of forward prices satisfies the following:*

- (a) *There is a horizon N such that for all $a_t \in \Omega$ and q_{t-1} , the conditional forward price variance, $\text{var}(F_{t+1,t+1+n} | a_t, q_{t-1})$, is decreasing in the horizon length n for all $n > N$.*
- (b) *There is a demand/inventory/contract-horizon combination (a_t, q^*, N) such that the conditional forward price volatility, $\text{var}(F_{t+1,t+1+n} | a_t, q_{t-1})$, is increasing in n for short horizons $n \leq N$ when inventory is sufficiently high, $q_{t-1} \geq q^*$.*

Part (a) says that, for sufficiently long horizons, forward price volatilities decline with maturity *regardless* of the current demand state and inventory. This follows immediately from Proposition 3 where the tightening bounds on

¹⁹ In Figure 4 the fact that the convenience yield can take two possible values at some intermediate spot price levels (e.g., around 0.42 for $y_{t,t+2}$) is an illustration that our convenience yield cannot be represented as a function of only the spot price as in Brennan (1991). A low state a_L and low inventory or a high state a_H and high inventory can both lead to the same spot price. However, since the forward prices are different in these situations, so are the implied convenience yields. Of course in our two-state example, very high spot prices are only consistent with a_H . Note this feature can also be observed in $y_{t,t+5}$ and $y_{t,t+7}$.

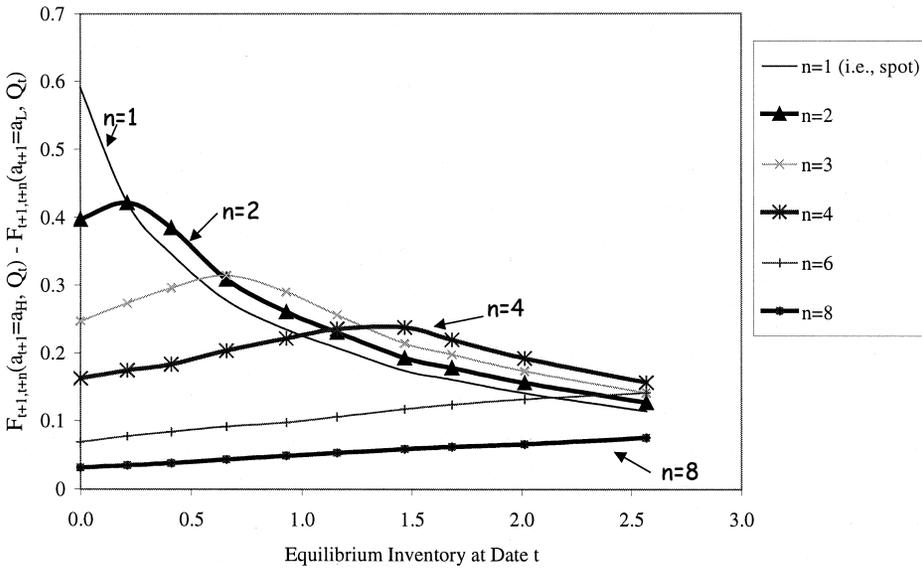


Figure 5. Conditional dispersions. The conditional dispersion, $F_{t+1,t+n}(a_H, q_t) - F_{t+1,t+n}(a_L, q_t)$, of date $t + 1$ futures prices is plotted for different horizons, n , as a function of the date t inventory, q_t , given the baseline parameters: $a_H = 1$, $a_L = 0$, $\pi(a_H|a_H) = \pi(a_L|a_L) = 0.75$, $r = 0$, and $\delta = 0.1$.

forward prices and existence of a limiting forward price, F_∞ , are consequences of the finite Markov structure (with $\Pi \gg 0$) and the regeneration property of inventory. However, part (b) states that conditional violations of the Samuelson effect can occur at shorter horizons. In particular, when current inventory is high, the spot price is less volatile than the short horizon forward. Fama and French (1988) document empirically such a relation between inventory and the conditional volatilities of different horizon forward prices.

The relation between volatility and inventory is easily seen in our two-state numerical example. Consider the difference $F_{t,t+n}(a_H, Q_{t-1}) - F_{t,t+n}(a_L, Q_{t-1})$ as a measure of *conditional dispersion*. Since the probability at $t - 1$ of the forward prices $F_{t,t+n}(a_H, Q_{t-1})$ and $F_{t,t+n}(a_L, Q_{t-1})$ at t are just $\pi(a_H|a_{t-1})$ and $\pi(a_L|a_{t-1})$ regardless of the particular forward horizon n , conditional dispersion is one-to-one increasing in volatility in this example. In Figure 5 the conditional dispersions are ordered by contract horizon n at low inventory levels. In particular, if inventory is zero ($Q_{t-1} = 0$), the dispersion (volatility) is determined solely by the Markov demand process and, therefore, declines with maturity. However, at high levels of current inventory Q_{t-1} , forward price volatilities are potentially increasing in maturity at several horizons. With enough inventory, stock-outs may not be possible for n periods so that the n -period forward price is simply the spot price scaled up by the cost-of-carry relation, $F_{t,t+n} = \theta^{-n} P_t$, where $\theta^{-n} > 1$ (see equation (3a)). In such cases, the longer-term (n pe-

riod) forward prices will be more volatile than the shorter-term ($n - 1$ period) forward prices. In Figure 5, for example, the volatility of the four-period forward price exceeds that of the one-period contract when inventory exceeds half of its endogenous maximum, Q_{max} . In this case, there is little probability of a stockout in the next three periods. In contrast, the volatility of the eight-period contract is always lower than the volatility of shorter horizon contracts. The probability of a stockout eight periods later, even if inventory is near its endogenous maximum, is always large. Finally, if storage is costly enough (large δ) so that inventory is always low (small Q_{max}), then expectations of future spot prices are not heavily influenced by inventory and conditional violations of the Samuelson effect do not occur (in particular, if $Q_{max} = 0$).²⁰

Inventory directly affects the unconditional volatility of forward prices. Deaton and Laroque (1992, 1996) and Chambers and Bailey (1996) show that storage smooths the impact of demand shocks on the spot price. This smoothing, in turn, influences the unconditional forward price volatilities. A high cost of storage, δ , leads to relatively few transmission of shocks via inventory across periods and forward price volatility declines rapidly in the forward horizon (due to the Markov property). In contrast, with a low δ considerable transmission of shocks from short horizons to longer horizons occurs via inventory, hence forward price volatility declines less rapidly. These intertemporal volatility smoothing effects of inventory (and their dependence on δ) can be significant as illustrated in Figure 6 for our numerical example. As discussed below in Section III, they are also important in calibrating our model to oil futures prices.

H. Conditional Hedging

The large losses of Metallgesellschaft in 1993 generated widespread interest in how to hedge long-horizon forward positions dynamically using short-horizon futures contracts.²¹ In our model, short-term and long-term forward prices differ in that long-term forward prices do not include the interim option to consume the good between the two delivery dates. Since the value of this timing option depends on the likelihood of a stockout during this time interval, hedge ratios are state dependent. They depend on the current demand state and inventory level.

²⁰ There are other explanations for violations of the Samuelson effect. For example, it is straightforward to generate violations of the Samuelson effect in the absence of inventory by introducing a seasonal component to the demand shocks (see, for example, Anderson and Danthine (1983)). Alternatively, Hong (2000) generates violations of the Samuelson effect in a model with exogenous, mean-reverting convenience yield and asymmetric information. He constructs a model where the informativeness (hence, the volatility) of the forward price decreases as contracts mature. This information effect offsets the convenience yield mean-reversion and can produce a U-shaped term-structure of volatility.

²¹ See Culp and Miller (1995), Edwards and Canter (1995), and Mello and Parsons (1995) for discussions of Metallgesellschaft and see Ross (1997), Neuberger (1999), Brennan and Crew (1997), and Ronn and Xuan (1997) for more generic treatments of hedging long-dated exposures with near-term contracts.

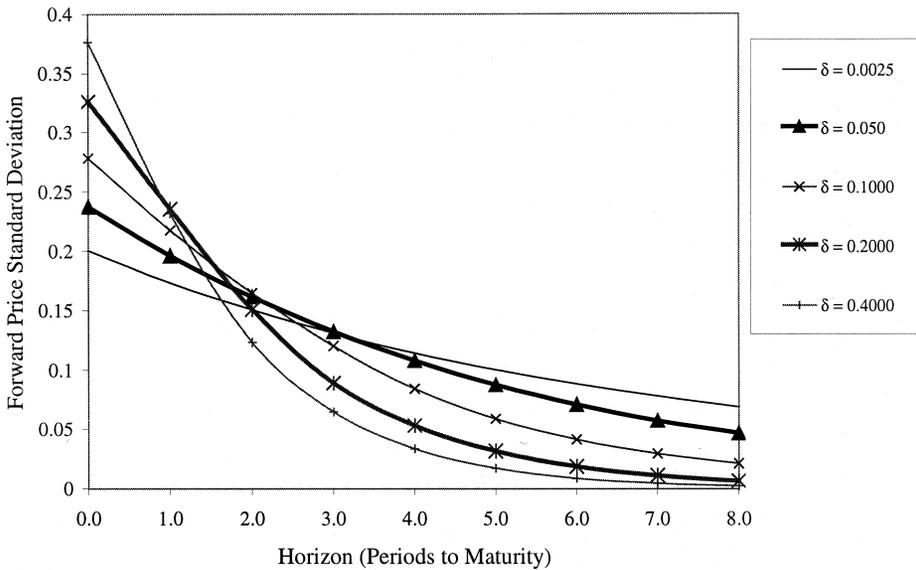


Figure 6. Spot and forward price standard deviations and storage costs. The unconditional standard deviation of spot and forward prices is plotted at different horizons. The baseline parameters are used, $a_H = 1$, $a_L = 0$, $\pi(a_H|a_H) = \pi(a_L|a_L) = 0.75$, and $r = 0$, except that the storage cost, δ , is varied.

This general point is easily illustrated in our two-state example. If a short-term forward (with delivery at date τ) and a riskless bond are used to hedge a long-dated forward (with delivery at $T > \tau$) between dates t and $t + 1$, then the number of short-dated contracts needed is

$$\frac{F_{t,T}(a_H, q) - F_{t,T}(a_L, q)}{F_{t,\tau}(a_H, q) - F_{t,\tau}(a_L, q)} (1 + r)^{\tau - T}. \tag{11}$$

The hedge ratio in equation (11) is just the ratio of the two conditional dispersions (from the previous section) with an adjustment for the different timing of cash flows for the two forward contracts. Properties of the hedge ratio are an immediate corollary of Propositions 4 and 8.

COROLLARY 8.1 (Hedge Ratios): *In a two-state economy with $\Omega = \{a_L, a_H\}$ and $\Pi \gg 0$, the hedge ratio in equation (11) at date t is:*

- (a) *positive if the economy satisfies (D) and (E)*
- (b) *less than one for large $T - \tau$*
- (c) *greater than one for a sufficiently small $T - \tau$ and large enough inventory, q .*

Figure 7 plots the hedge ratio for hedging an n -period forward using a one-period forward and a riskless bond (recall that $r = 0$ here). The declining long-horizon volatilities from Proposition 8(a) imply that hedge ratios for

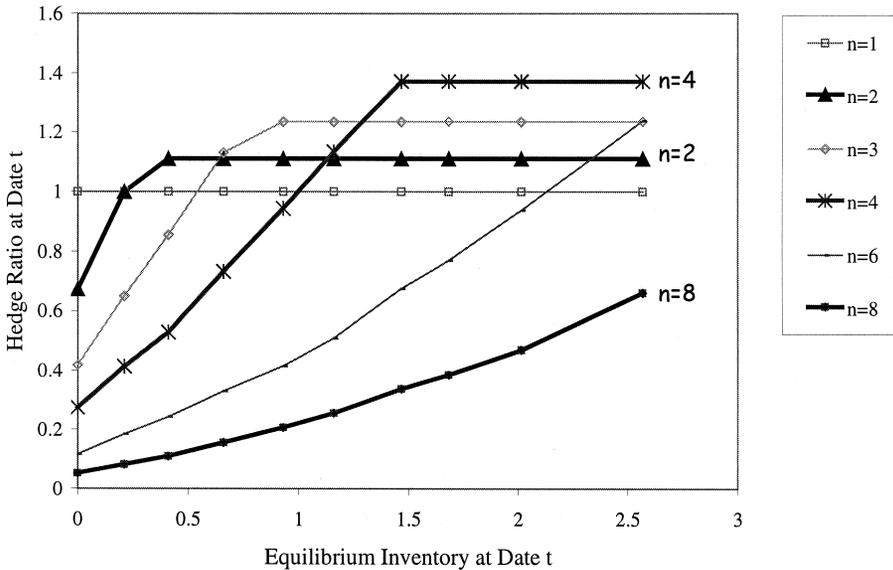


Figure 7. Hedge ratios. The figure plots the number of one-period forward contracts (with forward price $F_{t,t+1}$) needed to create a riskless hedge for one longer-horizon n -period contract (with forward price $F_{t,t+n}$) given the binomial tree induced by the baseline parameterization: $a_H = 1$, $a_L = 0$, $\pi(a_H|a_H) = \pi(a_L|a_L) = 0.75$, $r = 0$, and $\delta = 0.1$.

very long-dated contracts (e.g., $n = 8$ in Figure 7) are less than one. However, the hedge ratio for intermediate horizons can be *larger* than one with enough inventory. For example, the hedge ratio for a four-period ($n = 4$) forward exceeds one when $Q_t > 0.99$.

III. Calibration

This section illustrates how to calibrate two versions of our model to market data. The first version is a one-factor power specification that nests our ongoing numerical example. The second version is a multifactor generalization that allows for permanent as well as transitory shocks. The commodity used in this calibration is crude oil for the post-Gulf war period of 1992 through 1996. Our data are daily futures prices for the Nymex light sweet crude oil contract. The data are sorted by contract horizon with the “one-month” contract being the contract with the earliest delivery date, the “two-month” contract having the next earliest delivery date, etc. We consider contracts with up to 10 months to delivery since liquidity and data availability are good for these horizons. As is typical, spot prices are not available (see Schwartz (1997)).

Of the many possible calibration approaches, we follow Backus and Zin (1994) and calibrate the model to the historical means and standard deviations. In particular, we try to replicate the unconditional and conditional

moments of futures prices (see Figure 8). In calculating conditional moments we classify the futures curve at each date t as to whether one month (20 trading days) earlier the curve was in contango or backwardated (based on the then prevailing one-month and six-month prices). Not surprisingly, the conditional mean futures curve at t still tends to be in contango or backwardation a month later. More interestingly, the term structure of conditional futures price volatilities differs following backwardation or contango. It is precisely this type of heteroskedasticity which inventory induced in our model in Section II.

In calibrating the model, our goal is not to test the model formally. Rather we simply want to identify dimensions along which extremely parsimonious specifications of the model do and do not work well.

A. One-Factor Calibration

Our first specification is a slight generalization of our ongoing two-state numerical example. Instead of the linear specification, we allow for some curvature by assuming a power specification for the inverse net demand function

$$P_t = f(a_t, \Delta Q_t) = (a_t + \Delta Q_t)^\alpha. \quad (12)$$

We continue to use a simple two-state Markov process. However, in order to facilitate interpreting the calibrated parameters we choose the two-state Markov process to approximate an AR1 process, $\{A_t\}$.

$$A_t = (1 - \rho_A)\mu_A + \rho_A A_{t-1} + (1 - \rho_A)^{1/2} \sigma_A \epsilon_t, \quad (13)$$

where ϵ_t is an i.i.d. standard normal. The unconditional mean and variance for A_t are μ_A and σ_A^2 and the autocorrelation is ρ_A . The model parameters, a_H , a_L , Π , are calculated from equation (13) using the standard discretization method (see Tauchen and Hussey (1991)). This specification is easily extended to include more demand states.

From Proposition 1(c) any exponent $\alpha > 0$ is permissible as long as the discretized a values are all constrained to be positive. For the oil calibration parameters we considered, this constraint is not binding.²² The linear numerical example discussed previously is the special case of $\alpha = 1$.

Calibration involves picking the curvature α of the net demand and the mean μ_A , volatility σ_A^2 , and autocorrelation ρ_A for equation (13). For simplicity we fix the carrying charges r and δ . The interest rate r is fixed at a monthly rate of 0.04/12 or 0.33 percent. We justify this by noting that commodity prices are considerably more volatile than interest rates. Moreover,

²² Since the search algorithm may explore candidate (i.e., nonequilibrium) inventory functions which could result in $a_t + \Delta Q_t < 0$, the algorithm sets $\alpha = 1$ whenever it encounters $a_t + \Delta Q_t \leq 0$. This "off equilibrium" numerical assumption plays no role in the final equilibrium solution.

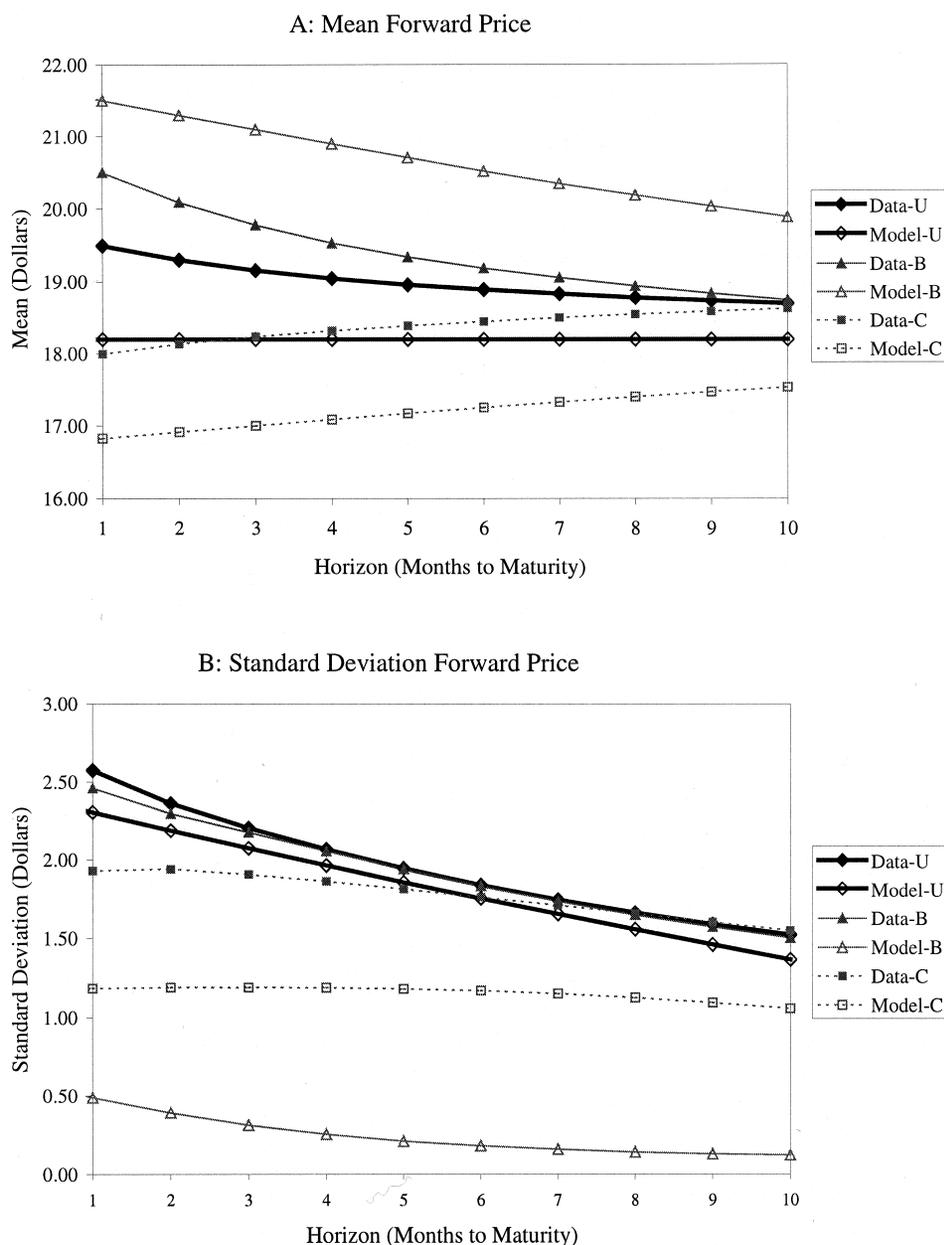


Figure 8. Unadjusted forward price (A) means and (B) standard deviations—empirical and “basic model.” The data are the daily Nymex crude oil futures prices from 1992 to 1996. For the data, the unconditional (U) moments are the sample averages and standard deviations and the conditional moments given backwardation (B) and contango (C) are conditioned on the shape of the forward curve one month (20 days) prior. For the “basic model,” the unconditional moments (U) are calculated from the limiting Markov distribution. The conditional moments (B) and (C) are calculated using the limiting Markov distribution, conditioning on the shape of the curve one month prior.

Schwartz (1997) finds that empirically an interest rate factor does not add much in explaining futures prices with less than a year to delivery. We also fix the storage cost δ to a monthly value of 0.0025 or 0.25 percent. This is significantly lower than what is implied by positively sloped forward curves in the data (using equation (3) to back out a bound on δ). However, a small δ is needed for inventory in our model to play an important role. Deaton and Laroque (1992) also found that storage costs implied by the data were quite low. We chose to fix the δ (after some preliminary exploration) to avoid extremely low values of δ where the numerical approximation of the inventory function is less precise.

The calibration problem is as follows: Let $\omega \in \{U, B, C\}$ index whether the futures price moments are unconditional (U) or are conditioned on the previous month's futures curve having been backwarddated (B) or in contango (C). For each ω let $\hat{\mu}_n(\omega)$ and $\hat{\sigma}_n(\omega)$ denote the mean and standard deviation of the n -period futures (i.e., forward) price in the model²³ and let $\mu_n(\omega)$ and $\sigma_n(\omega)$ denote their empirical counterparts from the Nymex data. Since prices are stationary in the model, these are price levels rather than returns (e.g., as in Schwartz (1997)). We then minimize the difference between the moments in the data and those implied by the model

$$\min_{\{\mu_A, \sigma_A^2, \rho_A, \alpha\}} \sum_{\omega \in \{U, B, C\}} \left(\sum_{n=1}^{10} [\hat{\mu}_n(\omega) - \mu_n(\omega)]^2 + \sum_{n=1}^{10} [\hat{\sigma}_n(\omega) - \sigma_n(\omega)]^2 \right). \quad (14)$$

The solution to equation (14) is reported in the unscaled “basic model” column of Table I along with various diagnostic statistics. Figure 8 plots the empirical and calibrated means and standard deviations. The model fits the unconditional volatility term structure extremely well. To do this, however, it needs a high level of autocorrelation ($\rho_A = 0.64$). This is similar to results in Deaton and Laroque (1992, 1996). However, the model has difficulty matching the conditional volatilities. Not only are they too low in the model, but their ordering (i.e., given backwardation and contango) is opposite from the data. The model also is unable to replicate the conditional means. Part of the problem is that inventory given backwardation is too low (with a mean of 0.97 units) and not variable enough (with a standard deviation of 1.75).²⁴ As a result, the price variability following backwardation is induced largely by just the shock process. Despite the high persistence of the calibrated shocks, the model volatilities die off too quickly relative to the empirical conditional volatilities. Thus, the model is unable to match both the conditional and unconditional moments of futures prices simultaneously.

²³ Also recall that constant interest rates imply forward and futures prices are identical (see Cox, Ingersoll, and Ross (1981)).

²⁴ Inventory is measured in units needed to move the spot price \$1 in the case of $\alpha = 1$. Thus, if an extra 1.75 units of inventory were “dumped” on the market all at once, prices would be depressed by \$1.75.

Table I
Calibration Results

Empirical moments are reported for daily Nymex crude oil futures prices from 1992 to 1996 and the model parameters and summary statistics calibrated to these data. In columns 1 and 2 the futures prices are unscaled for both the data and the “basic model.” In columns 3, 4, and 5, the futures prices for the data, “basic model,” and “augmented (full) model” have first been normalized by the 10-month futures price as follows: $(F_{t,t+n}/F_{t,t+10}) \times \text{mean}(F_{t,t+10})$. A forward curve is classified as being in backwardation (or contango) if the one-month and six-month forward prices satisfy $F_{t,t+6} - F_{t,t+1} < (>) 0$. A forward curve has a hump from the spot price if $P_t < F_{t,t+1}$ and $F_{t,t+1} > F_{t,t+2}$. Similarly, a forward curve has a hump from the one-period forward price if $F_{t,t+1} < F_{t,t+2}$ and $F_{t,t+2} > F_{t,t+3}$. The conditional moments for inventory (and for forward prices in Figures 8, 9, and 10) are conditioned on the shape of the forward curve one month prior.

		Normalization				
		Unscaled		By $F_{t,t+10}$		
		Data	Basic Model	Data	Basic Model	Augmented Model
Parameters						
Demand state mean	μ_A		16.1992		16.1992	17.7732
Demand volatility	σ_A		6.9988		6.9988	9.8742
Demand state autocorrelation	ρ_A		0.6370		0.6370	0.2462
Proportional storage cost	δ		0.0025		0.0025	0.0025
Curvature of net-demand	α		1.0092		1.0092	1.0172
Interest rate	r		0.33%		0.33%	0.33%
Forward prices						
Means reported in Figures		8(a)	8(a)	9(a) & 10(a)	10(a)	9(a)
Standard deviations reported in Figures		8(b)	8(b)	9(b) & 10(b)	10(b)	9(b)
Skewness of $F_{t,t+1}$		6.199	0.387	0.773	0.784	0.868
Kurtosis of $F_{t,t+1}$		1.792	-1.622	0.566	-1.253	0.025
Frequency backwardation		59.81%	32.00%	59.81%	32.00%	30.75%
Frequency of hump (from Spot)		n/a	4.44%	n/a	4.44%	13.34%
Frequency of hump (from $F_{t,t+1}$)		19.03%	3.02%	19.03%	3.02%	8.31%
Equilibrium inventory						
Unconditional mean			21.981		21.981	18.855
Mean conditional on backwardation			0.969		0.969	4.224
Mean conditional on contango			31.870		31.870	25.350
Unconditional standard deviation			17.102		17.102	12.457
Standard deviation conditional on backwardation			1.753		1.753	3.744
Standard deviation conditional on contango			11.096		11.096	8.979

B. Multifactor Calibration

The difficulty with the previous calibration stems from the long-horizon variance of futures prices. As documented in Schwartz (1997), the variance of long-horizon (e.g., seven-year) futures prices is substantial. To account for this, our second specification models the inverse “immediate use” net de-

mand as having two random shocks, one of which is permanent. The result is an equilibrium version of Schwartz and Smith (2000) in which prices have both permanent and transitory components, but where the impact of the transitory shocks is asymmetric depending on the level of inventory.

We add permanent shocks so that they are separable in that they do not affect the equilibrium condition equation (3) and therefore preserve the equilibrium inventory. This preserves many attractive features of the model. Consider two hypothetical economies. In the first the inverse net demand, f , is driven by a single transitory process, a_t , as in Section I. In the second economy, denoted with “*” superscripts, we augment the net-demand function with an additional random variable b_t . Define

$$P_t^* = f^*(a_t, b_t, \Delta Q_t) = b_t f(a_t, \Delta Q_t). \tag{15}$$

The goal is to specify a process for b_t which elevates long-horizon forward price volatility without changing the equilibrium inventory dynamics. In particular, we want $J^*(b, a, q)$ and $J(a, q)$ to be identical across the two economies.²⁵

PROPOSITION 9 (Augmentation): *If $b_t > 0$ and follows a random walk satisfying $E_t[b_{t+1}f(a_{t+1}, \Delta Q)] = b_t E_t[f(a_{t+1}, \Delta Q)]$ for all ΔQ , then inventory in the two economies is identical, $J^*(b, a, q) = J(a, q)$, for all a, b, q , and prices are*

$$P_t^* = b_t P_t \tag{16}$$

$$F_{t,t+n}^* = b_t F_{t,t+n}.$$

In our calibration we use a lognormal specification with $\ln(b_{t+1}/b_t) \sim N(-\frac{1}{2}\sigma_b^2, \sigma_b^2)$ which is independent of a_t . This choice is the discrete time analogue of the lognormal permanent process in Schwartz and Smith (2000) except that our drift is zero. Intuitively, this allows the transitory net demand shock, a_t , and endogenous inventory, Q_t , processes to match the shape of the forward curve while the permanent factor, b_t , accounts for changes in the curve’s level. From equation (16), the scaled slope of the forward curve in equation (8) satisfies $F_{t,t+n}^{*'} = F_{t,t+n}'$. Therefore, all of the properties concerning the shape of forward curves and their implied convenience yields are preserved.²⁶

²⁵ One numerical benefit of this construction is that it avoids the “curse of dimensionality” since solving for J^* in the multifactor economy is exactly the same as solving for J in the original one-factor economy.

²⁶ Another way to introduce permanent shocks is with an additive shock where $f^*(b'_t, a_t, \Delta Q_t) = b'_t + f(a_t, \Delta Q_t)$. The same logic establishes that $J^* = J$ if $b'_t \geq 0$ and $E_t[b'_{t+1}] = \theta^{-1}b'_t$. The resulting spot and forward/futures prices are $P_t^* = b'_t + P_t$ and $F_{t,t+n}^* = \theta^{-n} b'_t + F_{t,t+n}$. In this setting, long-dated forward prices are rising (rather than flat) in the horizon, n , because the additive permanent shock has positive drift ($\theta^{-1} > 1$). Finally, note that the additive and multiplicative shocks can be easily combined to form a three-factor model.

Prices are nonstationary in the augmented model so we break the calibration into two steps. The first step is to calibrate the short end of the futures curve relative to the level of the long end. We do this by defining *normalized forward prices* by dividing through by the contemporaneous 10-month forward price,

$$\tilde{F}_{t,t+n} = \frac{F_{t,t+n}}{F_{t,t+10}} \times \text{mean}(F_{t,t+10}), \quad (17)$$

for both the data and the model. Scaling the normalized prices by the constant mean is done solely to facilitate comparison to the basic model. Normalized prices and calibration parameters are of the same order of magnitude.

Analogously to equation (14) for the one-factor model, we choose the parameters of the stationary component of prices, curvature α , and the parameters for the discretized AR1 process, μ_A , σ_A^2 , and ρ_A . As before, the storage costs are held constant at $\delta = 0.0025$ and $r = 0.04/12$. The parameters are chosen to minimize equation (14) redefined for the normalized prices of equation (17).

The second step of the calibration is to set the volatility, σ_b^2 , of the permanent shock, b_t , to match the unconditional volatility of the raw (unnormalized) 10-month futures price log returns. Given separability, this pins down the variability of the long end of the futures curve. The initial level of the permanent factor, b_0 , is chosen to match the empirical level of the unadjusted forward prices. This two-step procedure uses the 10-month price as a numerical proxy for $F_{t,\infty}$, which in our basic model is constant. Looking at Figures 8a and 8b, for crude oil the conditional and unconditional moments in the data appear to have converged for the 10-month futures. This suggests that the pricing of contracts with 10 months or more until delivery is driven primarily by permanent factors rather than short-term shocks or inventory dynamics.

Table I reports the calibrated two-factor parameters in the Augmented Model column and Figure 9 plots the empirical and fitted moments. To compare the augmented and basic models, Figure 10 re-presents the results for the basic model (in Figures 8) in terms of normalized prices. The two-factor model closely matches both the unconditional and the contango volatilities. The fit of the backwardation volatilities is also dramatically improved. Inventory now plays a greater role in the backwardation state (with a mean of 4.2 and volatility of 3.7). This is a consequence of the fact that (with the permanent factor b_t) the persistence of the transitory shock, ρ_A , is greatly reduced (to 0.25 from 0.64). Some additional statistics are also reported in Table I which, although not used in the actual calibration equation (14), are useful diagnostics. Backwardation is still more frequent in the data than in the model (60 percent vs. 31 percent) which explains why the model undershoots the unconditional mean futures curve despite the good fit of the conditional means. Hump-shaped curves (i.e., initially rising, then falling with

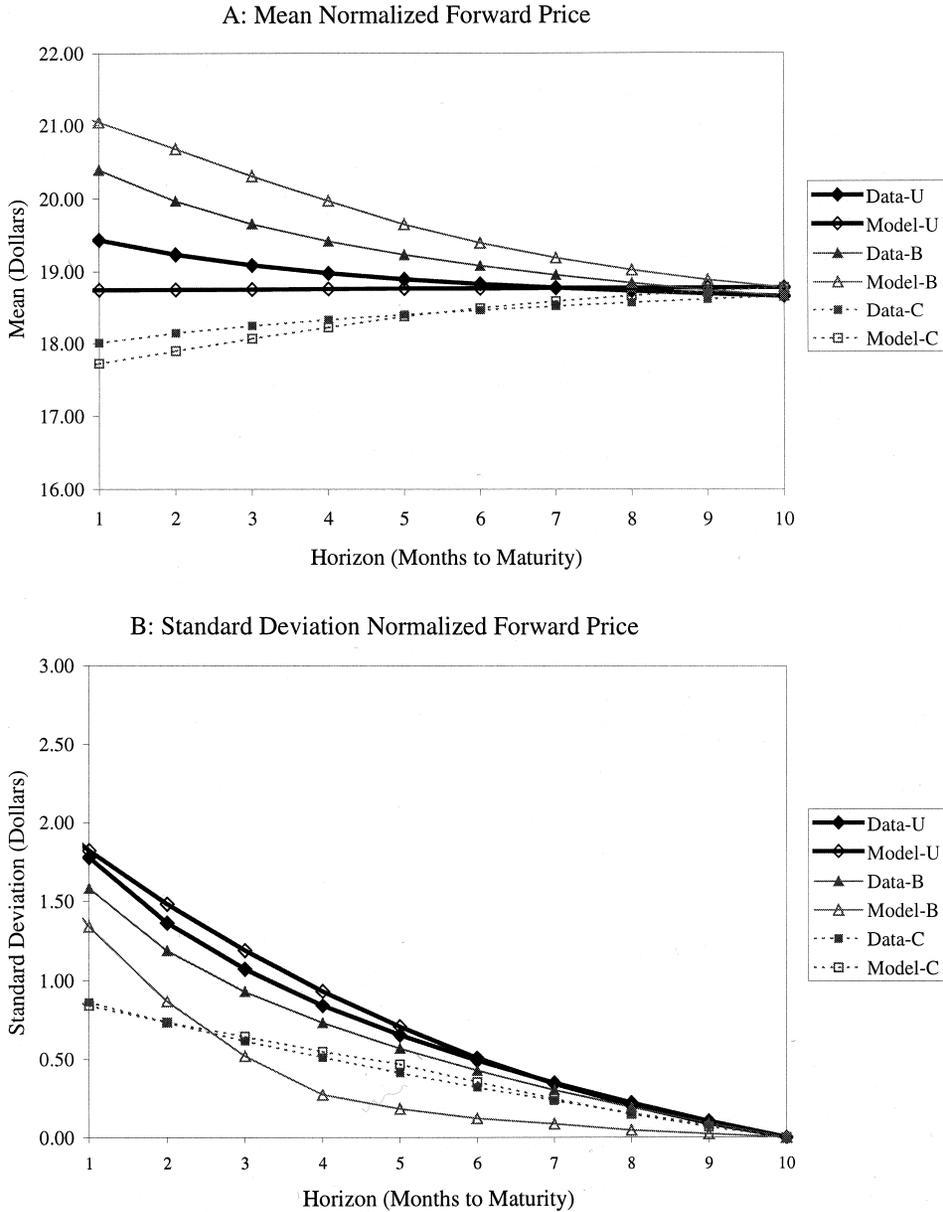


Figure 9. Adjusted forward price (A) means and (B) standard deviations—empirical and “full model.” The data are the daily Nymex crude oil futures prices from 1992 to 1996. All prices in the data and in the model have been normalized by the 10-month futures price as follows: $(F_{t,t+n}/F_{t,t+10}) \times \text{mean}(F_{t,t+10})$. For the data, the unconditional (U) moments are the sample averages and standard deviations and the conditional moments given backwardation (B) and contango (C) are conditioned on the shape of the forward curve one month (20 days) prior. For the “full model,” the unconditional moments (U) are calculated from the limiting Markov distribution. The conditional moments (B) and (C) are calculated using the limiting Markov distribution, conditioning on the shape of the curve one month prior.

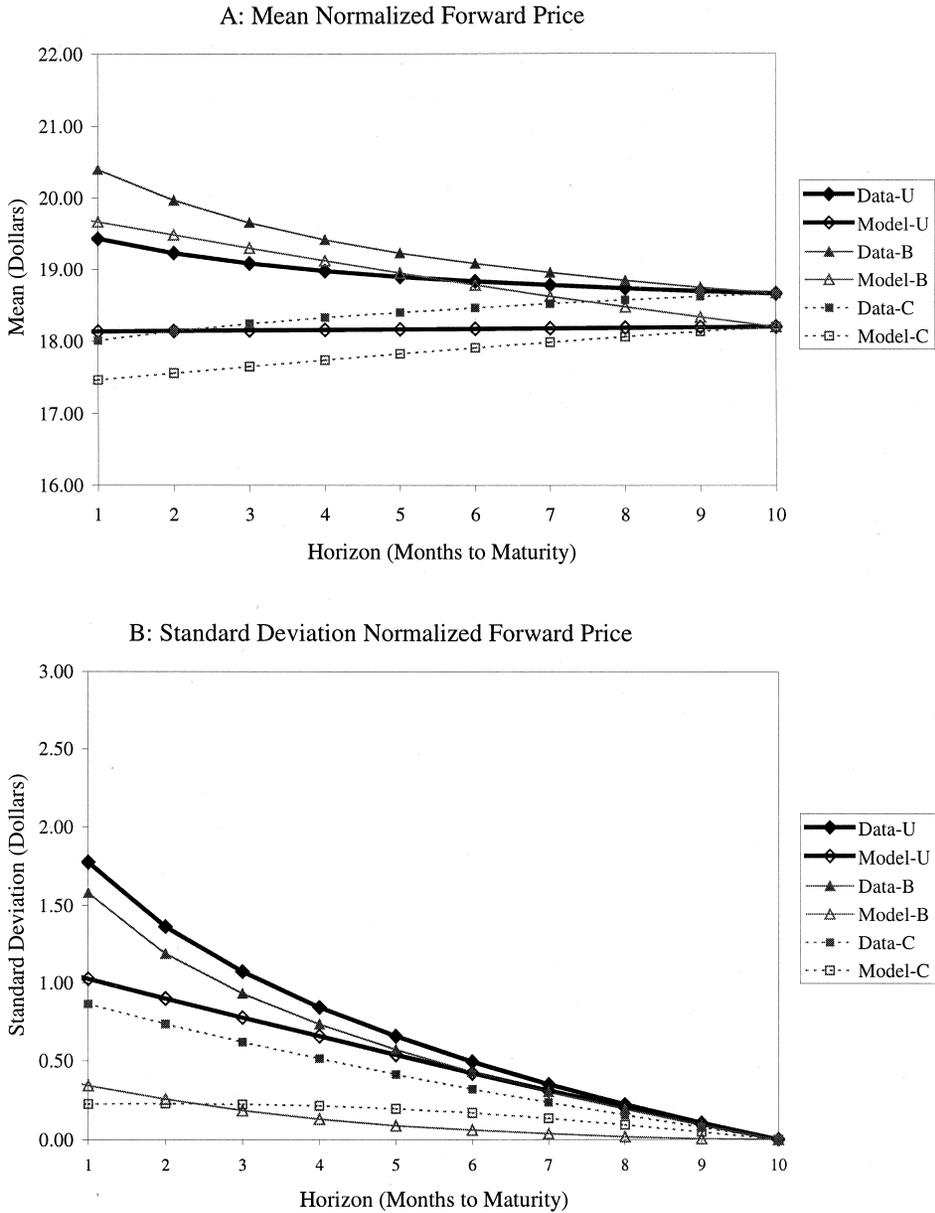


Figure 10. Adjusted forward price (A) means and (B) standard deviations—empirical and “basic model.” The data are the daily Nymex crude oil futures prices from 1992 to 1996. All prices in the data and in the model have been normalized by the 10-month futures price as follows: $(F_{t,t+n}/F_{t,t+10}) \times \text{mean}(F_{t,t+10})$. For the data, the unconditional (U) moments are the sample averages and standard deviations and the conditional moments given backwardation (B) and contango (C) are conditioned on the shape of the forward curve one month (20 days) prior. For the “basic model,” the unconditional moments (U) are calculated from the same limiting Markov distribution used in Figure 8. The conditional moments (B) and (C) are calculated using the limiting Markov distribution, conditioning on the shape of the curve one month prior.

time until delivery) are also less frequent in the model (8 percent vs. 19 percent) although this is improved from the basic model (with only 3 percent). The unconditional distribution of the one-month futures price is slightly more positively skewed in the model (0.87 vs. 0.77), but less kurtotic (0.03 vs. 0.57). The absence of “fat tails” is, in part, simply a consequence of the two-point distribution for the transitory shocks. Note also that the fitted α of 1.017 is effectively a linear specification. In light of the augmented model’s simplicity, the model fares surprisingly well.

IV. Conclusions

In this paper we have developed a simple equilibrium model of commodity forward curves. In particular, rather than exogenously assuming stochastic processes for spot prices and convenience yields (or, equivalently, forward prices), we derive the spot and forward price processes induced by a mean-reverting “immediate use” net-demand process and the resulting equilibrium inventory dynamics. We argue that forward curves for commodities differ from those for stocks or bonds because of a timing option in the ownership of the physical good. When inventory is stocked out (or low) and the current net consumption demand is high, forward prices in the model are backwarddated—or, equivalently, imply positive convenience yields—even in the absence of explicit service flows. We also investigated the volatility structure of forward prices and hedge ratios for forward positions at different horizons. Finally, we presented a tractable two-factor augmentation that helps in calibrating the model to both the unconditional and conditional empirical futures price volatility term structures.

Appendix

Preliminary to the proof of Proposition 1 (Equilibrium), we demonstrate existence of the equilibrium inventory function $J(a, q)$ as the limit of a sequence of finite-horizon economies as the finite horizon $T \rightarrow \infty$. Let $J(t, a, q; T)$ represent the equilibrium inventory function at date t in an economy with a finite horizon, T . It is helpful to extend the function so that $J(t, a, q; T) = 0$ for $q < 0$. We suppress the dependence on horizon T unless needed.

LEMMA A1: *Given a finite horizon, T , for all $a \in \Omega$, $q \in \mathbb{R}_+$ and $t = 1, \dots, T$, equilibrium inventory $J(t, a, q)$ exists, is continuous in q , and satisfies the finite-horizon analogue of equation (4),*

$$0 \leq J(t, a, q + \epsilon) - J(t, a, q) < (1 - \delta)\epsilon. \quad (\text{A1})$$

Proof of Lemma A1: The proof is by induction. Given a finite horizon T , $f(a, \Delta Q) > 0$ implies that $J(T, a, q) = 0$ for all $a \in \Omega$ and $q \in \mathbb{R}_+$. Note that $J(T, a, q)$ is continuous in q and satisfies equation (A1). Assume now that for

time $t + 1$, $J(t + 1, a, q) \geq 0$ exists, is continuous in q , and satisfies equation (A1). We construct $J(t, a, q)$ as follows. Consider the following version of the equilibrium condition given in equation (3a):

$$f(a_t, q^* - (1 - \delta)q) = \theta E_t[f(a, J(t + 1, a, q^*) - (1 - \delta)q^*)]. \quad (\text{A2})$$

The left-hand side is continuous and increasing without bound in q^* . On the right-hand side, for each a , $f(a, \Delta Q)$ is decreasing in q^* since $J(t + 1, a, q)$ satisfies equation (A1). Therefore if q^* exists, it is unique and finite. Define $J^*(t, a_t, q) = q^*$. If the left-hand side of equation (A2) exceeds the right-hand side for all q^* , define $J^*(t, a, q) = -\infty$. Since inventory is restricted to be nonnegative, this case has no effect on the equilibrium construction. To verify that J^* satisfies equation (A1) (when $J^* > -\infty$), perturb q by $\epsilon > 0$. Let $J^*(t, a_t, q + \epsilon) = J^*(t, a_t, q) + \xi = q^* + \xi$. Inserting this into equation (A2) and rearranging implies that

$$\begin{aligned} & f(a_t, q^* - (1 - \delta)q + [\xi - (1 - \delta)\epsilon]) \\ &= \theta E_t[f(a, J(t + 1, a, q^*) - (1 - \delta)q^* + [D(a, \xi) - (1 - \delta)\xi])], \end{aligned} \quad (\text{A3})$$

where $D(a, \xi) = J(t + 1, a, q^* + \xi) - J(t + 1, a, q^*) < (1 - \delta)\xi$ for all $a \in \Omega$ (from equation (A1)). At $\xi = 0$, the left-hand-side of equation (A3) is smaller than the right-hand-side. At $\xi = (1 - \delta)\epsilon$, the left-hand-side exceeds the right-hand-side. Therefore, $0 < \xi < (1 - \delta)\epsilon$ and $J^*(t, a_t, q)$ satisfies equation (A1). This also establishes that J^* is continuous in q . Finally note that if $J^*(t, a, q) < 0$ then

$$f(a_t, -(1 - \delta)q) > \theta E_t[f(a, J(t + 1, a, 0))]. \quad (\text{A4})$$

Therefore $J(t, a, q) = \max[J^*(t, a, q), 0]$ satisfies the equilibrium condition (3), is continuous in q , and satisfies equation (A1). Q.E.D.

LEMMA A2: *Given a finite horizon, T , for all $a \in \Omega$, $q \in \mathbb{R}_+$, and $t = 1, \dots, T - 1$, $J(t, a, q; T) \geq J(t + 1, a, q; T)$.*

Proof of Lemma A2: The proof is a straightforward induction using equation (A2) at t and $t - 1$. Q.E.D.

LEMMA A3: *Given a finite horizon, T , a time-independent upper bound, $K + (1 - \delta)q$, for inventory exists such that $J(t, a, q) \leq K + (1 - \delta)q$ for all $a \in \Omega$, $q \in \mathbb{R}_+$, and $t = 1, \dots, T$.*

Proof of Lemma A3: Define $P_{max} = \max_{a \in \Omega} f(a, 0)$. Using equation (A2) and the fact that $\theta < 1$ (costly storage), an induction argument establishes that $f(a, J(t, a, q) - (1 - \delta)q) \leq P_{max}$ for all dates $t \leq T$, states a , and inventory level q . For each state a , define K_a to solve $f(a, K_a) = P_{max}$. Since f is

increasing without bound in K_a , K_a is finite. Define $K = \max_{a \in \Omega} K_a$. Using equation (A2), $J(t, a, 0) \leq K$ (i.e., it is never an equilibrium to accumulate inventory such that the buying price exceeds the maximum possible selling price). Finally, by equation (A1), $J(t, a, q) \leq K + (1 - \delta)q$. Q.E.D.

LEMMA A4: For all $a \in \Omega$ and $q \in \mathbb{R}_+$, $J(a, q) = \lim_{T \rightarrow \infty} J(1, a, q; T)$.

Proof of Lemma A4: Given that we are taking the limit as $T \rightarrow \infty$, fixing $t = 1$ is arbitrary. Since $J(1, a, q; T)$ is nondecreasing in T (Lemma A2) and pointwise bounded (Lemma A3), the limit exists by the monotone convergence theorem. Q.E.D.

Proof of Proposition 1:

- (a) Established in Lemma A4. The limit preserves the relevant properties established in Lemma A1.
- (b) For each $a \in \Omega$, define \bar{Q}_a such that $\bar{Q}_a = J(a, \bar{Q}_a)$. If for some a , $J(a, 0) = 0$, then $\bar{Q}_a = 0$ and by equation (4), $J(a, q) < q$. If for some other a , $J(a, 0) > 0$ then by equation (4), $J(a, q)$ crosses the 45° line once from above. Hence for each a , \bar{Q}_a is unique. Define $Q_{max} = \max_{a \in \Omega} \bar{Q}_a$. By construction, for $q \in [0, Q_{max}]$, $J(a, q) \leq Q_{max}$ and for some a , $J(a, Q_{max}) = Q_{max}$.
- (c) By definition, $P(a, q) = f(a, J(a, q) - (1 - \delta)q)$. Therefore, by equation (4) and the properties of f , the equilibrium spot price is continuous and decreasing in q . This implies that $P_{min} < P(a, q) < P_{max}$ where $P_{min} = \min_{\{a \in \Omega\}} P(a, Q_{max}) > 0$ and $P_{max} = \max_{\{a \in \Omega\}} P(a, 0) < \infty$. Q.E.D.

Proof of Corollary 1.1 (Properties of J): Since $J(a, q)$ is continuous in q and satisfies equation (4), to establish (a) and (b) it is sufficient to show that there exists an $a \in \Omega$ such that $J(a, 0) = 0$. Suppose not, so that $J(a, 0) > 0$ for all $a \in \Omega$. This implies that, for all $a_t \in \Omega$,

$$\begin{aligned}
 f(a_t, J(a_t, 0)) &= \theta E_t [f(a_{t+1}, J(a_{t+1}, J(a_t, 0))) - (1 - \delta)J(a_t, 0)] \\
 f(a_t, 0) &< \theta E_t [f(a_{t+1}, J(a_{t+1}, 0))] \tag{A5} \\
 f(a_t, 0) &< \theta^n E_t [f(a_{t+n}, J(a_{t+n}, 0))].
 \end{aligned}$$

The first line follows since equation (3a) holds when $J(a, q) > 0$. The second line follows from the fact that the left-hand-side is increasing and the right-hand-side is decreasing in current storage (using equation (4)). The third line follows from recursively applying this step n times. Since $\theta < 1$ and $f(a, 0) > 0$, this is a contradiction because the inequality cannot hold for large n . To establish (b), recall that $J(a, q)$ is weakly increasing and continuous in q (see equation (4)). Q.E.D.

Proof of Corollary 1.2 (Regeneration): Part (a): Consider arbitrary a_t and Q_{t-1} . The proof is by contradiction. Suppose that $\text{Prob}(Q_\tau = 0 | a_t, Q_{t-1}) = 0$ for all $\tau = t, \dots, t+n$. This implies that the equilibrium condition (3a) holds for all $\tau = t, \dots, t+n$. Repeatedly substituting equation (3a) and applying the law of iterated expectations yields

$$\begin{aligned} & f(a_t, J(a_t, Q_{t-1}) - (1 - \delta)Q_{t-1}) \\ &= \theta E_t[f(a_{t+1}, J(a_{t+1}, Q_t) - (1 - \delta)Q_t)] \\ &= \theta^n E_t[f(a_{t+n}, J(a_{t+n}, Q_{t+n-1}) - (1 - \delta)Q_{t+n-1})]. \end{aligned} \tag{A6}$$

Since $f(a_t, \Delta Q_t) > 0$ and $\theta < 1$, equation (A6) cannot hold for large n . Therefore, $Q_\tau = 0$ with positive probability for some $\tau \geq t$.

Part (b): Choose an $a_s \in \Omega_s$ such that $J(a_s, q) = 0$ for some $q \leq q_s$. Define the random variable t^* as the first $t > 1$ such that $a_t = a_s$ and $Q_{t-1} \leq q_s$. By construction $Q_{t^*} = 0$.

$$\begin{aligned} & \text{Prob}(Q_T \leq q | a_t, Q_{t-1}) \\ &= \text{Prob}(Q_T \leq q | T \geq t^*, a_t, Q_{t-1}) \text{Prob}(T \geq t^* | a_t, Q_{t-1}) \\ &\quad + \text{Prob}(Q_T \leq q | T < t^*, a_t, Q_{t-1}) \text{Prob}(T < t^* | a_t, Q_{t-1}) \\ &= \text{Prob}(Q_T \leq q | T \geq t^*, a_t) \text{Prob}(T \geq t^* | a_t, Q_{t-1}) \\ &\quad + \text{Prob}(Q_T \leq q | T < t^*, a_t, Q_{t-1}) \text{Prob}(T < t^* | a_t, Q_{t-1}). \end{aligned} \tag{A7}$$

Note that conditional on $T > t^*$, Q_t does not affect the probability distribution of Q_T . By part (a), we can select an a_s , which defines t^* , such that $\lim_{T \rightarrow \infty} \text{Prob}(T \geq t^* | a_t, Q_{t-1}) = 1$ for any a_t and Q_{t-1} . The result for the spot price probabilities follows by a similar construction using the definition of $P(a, q)$.

Part (c): If the Markov demand process has a limiting invariant distribution then the probability distribution in equation (A7) does not depend on a_t for large T . A sufficient condition for the invariant distribution is $\Pi \gg 0$. Q.E.D.

Proof of Proposition 2 (Increased Inventory): Part (a) follows directly from Proposition 1(a). Part (b) follows from $P(a, q)$ decreasing in q (Proposition 1(c)). Since this holds path-by-path, the definition of forward prices in equation (6) implies that $F_{t, t+n}(a, q)$ also decreases in q . Part (c) is a direct consequence of (a). Part (d) follows from the nonzero stockout probability in Corollary 1.2 (renewal property). Once $Q_\tau^x = 0$ (which implies $Q_\tau = 0$), then $Q_t^x = Q_t$ for all $t \geq \tau$. Q.E.D.

Proof of Corollary 2.1 (Derivative Prices): The proof follows immediately from Proposition 2(b), which shows that forward prices are lower along every path of demand shocks. Q.E.D.

Proof of Proposition 3 (Forward Prices): The limiting forward price, F_∞ , follows immediately from Proposition 1(c) (spot prices are bounded) and, given $\Pi \gg 0$, Corollary 1.2(c) (existence of invariant distribution). Existence of monotone sequences of time and state independent bounds F_{min}^n and F_{max}^n is a consequence of bounded spot prices and the limiting forward price, F_∞ . Q.E.D.

Proof of Proposition 4 (Increased Demand): The proof of part (a) uses the limit of an induction argument in a finite horizon economy as the horizon $T \rightarrow \infty$. The notation is the same as for Lemmas A1–A4. Recall that $J(T, a_\ell, q) = J(T, a_h, q) = 0$; therefore at T prices are ordered, $P(T, a_1, q) < \dots < P(T, a_m, q)$. Assume $P(t + 1, a_1, q) < \dots < P(t + 1, a_m, q)$ at date $t + 1$. Consider two states $a_\ell, a_h \in \Omega$ such that $f(a_\ell, \Delta Q) < f(a_h, \Delta Q)$ as in property (D), Section II.B. If $J(t, a_h, q) \geq J(t, a_\ell, q)$, then $P(t, a_\ell, q) < P(t, a_h, q)$ follows directly. If, instead, $J(t, a_\ell, q) > J(t, a_h, q) \geq 0$, then

$$\begin{aligned} P(t, a_\ell, q) &= \theta E[P(t + 1, a, J(t, a_\ell, q)) | a_\ell] \\ &< \theta E[P(t + 1, a, J(t, a_\ell, q)) | a_h] \tag{A8} \\ &< \theta E[P(t + 1, a, J(t, a_h, q)) | a_h] \leq P(t, a_h, q). \end{aligned}$$

The first line is from $J(t, a_\ell, q) > 0$ which implies that equation (3a) holds. The first inequality holds by property (E) (first-order stochastic dominance). The second inequality holds from the maintained assumption that $J(t, a_\ell, q) > J(t, a_h, q) \geq 0$ and Proposition 1(c). The final line is equation (3) again. This ordering is preserved in the infinite horizon economy.

The proof of (b) for forward prices can be done directly for the infinite horizon economy using (a). If $J(a_\ell, q) \geq J(a_h, q)$ then

$$\begin{aligned} F_{t,\tau}(a_\ell, q) &= E[P(a_\tau, Q_\tau) | a_\ell, Q_t = J(a_\ell, q)] \\ &\leq E[P(a_\tau, Q_\tau) | a_\ell, Q_t = J(a_h, q)] \tag{A9} \\ &< E[P(a_\tau, Q_\tau) | a_h, Q_t = J(a_h, q)] = F_{t,\tau}(a_h, q). \end{aligned}$$

The first inequality follows from the result that prices are decreasing in initial storage (Proposition 1(c)). The second inequality is a consequence of the first-order stochastic dominance assumption (E). If, on the other hand, $J(a_h, q) > J(a_\ell, q) \geq 0$, the proof is more complicated as the effect on prices of the first-order stochastic dominance is offset by the higher future inventory. This case, therefore, is handled in the following lemma which completes the proof. Q.E.D.

LEMMA A5: In an economy satisfying (D) and (E) for any two states $a_\ell, a_h \in \Omega$ such that $f(a_\ell, \Delta Q) < f(a_h, \Delta Q)$ and $J(a_h, q) > J(a_\ell, q) \geq 0$, their forward prices are ordered as $F_{t,\tau}(a_\ell, q) < F_{t,\tau}(a_h, q)$ for all $\tau > t$.

Proof of Lemma A5: We first establish the result for the one-period horizon forward contract. If $J(a_h, q) > J(a_\ell, q) \geq 0$ then by the equilibrium equation (3),

$$P(a_h, q) = \theta E[P(a, J(a_h, q)) | a_h] = \theta F_{t,t+1}(a_h, q) \quad (\text{A10})$$

$$P(a_\ell, q) \geq \theta E[P(a, J(a_\ell, q)) | a_\ell] = \theta F_{t,t+1}(a_\ell, q).$$

Since $P(a_h, q) > P(a_\ell, q)$ by part (a), $F_{t,t+1}(a_h, q) > F_{t,t+1}(a_\ell, q)$. Next, we use induction to show that $F_{t,t+n}(a_h, q) > F_{t,t+n}(a_\ell, q)$ extends from $n = 1$ to longer horizons $n > 1$. Suppose that $F_{t,t+n-1}(a_h, q) > F_{t,t+n-1}(a_\ell, q)$.

- Define $P_{t+n-1}^h = P(a_{t+n-1}, Q_{t+n-2}^h)$ as the spot price at time $t + n - 1$ given the previous inventory, Q_{t+n-2}^h , on a (random) path of demand shocks, $\{a_{t+1}, \dots, a_{t+n-1}\}$ given $a_t = a_h$ at time t . Analogously, let P_{t+n}^h be the spot price one period later at $t + n$.
- Let E_t^h denote the conditional expectation at t given the transition probabilities on the future shocks a_τ given $a_t = a_h$. Let E_{t+n-1} denote the conditional expectation at time $t + n - 1$ given a particular a_{t+n-1} . Since the demand states are Markovian, given a_{t+n-1} , the earlier shock a_t does not affect the expectations at time $t + n - 1$.
- Define the random variable G_{t+n-1}^h for each path, $\{a_{t+1}, \dots, a_{t+n-1}\}$ after $a_t = a_h$ as:

$$G_{t+n-1}^h = \frac{E_{t+n-1}[P_{t+n}^h] - P_{t+n-1}^h}{P_{t+n-1}^h}. \quad (\text{A11})$$

- For $a_t = a_\ell$, define similar terms analogously.

Since the forward price is the expected future spot price, the forward prices in the a_h and a_ℓ states can be expressed in terms of the G_{t+n-1}^h and G_{t+n-1}^ℓ as follows:

$$F_{t,t+n}(a_h, q) = F_{t,t+n-1}(a_h, q) + E_t^h[P_{t+n-1}^h G_{t+n-1}^h] \quad (\text{A12})$$

$$F_{t,t+n}(a_\ell, q) = F_{t,t+n-1}(a_\ell, q) + E_t^\ell[P_{t+n-1}^\ell G_{t+n-1}^\ell].$$

Recall that $E_t^h[P_{t+n-1}^h] = F_{t,t+n-1}(a_h, q) > F_{t,t+n-1}(a_\ell, q) = E_t^\ell[P_{t+n-1}^\ell]$ (by the induction assumption). Therefore, a sufficient condition for $F_{t,t+n}(a_h, q) > F_{t,t+n}(a_\ell, q)$ is $G_{t+n-1}^h \geq G_{t+n-1}^\ell$ for each possible post- t path, $\{a_{t+1}, \dots, a_{t+n-1}\}$. Given a particular post- t path, the two values G_{t+n-1}^h and G_{t+n-1}^ℓ may differ

only because of different inventories at time $t + n - 1$, Q_{t+n-1}^h and Q_{t+n-1}^ℓ . In particular, since $Q_t^h = J(a_h, q) > J(a_\ell, q) = Q_t^\ell$ (i.e., the maintained assumption in this lemma), Proposition 2 says $Q_{t+n-1}^h \geq Q_{t+n-1}^\ell$.

There are two cases to consider to show $G_{t+n-1}^h \geq G_{t+n-1}^\ell$. In the first case, if $Q_{t+n-1}^h > 0$ and $Q_{t+n-1}^\ell \geq 0$, then

$$\begin{aligned} G_{t+n-1}^h &= \frac{E_{t+n-1}[P_{t+n}^h] - P_{t+n-1}^h}{P_{t+n-1}^h} \\ &= \frac{1}{\theta} - 1 \geq \frac{E_{t+n-1}[P_{t+n}^\ell] - P_{t+n-1}^\ell}{P_{t+n-1}^\ell} = G_{t+n-1}^\ell. \end{aligned} \tag{A13}$$

This follows from the equilibrium condition (3) where $Q_{t+n-1}^h > 0$ and $Q_{t+n-1}^\ell \geq 0$. In the second case, if $Q_{t+n-1}^h = Q_{t+n-1}^\ell = 0$ then

$$\begin{aligned} G_{t+n-1}^h &= \frac{E_{t+n-1}[P(a_{t+n}, 0)] - P_{t+n-1}^h}{P_{t+n-1}^h} \\ &\geq \frac{E_{t+n-1}[P(a_{t+n}, 0)] - P_{t+n-1}^\ell}{P_{t+n-1}^\ell} = G_{t+n-1}^\ell. \end{aligned} \tag{A14}$$

The inequality follows from two observations. First, conditional on a stockout at $t + n - 1$ ($Q_{t+n-1}^h = Q_{t+n-1}^\ell = 0$), the expected spot price for date $t + n$ depends only the realization for a_{t+n} whose conditional distribution depends only on a_{t+n-1} (and not on the demand state a_t at t). Second, the date $t + n - 1$ prices are ordered $P_{t+n-1}^\ell \geq P_{t+n-1}^h$ by Proposition 1(c) since $Q_{t+n-2}^h \geq Q_{t+n-2}^\ell$. Q.E.D.

Proof of Proposition 5 (Spanning): Consider a generic state (t, a_t, Q_{t-1}) . The period $t + 1$ payoff on a one-period forward contract is $P(a_{t+1}, Q_t) - F_{t,t+1}(a_t, Q_{t-1})$, where $Q_t = J(a_t, Q_{t-1})$ and a_{t+1} takes one of only two possible values, say $a_\ell, a_h \in \Omega$ at $t + 1$, for each realization $a_t \in \Omega$ at t . Assumptions (D) and (E) then imply that at $t + 1$, $P(a_h, Q_t) > P(a_\ell, Q_t)$ by Proposition 4(a). Since the payoffs on a one-period riskless bond are constant, the payoffs on a one-period bond and a one-period forward contract are linearly independent. Therefore one-period uncertainty is spanned for every (t, a_t, Q_{t-1}) and the market is dynamically complete (see Harrison and Kreps (1979)). Q.E.D.

Proof of Proposition 6 (Forward Slopes): $F'_{t,\tau} \leq (r + \delta)/(1 - \delta)$ follows from equilibrium condition (3) and the law of iterated expectations.

$$\begin{aligned} F_{t,\tau+1}(a_t, Q_{t-1}) - F_{t,\tau}(a_t, Q_{t-1}) &= E[P(a_{\tau+1}, Q_\tau) - P(a_\tau, Q_{\tau-1}) | a_t, Q_{t-1}] \\ &= E[E[P(a_{\tau+1}, Q_\tau) | a_\tau, Q_{\tau-1}] \\ &\quad - P(a_\tau, Q_{\tau-1}) | a_t, Q_{t-1}] \\ &\leq E[(\theta^{-1} - 1)P(a_\tau, Q_{\tau-1}) | a_t, Q_{t-1}] \\ &= \frac{r + \delta}{1 - \delta} F_{t,\tau}(a_t, Q_{t-1}). \end{aligned} \tag{A15}$$

The inequality is strict if $\text{Prob}(Q_\tau = 0) > 0$. Part (a) follows from $\Pi \gg 0$ which implies that $\pi(a_s|a) > 0$ for $a_s \in \Omega$, hence $\text{Prob}(Q_{\tau+1} = 0|Q_\tau = 0) > 0$. For part (b) define a_H such that $P(a_H, 0) = P_{max}$. Note that $J(a_H, 0) = 0$. Since $\Pi \gg 0$, $F_{\tau, \tau+1}(a_H, 0) = E_\tau[P(a_{\tau+1}, 0)] < P(a_H, 0)$. Therefore the demand state and previous inventory pair $(a_H, 0)$ is sufficient for a negative slope (backwardation). Since the probability of a stockout is nonzero and $\pi(a_H|a) > 0$ for all a , the probability of $(a_H, 0)$ at any $\tau > \tau^*$ is nonzero. Finally, (c) follows directly from the limiting forward price in Proposition 3. Q.E.D.

Proof of Proposition 7 (Convenience Yield): Parts (a) and (b) follow by rearranging equation (A15) where $\tau = t + n$:

$$\frac{F_{t, t+n+1}}{F_{t, t+n}} \leq \frac{(1+r)}{(1-\delta)}. \quad (\text{A16})$$

As in Proposition 6, the inequality is strict only if $\text{Prob}(Q_{t+n} = 0) > 0$. Part (c)(i) follows from $\Pi \gg 0$ which implies that $\text{Prob}(Q_{t+n+1} = 0|Q_{t+n} = 0) > 0$. Finally part (c)(ii) follows from equation (10) and the Proposition 6 implication that $\lim_{n \rightarrow \infty} (F_{t, t+n+1}/F_{t, t+n}) = 1$. Q.E.D.

Proof of Proposition 8 (Volatility): The conditional forward-price volatility must decline in maturity for forwards of sufficiently long maturity due to the collapsing bounds, $\{F_{min}^n\}$ and $\{F_{max}^n\}$, and the limiting forward price, F_∞ (from Proposition 3). To construct the violation of the Samuelson effect in (b), consider a and q such that $J(a, q) > 0$. By the equilibrium condition (3a) and the definition of forward prices in equation (6), there exists an $N \geq 1$ such that for all $n \leq N$, $F_{t, t+n}(a, q) = \theta^{-n}P(a, q)$. Since $\theta < 1$, this implies that the variance of $F_{t, t+n}(a, q)$ is larger than the variance of the spot price and, hence, increasing in n . Q.E.D.

Proof of Corollary 8.1 (Hedge Ratios): The proof follows immediately from Propositions 4 and 8. Q.E.D.

Proof of Proposition 9 (Augmentation): Under the conditions $b_t > 0$ and $E_t[b_{t+1}f(a_{t+1}, \Delta Q)] = b_t E_t[f(a_{t+1}, \Delta Q)]$, the equilibrium condition (3) in the two economies is identical. Therefore, $J^* = J$. Q.E.D.

REFERENCES

- Amin, Kaushik, Victor K. Ng, and S. Craig Pirrong, 1995, Valuing energy derivatives, in Robert Jameson, ed.: *Managing Energy Price Risk* (Risk Publications, London).
- Anderson, Ronald W., and Jean-Pierre Danthine, 1983, The time pattern of hedging and the volatility of futures prices, *Review of Economic Studies* 50, 249–266.
- Backus, David K., and Stanley E. Zin, 1994, Reverse engineering the yield curve, Working paper #W4676, NBER.
- Breeden, Douglas T., and Robert H. Litzenberger, 1978, Prices of state-contingent claims implicit in option prices, *Journal of Business* 51, 621–651.
- Brennan, Michael J., 1991, The price of convenience and the valuation of commodity contingent claims, in Diderik Lund, and Bernt Oksendal, eds.: *Stochastic Models and Option Models* (North Holland, Amsterdam).

- Brennan, Michael J., and Nicholas I. Crew, 1996, Hedging long maturity commodity commitments with short-dated futures contracts, in Michael Dempster and Stanley Pliska, eds.: *Mathematics of Derivative Securities* (Cambridge University Press, Cambridge).
- Bresnahan, Timothy F., and Pablo T. Spiller, 1986, Futures market backwardation under risk neutrality, *Economic Inquiry* 24, 429–441.
- Chambers, Marcus J., and Roy E. Bailey, 1996, A theory of commodity price fluctuations, *Journal of Political Economy* 104, 924–957.
- Cox, John C., Jonathan E. Ingersoll, Jr., and Stephen A. Ross, 1981, The relation between forward prices and futures prices, *Journal of Financial Economics* 9, 321–346.
- Cox, John C., Jonathan E. Ingersoll, Jr., and Stephen A. Ross, 1985, A theory of the term structure of interest rates, *Econometrica* 53, 385–407.
- Culp, Christopher L., and Merton H. Miller, 1995, Metallgesellschaft and the economics of synthetic storage, *Journal of Applied Corporate Finance* 7, 62–76.
- Deaton, Angus, and Guy Laroque, 1992, On the behaviour of commodity prices, *Review of Economic Studies* 59, 1–23.
- Deaton, Angus, and Guy Laroque, 1996, Competitive storage and commodity price dynamics, *Journal of Political Economy* 104, 896–923.
- Duffie, Darrell, and Stephen Gray, 1995, Volatility in energy prices, in Robert Jameson, ed.: *Managing Energy Price Risk* (Risk Publications, London).
- Edwards, Franklin R., and Michael S. Canter, 1995, The collapse of Metallgesellschaft: Unhedgeable risks, poor hedging, or just bad luck, *Journal of Futures Markets* 15, 211–264.
- Fama, Eugene F., and Kenneth R. French, 1988, Business cycles and the behavior of metals prices, *Journal of Finance* 43, 1075–1093.
- Gibson, Rajna, and Eduardo S. Schwartz, 1990, Stochastic convenience yield and the pricing of oil contingent claims, *Journal of Finance* 45, 959–976.
- Harrison, J. Michael, and David M. Kreps, 1979, Martingales and arbitrage in multiperiod securities markets, *Journal of Economic Theory* 20, 381–408.
- Heath, David, Robert Jarrow, and Andrew Morton, 1992, Bond pricing and the term structure of interest rates: A new methodology, *Econometrica* 60, 77–105.
- Heinkel, Robert, Maureen E. Howe, and John S. Hughes, 1990, Commodity convenience yields as an option profit, *Journal of Futures Markets* 10, 519–533.
- Hong, Harrison, 2000, A model of returns and trading in futures markets, *Journal of Finance* 55, 959–988.
- Hull, John C., 1997, *Options, Futures and Other Derivatives*, 3rd Ed. (Prentice-Hall, Upper Saddle River, N.J.).
- Jagannathan, Ravi, 1985, An investigation of commodity futures prices using the consumption-based intertemporal capital asset pricing model, *Journal of Finance* 60, 175–191.
- Jaillet, Patrick, Ehud I. Ronn, and Stathis Tompaidis, 1997, Valuation of commodity-based “swing” options, Working paper, University of Texas at Austin.
- Kaldor, Nicholas, 1939, Speculation and economic stability, *Review of Economic Studies* 7, 1–27.
- Karlin, Samuel, and Howard M. Taylor, 1975, *A First Course in Stochastic Processes* (Academic Press, New York).
- Keynes, John M., 1930, *A Treatise on Money* (Harcourt Brace, New York).
- Litzenberger, Robert H., and Nir Rabinowitz, 1995, Backwardation in oil futures markets: Theory and empirical evidence, *Journal of Finance* 50, 1517–1545.
- Mello, Antonio S., and John E. Parsons, 1995, The maturity structure of a hedge matters: Lessons from the Metallgesellschaft debacle, *Journal of Applied Corporate Finance* 8, 106–120.
- Miltersen, Kristian R., and Eduardo S. Schwartz, 1998, Pricing of options on commodity futures with stochastic term structures of convenience yields and interest rates, *Journal of Financial and Quantitative Analysis* 33, 33–59.
- Neuberger, Anthony J., 1999, Hedging long-term exposures with multiple short-term futures contracts, *Review of Financial Studies* 12, 429–459.
- Ng, Victor K., and S. Craig Pirrong, 1994, Fundamentals and volatility: Storage, spreads and the dynamics of metals prices, *Journal of Business* 67, 203–230.

- Pilipovic, Dragna, and J. Wengler, 1998, Getting into the swing, *Energy & Power Risk Management* 2, 22–24.
- Pirrong, S. Craig, 1998, Price dynamics and derivative prices for continuously produced, storable commodities, Working paper, Washington University.
- Raab, Martin, and Robert Schwager, 1993, Spanning with short-selling restrictions, *Journal of Finance* 68, 791–793.
- Richard, Scott F., and M. Sundaresan, 1981, A continuous time equilibrium model of forward prices and futures prices in a multigood economy, *Journal of Financial Economics* 9, 347–371.
- Ronn, Ehud I., and Chang N. Xuan, 1997, Hedging long-dated oil futures contracts: An empirical investigation, Working paper, University of Texas at Austin.
- Ross, Stephen A., 1976, Options and efficiency, *Quarterly Journal of Economics* 90, 75–89.
- Ross, Stephen A., 1997, Hedging long run commitments: Exercises in incomplete market pricing, *Economic Notes by Banca Monte* 26, 99–132.
- Routledge, Bryan R., Duane J. Seppi, and Chester S. Spatt, 1999, The “spark spread:” An equilibrium model of cross-commodity price relationships in electricity, Working paper, Carnegie Mellon University.
- Samuelson, Paul A., 1965, Proof that properly anticipated prices fluctuate randomly, *Industrial Management Review* 6, 41–49.
- Schwartz, Eduardo S., 1997, The stochastic behavior of commodity prices: Implications for valuation and hedging, *Journal of Finance* 52, 923–973.
- Schwartz, Eduardo S. and James E. Smith, 2000, Short-term variations and long-term dynamics in commodity prices, *Management Science*, forthcoming.
- Tauchen, George, and Robert Hussey, 1991, Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models, *Econometrica* 59, 371–396.
- Telser, Lester G., 1958, Futures trading and the storage of cotton and wheat, *Journal of Political Economy* 66, 133–144.
- Williams, Jeffrey C., and Brian D. Wright, 1991, *Storage and Commodity Markets* (Cambridge University Press, Cambridge, England).
- Working, Holbrook, 1948, Theory of the inverse carrying charge in futures markets, *Journal of Farm Economics* 30, 1–28.
- Working, Holbrook, 1949, The theory of the price of storage, *American Economic Review* 39, 1254–1262.
- Wright, Brian D., and Jeffrey C. Williams, 1989, A theory of negative prices for storage, *Journal of Futures Markets* 9, 1–13.