Commodity Price Behavior With Storage Frictions

Lewis Evans

Victoria University of Wellington

Graeme Guthrie*

Victoria University of Wellington

April 18, 2008

*Corresponding author: School of Economics and Finance, PO Box 600, Victoria University of Wellington, Wellington, New Zealand. Ph: 64-4-4635763. Fax: 64-4-4635014. Email: graeme.guthrie@vuw.ac.nz

Commodity Price Behavior With Storage Frictions

Abstract

We present a competitive storage model of commodity prices featuring frictions that introduce an element of irreversibility into storage decisions. This leads to situations in which speculators do not trade in the spot market even though total storage is positive. As a result the market value of the stored commodity, which is determined in the (financial) market for ownership of firms operating storage facilities, can diverge from the spot price. Such price separation leads to the existence of an endogenous convenience yield, which we show equals the expected excess return on a real option embedded in each unit of the stored commodity. The outputs of our model are consistent with the stylized facts regarding commodity price distributions, including serial correlation and GARCH characteristics. Samuelson's hypothesis — that forward prices are less volatile than spot prices — does not hold in general.

JEL Classification code: G1, G13, Q1, Q4.

Keywords: backwardation, commodities, convenience yield, GARCH, real options, Samuelson's hypothesis, storage.

Commodity Price Behavior With Storage Frictions

1 Introduction

There is now an extensive empirical literature that documents the properties of commodity spot and futures prices. In contrast, the fundamental economic origins of the behavior that this literature has identified are less well understood. What insights there are have typically been derived from a competitive storage model in which speculators buy a commodity from producers and store it for subsequent sale to consumers. While the early part of the competitive storage literature concentrates on establishing the existence and uniqueness of rational expectations equilibria (Muth, 1961; Samuelson, 1971; Kohn, 1978; Wright and Williams, 1982, 1984; Scheinkman and Schechtman, 1983), more recent work has shown that the empirical performance of the model leaves much to be desired (Deaton and Laroque, 1992, 1996; Chambers and Bailey, 1996; Ng, 1996; Michaelides and Ng, 2000; Routledge et al., 2000). In short, the standard competitive storage model has trouble explaining much of the behavior documented in the empirical literature. In this paper we introduce a simple friction into the storage process — a cost is incurred each time a unit of the commodity is moved into or out of storage — and show that this leads to quite dramatically altered predictions along a number of dimensions.

The storage friction alters the equilibrium behavior of speculators because it introduces an element of irreversibility into the decision to change storage levels — it is costly to move the commodity into and then immediately out of storage (and also to empty and then immediately refill a storage facility). Since storage changes are no longer costlessly reversible, in some states of nature the spot price is simultaneously too low for the payoff from selling from storage to compensate speculators for the destruction of the real option to wait and sell at a future date and too high for the payoff from raising storage to compensate speculators for the destruction of the real option to wait and buy at a future date.[1] As a result, there is a region in which speculators do not trade in the spot market even though total storage is positive. The existence of this region is the source of all of our results.[2]

Firstly, it leads to GARCH-like behavior in commodity prices. As long as speculators are trading in the spot market, they will help dampen shocks (in our model, supply shocks). However, when speculators withdraw, the spot price will bear the full impact of such shocks. This shows up in our model as periods of relatively low price volatility while speculators are trading

[1] Several authors have studied the call option associated with holding a commodity in storage, but usually in very restrictive ways. For example, Heinkel et al. (1990) analyze the call option associated with holding a commodity in storage, while Litzenberger and Rabinowitz (1995) build an equilibrium model of oil extraction, but in both papers the decision to release the commodity from storage is completely irreversible.

[2] Since the existence of this region is an immediate result of our assumption that changes of inventory are not costlessly reversible, our model could be recast in terms of a wide variety of different frictions and our main results would still hold. For example, if moving the commodity into or out of storage takes time then this delay makes immediate reversal impossible, which is sufficient for the real option to delay changing storage to have value.

in the spot market, interrupted by relatively brief periods of much higher volatility when they are not trading in the spot market. Such behavior is suggestive of the GARCH and similar models that are frequently fitted to commodity price data (Baillie and Myers, 1991; Ng and Pirrong, 1994, 1996). Indeed, our simulations show that our model is able to mimic the results of fitting GARCH models to actual commodity price data, provided that the friction is present. In contrast, when the friction is omitted from the model, so that the no-trade region does not arise, simulated prices exhibit no GARCH-like behavior.

Secondly, in the region where speculators withdraw from the spot market, the market value of the stored commodity diverges from the spot price. Indeed, there are really two distinct goods appearing in our model: the commodity being traded in the spot market and the commodity being held in storage. Their respective prices are the spot price and the market value of a firm holding one unit of the commodity in storage.[3] Equilibrium in the market for ownership of such firms imposes intertemporal restrictions on the price of the stored commodity, and the spot price inherits these intertemporal restrictions whenever storage operators are active in the spot market. However, whenever storage operators do not participate in the spot market, the items traded there are created and consumed instantaneously. As a result, these items cannot be used to transfer wealth across time, so there are no intertemporal restrictions on the spot price. However, there remains a price for the stored commodity in the market value for storage firms. The difference between the two prices reflects the value of the real option to sell the stored commodity in the future. When storage decisions are costlessly reversible, this real option has zero value, so that the two prices are always equal.

The friction-induced price divergence leads to the existence of an endogenous convenience yield. A convenience yield is usually motivated as a flow of nonpecuniary benefits accruing to the owner of the stored commodity but not the owner of a contract for future delivery of the commodity (Kaldor, 1939; Working, 1948, 1949). It is suggested that this flow compensates for an inadequate expected capital gain and thereby induces firms to hold the commodity even though the expected net return from doing so appears to be negative. A variety of sources have been suggested for the benefits underlying the convenience yield. For example, Telser (1958) suggests that inventory allows producers to reduce the costs of producing at a given level and of varying output. Brennan (1958) suggests sources for the convenience yield that include avoiding costs associated with frequent deliveries of inputs, avoiding delays and costs of varying the production schedule to meet fluctuating demand, and the ability to take advantage of price changes at short notice.

Of course, because we explicitly model the costs and benefits of storing the commodity in this paper, there should be no convenience yield. And, in fact, as long as the market value of the stored commodity is used to calculate the expected return from storage, there is no expected return shortfall. That is, when returns and opportunity costs are correctly measured, there is

[3]This follows from our assumption of free entry into the storage market.

no convenience yield in this model. However, when we use the widely-accepted approach of calculating the expected return from storage using the spot price of the commodity, there may be an apparent shortfall, implying a nonzero convenience yield. In our model, this convenience yield is zero when speculators are active in the spot market (and the spot price inherits the intertemporal properties of the market value of the stored commodity). Outside this region, the convenience yield is nonzero. Therefore, at least in our model, the convenience yield is simply a consequence of mis-calculating the return to storage by using the spot price rather than the true market value of the stored commodity.

Since the convenience yield arises due to the divergence between the spot price and the market value of the stored commodity, and since this divergence is due to the presence of a real option in each unit of the stored commodity, it follows that the convenience yield somehow reflects the existence of this real option. Although the convenience yield is commonly defined to be the flow of nonpecuniary benefits that accrue to the owner of physical inventory, the usual explanations for the source of the convenience yield ultimately originate from various cash flows. Some (for example, reducing the cost of producing at a given level or of varying output to meet fluctuating demand) involve the avoidance of future cash outflows. Others (for example, the ability to take advantage of price changes at short notice) involve future cash inflows that result from ownership of a commodity. At least in our model, these cash flows can be thought of as the payoff from exercising the real option embedded in the stored commodity. The expected return shortfall, and hence the convenience yield, arises simply because the usual expected return calculation ignores the value of this embedded option. Indeed, we show that the convenience yield equals the amount by which the capital gain on the embedded real option is expected to exceed the opportunity cost of ownership.

When the friction is absent, the instantaneous convenience yield is zero except during stock-outs. However, when the friction is included in the model the convenience yield can take nonzero values for positive storage levels. When this happens the convenience yield will be positive for low harvest levels and negative for high harvest levels, reflecting the mean-reverting nature of the harvest in our model. For example, if the current harvest is higher than the long-run average, the harvest is expected to fall in the future. In the region where speculators have withdrawn from the spot market, this implies that the spot price is expected to rise in the future, so that the expected return from storage will be relatively high; if the harvest is sufficiently high, the implied convenience yield will be negative. The possibility that the convenience yield can take negative values has two important implications for the behavior of commodity spot and forward prices.

The first involves the relationship between spot and forward prices. This relationship is summarized by the extent of backwardation, which we define as the amount by which the spot price exceeds the forward price, adjusted for interest and storage costs.[4] The resulting measure,

[4] Alternative measures include the unadjusted spread between spot and forward prices (essentially just the

which actually equals the present value of the flow of the convenience yield prior to the forward contract's delivery date,[5] gives the amount by which the payoff from selling the stored commodity immediately exceeds that from selling the commodity at a fixed date in the future. With negative storage allowed and frictions omitted, this measure would always equal zero in equilibrium. With negative storage impossible, when storage is low firms hold inventory at a level such that our backwardation measure is positive; that is, the marginal storage operator is holding inventory even though the return from selling immediately is greater than the (apparent) return from continuing to hold the commodity in storage. When a storage friction is added, for all but very low storage levels firms hold inventory at a level such that our measure of backwardation is negative; that is, the marginal storage operator is holding inventory when the (apparent) return from adding even more of the commodity to storage is positive. Nevertheless, for both positive and negative values of our backwardation measure the marginal storage operator is behaving rationally. In the first case, the best alternative to selling immediately is not necessarily to commit to sell a fixed number of years in the future, but rather to retain the option to sell at any time; the value of this option is omitted from the backwardation calculation. In the second case, the best alternative to buying immediately is not necessarily to commit to buy the commodity a fixed number of years in the future, but rather to retain the option to buy at any time (or not at all); in this case the extent of backwardation overstates the cost of the alternative to buying the commodity immediately.

Our prediction that the apparent return from storage will often be negative in equilibrium is not new.[6] Indeed, an extensive literature documents, and attempts to explain, evidence that firms often hold considerable amounts of inventory even though markets are backwardated. The earliest explanation relies on the existence of a flow of nonpecuniary benefits accruing to the owner of the stored commodity but not the owner of a contract for future delivery of the commodity (that is, a convenience yield), which compensates for an inadequate expected capital gain and thereby induces firms to hold the commodity during backwardations (Kaldor, 1939; Working, 1948, 1949). Unfortunately, relatively little progress has been made in developing the theoretical foundations of the convenience yield. The most popular alternative explanation is that ownership of the stored commodity confers a valuable timing option that can be exercised during stock-outs (Routledge et al., 2000).[7] One contribution of this paper is to integrate these

slope of the term structure of forward prices) and the amount by which the spot price exceeds the forward price adjusted for interest costs only (Litzenberger and Rabinowitz, 1995).

[5]Since the convenience yield can be nonzero when inventory is positive, the extent of backwardation does not depend solely on stock-outs as it does in the frictionless model.

[6]Our other prediction — that the apparent return from storage will often be positive in equilibrium — does seem to be new and is worthy of further investigation.

[7]Along similar lines, Wright and Williams (1989) argue that commodities that are aggregated for reporting purposes are often economically distinct. They show that if the cost of transforming one commodity into another is different when carried out in a later period, then one commodity may be stored in positive quantities even though (apparent) excess returns are available from storing the other commodity. Such a situation will show

two approaches, showing how the convenience yield can be interpreted as the amount by which the expected return from holding a timing option embedded in each unit of the stored commodity exceeds the opportunity cost.

Secondly, although we find that the Samuelson (1965) hypothesis that forward prices are less volatile than spot prices holds in the frictionless case, it does not hold in general when a storage friction is added to the model. In fact, in this latter case, Samuelson's hypothesis is violated during moderately good harvests (while the introduction of a storage friction actually strengthens Samuelson's hypothesis for moderately poor harvests).[8] For moderately high harvest levels, an increase in the harvest lowers the spot price, and moves the system further away from the region where the convenience yield is negative. Because the spot price is expected to grow rapidly in this region, the higher harvest level makes imminent rapid increases in the spot price less likely than before. Thus, the forward price falls even further than the spot price.[9]

Although they do not investigate the role of storage, three recent papers also use continuous-time equilibrium models of spot prices to analyze the behavior of forward prices. Carlson et al. (2006) analyze an exhaustible resource and find, amongst other results, that there is a U-shaped relationship between spot price volatility and the slope of the term structure of forward prices. Kogan et al. (2005) obtain a similar result for futures price volatility in a model that features irreversible investment and a capacity constraint, a result that they claim cannot be captured by standard storage-based models of commodity prices. Casassus et al. (2005) build an equilibrium model involving a commodity that is used as an input into a production process. There is no storage in their model, but a friction in the commodity extraction process is sufficient to induce an endogenous convenience yield.

We describe our model set-up in Section 2. Then, in Section 3, we show how a competitive equilibrium spot price can be derived. We also describe some key properties of equilibrium storage policies and explain why the market value of the stored commodity deviates from the spot price. Our discussion of the convenience yield and its origins in the real option embedded in the stored commodity is contained in Section 4, while the next section explains how the frictions in our model affect the extent of backwardation observed in commodity prices. The results of our numerical analysis are presented in Section 6, where we confirm the predictions of the earlier sections as well as describing the high-frequency price dynamics implied by our model. Section 7

---

up in the data as positive industry-wide storage with backwardated prices. Brennan et al. (1997) analyze the wheat marketing system of Western Australia and find behavior consistent with this argument, while Frechette and Fackler (1999) present contradictory empirical evidence from the US corn market.

[8] Existing theoretical explanations for violations of Samuelson's hypothesis include seasonal demand shocks (Anderson and Danthine, 1983) and information asymmetry among investors (Hong, 2000). We achieve the same result with a time-homogeneous information flow and all agents sharing a common information set.

[9] This is consistent with the observation by Bessembinder et al. (1996) that Samuelson's hypothesis is driven by a negative correlation between changes in the spot price and changes in the slope of the term structure of forward prices. For moderately high harvest levels, an increase in the harvest reduces both the spot price and the slope of the term structure of forward prices, so that the correlation is positive.

offers some concluding remarks. All proofs can be found in Appendix A.

2 Basic model structure

There are three types of agents: consumers, producers, and speculators. Speculators purchase the commodity from producers and store it for later sale to consumers. Producers can also sell directly to consumers, but these sales occur immediately after production (that is, only speculators can store the commodity). Let s_t denote the total quantity of the commodity held in storage at date t and suppose that inventories decay at the rate ε, for some constant $\varepsilon > 0$. Some of the commodity is destroyed whenever it is moved into or out of storage. In particular, each unit of the commodity purchased by speculators only increases inventory by $1 - k_i$ units, while each unit sold to consumers reduces inventory by $1 + k_o$ units, where $k_i + k_o > 0$.[10] Thus, total storage evolves according to

$$ds_t = -(\pi(z_t) + \varepsilon s_t)dt, \tag{1}$$

where z_t is the rate at which speculators sell the commodity to consumers[11] and

$$\pi(z) = \begin{cases} (1 - k_i)z, & \text{if } z < 0, \\ (1 + k_o)z, & \text{if } z \geq 0. \end{cases}$$

The market-clearing spot price at date t is $p_t = \psi(y_t + z_t)$, where the consumers' inverse demand function ψ is assumed to satisfy $\psi' < 0$ and $\lim_{q\to\infty} \psi(q) = 0$, and y_t is the rate at which the commodity is produced by producers. We suppose that y_t evolves according to the diffusion process

$$dy_t = \nu(y_t)dt + \phi(y_t)d\xi_t, \tag{2}$$

for some functions ν and ϕ.[12] If risk premia are determined by the Capital Asset Pricing Model, then the appropriate 'risk-neutral process' for y_t is

$$dy_t = (\nu(y_t) - \lambda\phi(y_t))dt + \phi(y_t)d\xi_t, \tag{3}$$

where λ (which we assume to be constant) is the product of the market price of risk (the market risk premium divided by the standard deviation of the rate of return on the market

[10]This particular friction is the specific source of the main results in our model. However, anything that introduces an element of irreversibility into the decision to change storage levels will lead to valuable timing options embedded in the stored commodity, and therefore to qualitatively similar results. For example, the friction could arise from cash outflows incurred whenever moving the commodity into and out of storage, or from short-run inflexibility in processing and transporting commodities. Either friction would generate the embedded options that are key to our results, albeit at the cost of a much more complicated model in the case of the latter alternative.

[11]If $z_t < 0$, then $-z_t$ is the rate at which speculators purchase the commodity from producers.

[12]Seasonality could be introduced into the harvest process by allowing ν and ϕ to depend explicitly on t, although this would complicate the analysis by introducing another dimension (time) into the partial differential equation that we need to solve when analyzing the model.

portfolio) and the correlation coefficient between the rate of return on the market portfolio and the exogenous harvest (Trigeorgis, 1996, pp. 95–101). As long as we use the risk-neutral process, we can calculate present values by discounting expected cash flows using the risk-free interest rate r, which we assume is constant.

3 Competitive equilibrium spot prices

3.1 Socially optimal management of storage

In this section we derive an inventory management policy chosen by a social planner who seeks to maximize the present value of the flow of total surplus, comprising the sum of the surpluses flowing to consumers, producers, and speculators.

Proposition 1 *The flow of total surplus at date t equals*

$$TS(z_t; y_t) = \int_0^{y_t+z_t} \psi(q)\, dq. \tag{4}$$

■

The social planner faces the following problem: given current storage s and the current harvest y, choose the level of sales from storage z in order to maximize the present value of total surplus,

$$W(s,y) = E_0^* \left[\int_0^\infty e^{-rt}\, TS(z(s_t,y_t); y_t)\, dt \mid (s_0,y_0) = (s,y) \right],$$

where E^* denotes expectations with respect to the risk-neutral process in (3). The corresponding Hamilton-Jacobi-Bellman equation is

$$0 = \max_z \ -(\pi(z) + \varepsilon s)\frac{\partial W}{\partial s} + (\nu - \lambda\phi)\frac{\partial W}{\partial y} + \frac{1}{2}\phi^2 \frac{\partial^2 W}{\partial y^2} - rW + TS(z;y). \tag{5}$$

The choice of z in (5) is constrained by the requirements that (i) $y + z \geq 0$ for all (s, y) and (ii) $z \leq 0$ whenever $s = 0$. The first constraint reflects the fact that consumption cannot be negative; the second reflects the fact that storage cannot be negative.

Given the solution for W, the optimal inventory management policy, $z^*(s, y)$, is straightforward to determine.

Proposition 2 *During a stock-out, the social planner will store the entire harvest if $\psi(0) < (1-k_i)\frac{\partial W}{\partial s}$, sell the entire harvest to consumers if $\psi(y) > (1-k_i)\frac{\partial W}{\partial s}$, and otherwise choose $z^* \in [-y, 0]$ defined implicitly by $\psi(y+z^*) = (1-k_i)\frac{\partial W}{\partial s}$. Whenever inventory is positive,*

- *if $\psi(0) < (1-k_i)\frac{\partial W}{\partial s}$ then the social planner will store the entire harvest, so that $z^* = -y$;*
- *if $\psi(y) \leq (1-k_i)\frac{\partial W}{\partial s} \leq \psi(0)$ then the social planner will store some of the harvest, choosing $z^* \in [-y, 0]$ defined implicitly by*

$$\psi(y+z^*) = (1-k_i)\frac{\partial W}{\partial s}; \tag{6}$$

- *if*

$$(1-k_i)\frac{\partial W}{\partial s} \le \psi(y) \le (1+k_o)\frac{\partial W}{\partial s}, \tag{7}$$

  then the social planner holds inventory constant, so that $z^* = 0$*;*

- *if* $\psi(y) \ge (1+k_o)\frac{\partial W}{\partial s}$ *then the social planner will lower storage, choosing* $z^* \ge 0$ *defined implicitly by*

$$\psi(y+z^*) = (1+k_o)\frac{\partial W}{\partial s}. \tag{8}$$

■

Our assumption that $k_i + k_o > 0$ ensures that condition (7) will occur over some nontrivial range of harvest levels. Thus, there will be a range of harvest levels for which it is socially optimal to hold inventory constant. There will be another region, where (8) holds, where it is socially optimal to sell the commodity out of storage. Sales will be set at a level such that the opportunity cost of not consuming one additional unit immediately, $\psi(y+z)$, equals the increase in overall welfare from storing $1+k_o$ additional units, $(1+k_o)\frac{\partial W}{\partial s}$. In the remaining region, where (6) holds, it is socially optimal to purchase the commodity and store it. Purchases will be set at a level such that the opportunity cost of not consuming an additional unit immediately, $\psi(y+z)$, equals the increase in overall welfare from storing $1-k_i$ additional units, $(1-k_i)\frac{\partial W}{\partial s}$.

Given this policy, we should observe discrete intervals between periods when the social planner raises and lowers inventory. That is, it will raise inventory for a while (perhaps including periods when inventory is held constant), and then after some delay will lower inventory. However, the social planner will not immediately switch from raising inventory to lowering it. This delay is required for the state of the system, (s_t, y_t), to move from the inventory-lowering region (where $\psi(y) > (1+k_o)\frac{\partial W}{\partial s}$) to the inventory-raising region (where $\psi(y) < (1-k_i)\frac{\partial W}{\partial s}$).

It follows from (5) that in order to find a surplus-maximizing inventory management policy we need to solve

$$0 = -(\pi(z^*(s,y)) + \varepsilon s)\frac{\partial W}{\partial s} + (\nu - \lambda\phi)\frac{\partial W}{\partial y} + \frac{1}{2}\phi^2\frac{\partial^2 W}{\partial y^2} - rW + \int_y^{y+z^*(s,y)} \psi(q)\, dq \tag{9}$$

for $W(s,y)$, where the function $z^*(s,y)$ is given in Proposition 2. We describe how we solve this problem, which must be done numerically due to its highly nonlinear nature, in Appendix B.

3.2 Competitive equilibrium

We now show that the solution to the social planner's problem described in Section 3.1 corresponds to a competitive equilibrium with spot price equal to

$$p_t = P(s_t, y_t) \equiv \psi(y_t + z^*(s_t, y_t)) \tag{10}$$

at date t.[13] This involves answering two questions. Firstly, what is the behavior of a storage firm that takes this spot price process as given? Secondly, if all storage firms behave in this way,

[13] Deriving a competitive equilibrium by solving an associated 'surplus' maximization problem has a long history in studies of commodity markets (Samuelson, 1971; Scheinkman and Schechtman, 1983).

what would the spot price be? We find that the answer to the second question is the spot price in (10), confirming that this spot price arises in a competitive equilibrium.

Consider a firm that currently holds x units of the commodity in storage and takes the spot price process, $P(s,y)$, and the aggregate sales by speculators, $z^*(s,y)$, as given. The firm chooses its own storage policy in order to maximize its market value. Specifically, the firm chooses its rate of sales from storage, $w(x;s,y)$, in order to maximize

$$G(x;s,y) = E_0^*\left[\int_0^\infty e^{-rt} w(x_t;s_t,y_t)P(s_t,y_t)\,dt \mid (x_0,s_0,y_0) = (x,s,y)\right],$$

where its total inventory evolves according to

$$dx_t = -(\pi(w(x_t;s_t,y_t)) + \varepsilon x_t)dt.$$

The corresponding Hamilton-Jacobi-Bellman equation is

$$0 = \sup_w \; -(\pi(w)+\varepsilon x)\frac{\partial G}{\partial x} - (\pi(z^*)+\varepsilon s)\frac{\partial G}{\partial s} + (\nu - \lambda\phi)\frac{\partial G}{\partial y} + \frac{1}{2}\phi^2\frac{\partial^2 G}{\partial y^2} - rG + wP(s,y).$$

The choice of w is constrained by the requirement that $w \le 0$ whenever $x = 0$; that is, the firm's storage cannot be negative.[14] Solutions to this equation, one of which is described in the following proposition, correspond to optimal storage policies for a price-taking storage operator.

Proposition 3 *An optimal storage policy for a price-taking storage operator with x units of the commodity in storage is to sell from storage at the rate*

$$w(x;s,y) = \frac{x}{s}z^*(s,y). \tag{11}$$

(Negative values of w correspond to spot market purchases.) The market value of such a storage operator equals

$$G(x;s,y) = x\frac{\partial W}{\partial s}. \tag{12}$$

■

Suppose that each storage operator i, who has x_i units of the commodity in storage, follows the storage policy described in Proposition 3. Then total storage is $\sum_i x_i = s$ and the total sales from storage equals

$$\sum_i w_i = \sum_i \frac{x_i}{s}z^*(s,y) = \frac{\sum_i x_i}{s}z^*(s,y) = z^*(s,y).$$

Therefore if individual firms follow the storage policy in Proposition 3, the aggregate behavior matches the evolution of total inventory under the social planner's optimal policy. Moreover, Proposition 3 shows that this behavior is optimal for individual price-taking firms. It follows that the solution to the social planner's problem described in Section 3.1 corresponds to a particular competitive equilibrium as described in the next proposition.

[14] In keeping with standard competitive equilibrium models, we assume that the firm believes it can always buy and sell as much of the commodity as it wishes. In particular, the firm's choice of w is unaffected by the nonnegativity constraints on consumption and aggregate storage.

Proposition 4 *The spot price $P(s,y)$ described in equation (10) is the result of a competitive equilibrium in which the market value of each unit of the stored commodity equals*

$$V(s,y) = \begin{cases} \frac{P(s,y)}{1-k_i}, & \textit{if } z^*(s,y) < 0, \\ \frac{\partial W}{\partial s}, & \textit{if } z^*(s,y) = 0, \\ \frac{P(s,y)}{1+k_o}, & \textit{if } z^*(s,y) > 0, \end{cases}$$

and a storage operator with x units of the commodity in storage sells from storage at the rate

$$w(x;s,y) = \frac{x}{s} z^*(s,y).$$

■

There are two distinct goods in our model — the commodity being traded in the spot market and the commodity being held in storage. Their respective prices are the spot price and the market value of a firm holding one unit of the commodity in storage, and — as Proposition 4 shows — these prices are not equal. However, speculators' ability to transform one good into the other explains the links between the two prices that are evident from Proposition 4. When speculators are selling the commodity in the spot market (that is, when $z^*(s,y) > 0$), the market value of a firm holding one unit of the commodity equals the net proceeds from selling the commodity. When speculators are buying the commodity from producers (that is, when $z^*(s,y) < 0$), the market value of a firm holding one unit of the commodity equals the expenditure that is avoided by already owning the commodity. In the remaining case, where speculators are holding the commodity but not trading in the spot market, the market value of the firm is equal to the marginal value that the social planner attributes to inventory.[15]

The wedge between the market value of the stored commodity and the spot price (adjusted for the cost of moving the commodity out of storage) has a natural interpretation in terms of a real option embedded in each unit of the stored commodity. This option, which gives the firm the ability to choose the time at which the stored commodity (or what is left of it) is sold, is worth

$$U(s,y) \equiv V(s,y) - \frac{P(s,y)}{1+k_o}.$$

Because a firm that currently holds one unit of the commodity in storage has the option to sell this unit immediately, for a payoff of $P(s,y)/(1+k_o)$, it must be the case that $V(s,y) \geq P(s,y)/(1+k_o)$. Thus, the market value of the embedded real option is never negative. Similarly, a firm with no current stock-holding always has the option to buy $1/(1-k_i)$ units of the commodity and move them into storage. This costs $P(s,y)/(1-k_i)$ and results in the firm being worth $V(s,y)$, so that the firm must currently be worth at least $V(s,y) - P(s,y)/(1-k_i)$. Since

[15] Notice that if $k_i + k_o = 0$, so that there are no storage frictions, the market value of one unit of the stored commodity always equals $P(s,y)/(1+k_o)$. We will see that this result plays an important role in explaining why the convenience yield is zero (absent stock-outs) when the storage technology is frictionless.

the firm with zero inventory has market value of zero, we must have $V(s,y) - P(s,y)/(1-k_i) \leq 0$. Thus, the market value of the embedded option must also satisfy

$$U(s,y) = V(s,y) - \frac{P(s,y)}{1+k_o} \leq \frac{P(s,y)}{1-k_i} - \frac{P(s,y)}{1+k_o} = \frac{(k_i+k_o)P(s,y)}{(1-k_i)(1+k_o)}.$$

Since this confirms that the real option is worthless in the frictionless case, it shows that the option value arises because of the element of irreversibility in the storage technology. Not surprisingly, therefore, we will see that this real option embedded in the stored commodity plays an important part in our interpretation of the convenience yield, which is the subject of the next section.

4 The convenience yield: valuation of the stored commodity

In this section we demonstrate how the presence of the friction in our competitive storage model gives rise to an endogenous convenience yield. The usual interpretation of the convenience yield is that it is a flow of nonpecuniary benefits that firms receive as long as they store the commodity and which can be large enough for firms to hold the commodity in storage even though the apparent expected return from doing so is negative.[16] The value of these benefits is such that when they are incorporated into the expected return calculation, the owner of the stored commodity just breaks even. That is, the sum of the expected capital gain and the convenience yield equals the sum of opportunity and storage costs. Equivalently, the convenience yield equals the expected return shortfall.

However, because we explicitly model all of the costs and benefits associated with holding the commodity in storage it follows that there are no such benefits in our model, and therefore no expected return shortfall. Indeed, this is the case as long as the market value of the stored commodity (V in Proposition 4) is used to measure both the opportunity cost and the capital gain.

Proposition 5 *The expected return shortfall from holding the commodity in storage is always equal to zero:*

$$(r+\varepsilon)V - \frac{E^*[dV]}{dt} = 0.$$

■

This is unsurprising, since V is determined in the market for ownership shares in storage operators. Equilibrium in this market, like that in any financial market, will ensure that the owners of storage operators never experience shortfalls or windfalls in the expected rate of return on their investment.

[16]This is often expressed in terms of forward prices, but the arguments apply just as well using the expected future spot price.

Of course, the usual approach to calculating the convenience yield uses the commodity spot price as a proxy for the market value of the stored commodity. In this approach, the commodity is valued at the spot price, so that the convenience yield CY is

$$CY = (r+\varepsilon)P - \frac{E^*[dP]}{dt}. \tag{13}$$

The precise form of this convenience yield is reported in Proposition 6.

Proposition 6 *The convenience yield is*

$$CY(s,y) = \begin{cases} 0, & \text{if } z^*(s,y) < 0, \\ (r+\varepsilon)\psi(y) - (\nu - \lambda\phi)\psi'(y) - \frac{1}{2}\phi^2\psi''(y), & \text{if } z^*(s,y) = 0, \\ 0, & \text{if } z^*(s,y) > 0. \end{cases} \tag{14}$$

■

In our model — and we argue that the same is true more generally — the convenience yield merely reflects an *apparent* expected return shortfall. It arises above solely because the stored commodity is being misvalued when the spot price, rather than the true market value of the stored commodity, is used to calculate the expected return. The correct value is given by the market value of the storage firm.

Proposition 6 shows that the convenience yield is nonzero only when storage operators are not active in the spot market; speculative trade in the commodity (that is, $z^*(s,y) \neq 0$) is sufficient to keep the convenience yield zero. While the misvaluation implicit in the common construction of the convenience yield can explain the existence of a nonzero convenience yield, why is the convenience yield zero whenever $z^*(s,y) \neq 0$? The answer to this question can be found in Proposition 4. When $z^*(s,y) < 0$, speculators are buying the commodity in the spot market and their presence means that the spot price equals $(1-k_i)V$. Thus, the intertemporal restrictions imposed on the market value of the stored commodity in financial markets carry over to the spot market — there is no expected return shortfall in either market. Similarly, when $z^*(s,y) > 0$, speculators are selling the commodity in the spot market, ensuring that the spot price equals $(1+k_o)V$ and thereby inherits the intertemporal restrictions imposed on V in financial markets. Again, there is no expected return shortfall in either market. Only when $z^*(s,y) = 0$, so that speculators are not active in the spot market, does the linkage between P and V break down. Without this linkage, equilibrium in financial markets cannot impose any intertemporal restrictions on the spot price.

The convenience yield has a natural interpretation in terms of the real option embedded in the stored commodity, which is omitted when the spot price is used in place of the correct market value. Since the spot price satisfies $P = (1+k_o)(V-U)$, the convenience yield can be written as

$$\begin{aligned} CY &= \left((r+\varepsilon)P - \frac{E^*[dP]}{dt}\right) \\ &= (1+k_o)\left(\left((r+\varepsilon)V - \frac{E^*[dV]}{dt}\right) - \left((r+\varepsilon)U - \frac{E^*[dU]}{dt}\right)\right). \end{aligned} \tag{15}$$

As shown in Proposition 5, financial markets will ensure that investors cannot earn excess returns by trading in shares of storage operators, which ensures that the first term in brackets on the right hand side of (15) equals zero. It follows that the convenience yield equals

$$CY = (1+k_o)\left(\frac{E^*[dU]}{dt} - (r+\varepsilon)U\right). \tag{16}$$

That is, the convenience yield is actually the expected excess return on the real option embedded in the stored commodity.

Equation (16) gets at the very essence of the convenience yield. It is usually interpreted as the flow of benefits that the storage operator receives while holding the commodity in storage. However, consideration of all of the usual stories that motivate these benefits reveals that they are only received when the commodity is consumed, transformed as part of a production process, or otherwise destroyed — they are not received while the commodity continues to be stored. Put simply, the traditional interpretation of the convenience yield involves the storage operator 'having its cake and eating it too.' In contrast, equation (16) shows that the convenience yield is actually capturing the change in the present value of the timing option embedded in the stored commodity. Different motivations for the convenience yield ultimately reduce to different reasons why timing flexibility might be valuable. In our model the source of the timing option's value is the friction in storage, but it also could arise from frictions in production, for example. Thus, equation (16) can be interpreted as showing that the convenience yield reflects the expected change in the *present value* of those *future* benefits that occurs while the commodity is being held in storage.

5 How frictions affect backwardation

A large part of the literature on commodity markets has focussed on the relationship between spot and forward prices. This relationship is summarized by the extent of backwardation — the amount by which the spot price exceeds the (appropriately discounted) forward price. Standard market practice is to measure the extent of backwardation using the amount by which the spot price exceeds a given forward price (essentially the slope of the term structure of forward prices). In contrast, Litzenberger and Rabinowitz (1995) discount the forward price by an amount reflecting the interest carry charges in their measure of what they call 'weak backwardation'. We go one step further and also adjust the forward price to incorporate the physical depreciation of the commodity that occurs during storage. The resulting measure, which actually equals the present value of the flow of the convenience yield prior to the forward contract's delivery date, offers significant insights into the impact of storage frictions.[17]

Following Routledge et al. (2000), we assume that the introduction of a forward market has no effect on equilibrium in the spot market.[18] We value any forward position simply by calculating the present value of the cash flows it generates. For example, a long position in a forward contract generates a cash flow of $p_T - f$ on its delivery date T, where f is the contracted forward price. By setting the date t market value of this position,

$$e^{-r(T-t)}E_t^*[p_T - f] = e^{-r(T-t)}(E_t^*[p_T] - f),$$

equal to zero, we find that the forward price at date t for delivery at date T is[19]

$$f_{t,T} = E_t^*[p_T].$$

Suppose we currently hold $1 + k_o$ units of the commodity in storage. If we move these units out of storage we receive an immediate cash flow of $P(s, y)$. Suppose, instead, that we hold them in storage for τ years and then sell the units that remain. We can hedge spot price risk by taking $e^{-\varepsilon\tau}$ long forward positions, ensuring that τ years from now we receive $e^{-\varepsilon\tau}F(\tau, s, y)$, where $F(\tau, s, y)$ is the current forward price for delivery τ years in the future. In present value terms, the gain from selling immediately equals

$$B(\tau, s, y) = P(s, y) - e^{-(r+\varepsilon)\tau}F(\tau, s, y). \tag{17}$$

We refer to this as the extent of backwardation in the forward market. An equivalent measure of backwardation is the interest- and storage-adjusted basis[20]

$$I(\tau, s, y) = \frac{1}{\tau}\log\left(\frac{F(\tau, s, y)}{P(s, y)}\right) - r - \varepsilon = \frac{1}{\tau}\log\left(\frac{P(s, y) - B(\tau, s, y)}{P(s, y)}\right). \tag{18}$$

It can be interpreted as the flow of benefits (expressed as a constant proportion of the spot price) that would be required for the return from holding one unit of the commodity in storage

[17] The other backwardation measures can also be expressed as present values of flows. For example, the measure of weak backwardation in Litzenberger and Rabinowitz (1995) equals the present value of the flow of the convenience yield net of ongoing storage costs. However, our measure is derived from the convenience yield alone, and so captures the effects of storage frictions without being contaminated by factors such as ongoing storage costs.

[18] In particular, the introduction of a forward market cannot affect the behavior of consumers, producers, or storage operators. While this assumption is strong, it greatly simplifies the analysis. The (much more complicated) alternative is to explicitly include forward markets in the model, derive demand and supply functions for forward contracts with various delivery dates, and calculate market-clearing prices. For example, Richard and Sundaresan (1981) and Kawai (1983) model simultaneous equilibria in frictionless spot and forward markets.

[19] Note that we use the risk-neutral harvest process to calculate the expectation in this equation. Thus the forward price is not necessarily an unbiased predictor of the future spot price when the actual harvest process is used.

[20] $-I$ is the usual expression for the convenience yield extracted from forward and spot prices in textbook treatments of commodity futures markets (Hull, 2005, p. 117).

for τ years to equal the payoff from selling that unit immediately.[21] The interest- and storage-adjusted basis reveals similar information to B. For example, B and I always have opposite signs, so that the market is backwardated if and only if I is negative.

Cash-and-carry arbitrage is possible if B is too large and negative (or I is too large and positive). If, for example, we own an empty storage facility, we can buy $1 + k_o$ units of the commodity, immediately move $(1 - k_i)(1 + k_o)$ units into storage (the remainder are lost due to the storage friction) and sell the units that remain after τ years have elapsed. We can hedge spot price risk by taking $(1 - k_i)e^{-\varepsilon\tau}$ long forward positions, ensuring that τ years from now we receive $(1 - k_i)e^{-\varepsilon\tau}F(\tau, s, y)$. The immediate profit from this strategy equals

$$(1 - k_i)e^{-(r+\varepsilon)\tau}F(\tau, s, y) - (1 + k_o)P(s, y).$$

If cash-and-carry arbitrage is to be eliminated (which is a necessary condition for equilibrium) then this profit must be less than or equal to zero. Equivalently, the forward price must satisfy

$$F(\tau, s, y) \leq \left(\frac{1 + k_o}{1 - k_i}\right) e^{(r+\varepsilon)\tau} P(s, y),$$

so that the extent of backwardation satisfies

$$B(\tau, s, y) \geq -\left(\frac{k_i + k_o}{1 - k_i}\right) P(s, y)$$

and the interest- and storage-adjusted basis satisfies

$$I(\tau, s, y) \leq \frac{1}{\tau} \log\left(\frac{1 + k_o}{1 - k_i}\right).$$

In particular, in the frictionless case (in which $k_i + k_o = 0$), the extent of backwardation can never be negative and the interest- and storage-adjusted basis can never be positive.

The following proposition shows that $B(\tau, s, y)$ equals the present value of the flow of the convenience yield in (14).

Proposition 7 *The extent of backwardation and the convenience yield are related according to*

$$B(\tau, s, y) = E_0^*\left[\int_0^\tau e^{-(r+\varepsilon)t} CY(s_t, y_t)\, dt | (s_0, y_0) = (s, y)\right].$$

■

Figure 1 offers a stylized representation of the behavior of the convenience yield in our model when the friction is present. The commodity is sold from storage ($z^* > 0$) in the lower unshaded region, while it is bought and stored in the upper unshaded region. The convenience yield is zero in both regions. The shaded regions show where inventory is held constant ($z^* = 0$). The convenience yield is negative in the gray region and positive in the black region.[22] If the friction

[21] Of course, this calculation ignores the value of any real options embedded in the stored commodity.

[22] When the harvest is sufficiently high, mean reversion imparts a negative trend in the harvest, and therefore a positive trend in the spot price — the expected return shortfall (the convenience yield) is negative. Similarly, when the harvest is relatively low, the positive trend in the harvest implies a positive expected return shortfall, and therefore a positive convenience yield.

Figure 1: Convenience yield

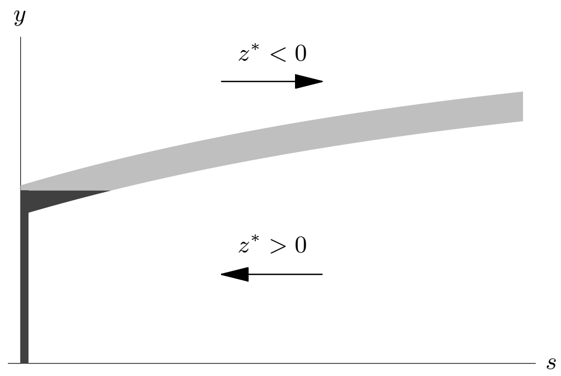

Notes. The commodity is sold from storage ($z^* > 0$) in the lower unshaded region, while it is bought and stored in the upper unshaded region. The convenience yield is zero in both regions. The shaded regions show where inventory is held constant ($z^* = 0$). The convenience yield is negative in the gray region and positive in the black region. If the friction is removed from the model, the shaded region is restricted to that part of the vertical axis where the harvest is so low during a stock-out that none of it is diverted into storage.

is removed from the model, the shaded region is restricted to that part of the vertical axis where the harvest is so low during a stock-out that none of it is diverted into storage.

From Proposition 7, the extent of backwardation equals the present value of the flow of the convenience yield over the period until the relevant forward contract's delivery date. In the frictionless case, this flow is zero except when a stock-out occurs. Based on our numerical analysis in Section 6, the extent of backwardation is approximately zero until a stock-out is imminent; consistent with the impossibility of cash-and-carry arbitrage, it is never negative. As Figure 1 suggests, adding a storage friction to the competitive storage model introduces a small region where the convenience yield is positive and a larger region where it is negative. The former increases the extent of backwardation relative to the frictionless case, but the effect will be greatest when a stock-out is imminent. The latter, however, reduces the extent of backwardation and — as we will see in Section 6 — makes it negative for all but low storage levels.

Therefore, two things happen in equilibrium. With or without a storage friction in the model, when inventory is low firms hold inventory at a level such that $B > 0$; that is, the marginal storage operator is holding inventory even though the (apparent) return from selling immediately is greater. When a storage friction is added, for all but very low levels of inventory firms hold inventory at a level such that $B < 0$; that is, the marginal storage operator is holding inventory when the (apparent) return from adding even more of the commodity to storage is positive. However, in both cases the marginal storage operator is behaving rationally. In the first case, making decisions on the basis of B alone ignores the value of the real option to wait and sell the stored commodity at a later date. That is, the best alternative to selling immediately is not necessarily to commit to sell τ years in the future, but rather to retain the option to sell at any time; the value of this option is omitted from the B calculation. In the second case, making decisions on the basis of B alone ignores the value of the real option to wait and buy more units of the commodity at a later date. That is, the best alternative to buying immediately is not necessarily to commit to buy τ years in the future, but rather to retain the option to buy at any time (or not at all); thus the term $e^{-(r+\varepsilon)\tau}F(\tau, s, y)$ in B overstates the cost of the alternative to buying the commodity immediately.

6 Storage and price behavior

Although the description of the model in this paper is relatively straightforward, the key equations are highly nonlinear and (as we will see) the implied price behavior is complex. Further investigation of the role of storage frictions and the extent to which our model matches the stylized facts of commodity price distributions requires numerical analysis. In order to maintain the focus of the paper, we do not attempt a formal calibration of the sort implemented for the frictionless model by other authors (Chambers and Bailey, 1996; Deaton and Laroque, 1992, 1996; Routledge et al., 2000). Instead, we assign representative values for the model's underlying parameters and analyze the impact of varying these values for the most important parameters.

6.1 Parameter specification

We assume that the harvest evolves according to

$$dy_t = \eta(\mu - y_t)dt + \sigma y_t^{1/2} d\xi_t$$

and that the inverse demand curve is described by

$$\psi(x) = \gamma e^{\alpha(\mu - x)},$$

where η, μ, σ, γ and α are positive constants. Harvests generated by this process are mean-reverting with the long-run harvest being drawn from the gamma distribution with shape and scale parameters $2\eta\mu/\sigma^2$ and $\sigma^2/(2\eta)$ respectively. As a consequence, the unconditional mean and variance of the harvest are μ and $\sigma^2\mu/(2\eta)$ respectively. Deviations of harvests from the long-run level are expected to decline at rate η, the harvest can never become negative and, provided $2\eta\mu \geq \sigma^2$, there is sufficient upward drift to make the origin inaccessible, in which case the harvest will always be positive.[23] We choose the units in which the harvest is measured so that the long-run average harvest level is unity, while we choose the units in which the price is measured so that the price is unity when consumption equals the long-run average harvest. That is, without loss of generality, we set $\gamma = \mu = 1$.

In order to complete the specification of harvest and demand we need to select values for η, σ, and α. Setting $\eta = \log 2 = 0.693$ implies that the half-life of harvest shocks is one year, while

[23] Cox et al. (1985) use this process to model interest rates, where properties of mean reversion and non-negativity are also important.

Table 1: Baseline parameters

Harvest	$dy_t = \eta(\mu - y_t)dt + \sigma y_t^{1/2} d\xi_t$
μ	1.000
η	0.693
σ	0.589
Inverse demand	$\psi(x) = \gamma e^{\alpha(\mu - x)}$
γ	1
α	2
Storage cost	
ε	0.030
k_i	0.025
k_o	0.025
Other	
r	0.04
λ	0.04

setting $\sigma = (\eta\mu/2)^{1/2} = 0.589$ implies that the unconditional standard deviation of the harvest is half the unconditional mean.[24] When consumption equals the long-run average harvest level, demand elasticity equals

$$\frac{d\psi(x)}{dx} \cdot \frac{x}{\psi(x)}\bigg|_{x=\mu} = -\alpha\mu.$$

We consider a range of values for α in the analysis that follows, but we ultimately settle on $\alpha = 2$, so that the demand elasticity equals two when consumption equals the long-run average harvest level.

We set the risk-free interest rate equal to $r = 0.04$. In order to specify the market-price of harvest risk we suppose that the market risk premium is 0.08 and that the rate of return on the market portfolio has standard deviation 0.2. We set the correlation coefficient between the rate of return on the market portfolio and the exogenous harvest equal to 0.1, implying a market price of harvest risk of $\lambda = 0.04$.

This leaves the costs associated with storage. Along with Routledge et al. (2000), we suppose that inventory decays at a rate of $\varepsilon = 0.03$.[25] We consider several different values for the costs of moving the commodity into and out of storage, but in all cases we assume that $k_i = k_o$. Our parameter choices are summarized in Table 1.

6.2 Optimal storage policy and the implied price behavior

We begin by solving for the social planner's optimal storage policy, which is illustrated in Figure 2. The top panel plots $z^*(s, y)$ for two different levels of frictions for the case where

[24]Note that in their one-factor calibrated model, Routledge et al. (2000) find that net demand has volatility equal to approximately half of the mean, so this assumption is consistent with their model.

[25]This value for ε is much lower than those estimated by Deaton and Laroque (1992, 1996), but Routledge et al. (2000) argue that a small value is necessary if storage is to be significant.

Figure 2: Social planner's storage policy

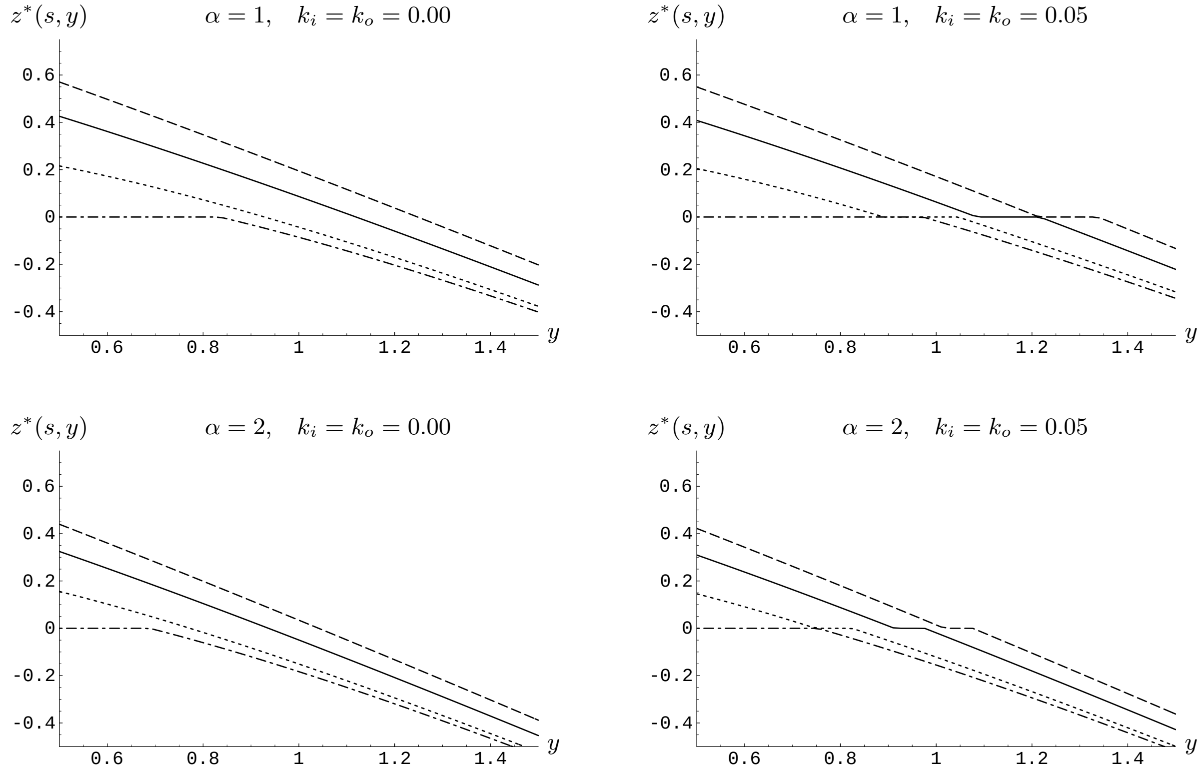

Notes. Each graph plots $z^*(s, y)$ as a function of the harvest (y) for four different levels of total storage: the dashed-dotted curves correspond to $s = 0$, the dotted curves to $s = 0.1$, the solid curves to $s = 0.5$, and the dashed curves to $s = 1$. The top panel plots $z^*(s, y)$ for two different levels of frictions for the case where $\alpha = 1$, while the bottom panel repeats this exercise for $\alpha = 2$. In each panel the left graph corresponds to the case where $k_i = k_o = 0$, while the right graph corresponds to the case where $k_i = k_o = 0.05$. All other parameter values are given in Table 1.

$\alpha = 1$, while the bottom panel repeats this exercise for $\alpha = 2$. In each panel, $k_i = k_o = 0$ in the left graph and $k_i = k_o = 0.05$ in the right graph. The four curves plot $z^*(s, y)$ as a function of the harvest (y) for different levels of total storage: the dashed-dotted curves correspond to $s = 0$, the dotted curves to $s = 0.1$, the solid curves to $s = 0.5$, and the dashed curves to $s = 1$. The graphs in the left column show that when the harvest is poor speculators sell the commodity from storage (that is, $z^* > 0$); when the harvest is good, speculators buy the commodity and store it (that is, $z^* < 0$). Inspection of the four curves in each graph shows that when inventory is high, inventory is either sold more aggressively (that is, positive values of z^* are larger) or bought more cautiously (that is, negative values of z^* are smaller). The graphs in the right column show that when a friction is introduced into the model the most noticeable change is the appearance of a range of harvest values for which speculators temporarily withdraw from the spot market — they neither buy nor sell the commodity until the harvest moves out of this region. This is consistent with the behavior noted in Proposition 2.

Comparison of the two panels in Figure 2 reveals the importance of the shape of the demand curve. As α increases, so that stock-outs become more costly in terms of overall welfare, the curves shift left, implying that inventory is either bought more aggressively (that is, negative values of z^* become larger) or sold more cautiously (that is, positive values of z^* become smaller). As our simulations will confirm, this leads to greater storage overall and less frequent stock-outs. The intuition for this behavior is straightforward: if the demand curve is sufficiently steep for low levels of consumption, the welfare losses from a future stock-out will be so high that the social planner will hold back enough of the harvest in order to eke out whatever inventory remains. Thus, we expect to see fewer stock-outs when the demand curve is very steep for low levels of demand. In line with the observation that stock-outs are rare in actual commodity markets, from now on we concentrate on the case where $\alpha = 2$, but point out the effects of other values in the next subsection.

Figure 3 reveals some properties of the spot price that is implied by the social planner's optimal storage policy. The top panel plots the spot price using the same format as Figure 2; that is, the four curves in each graph plot $P(s, y)$ as a function of the harvest (y) for different levels of total storage (s). The graphs show that the spot price is decreasing in both the level of the harvest and total storage. When the storage technology contains a friction, the spot price is especially sensitive to the harvest in the region where $z^* = 0$ — this corresponds to the kinks in each of the curves in the top panel of Figure 3. To see why, recall that the spot price equals $P(s, y) = \psi(y + z^*(s, y))$, so that

$$\frac{\partial P}{\partial y} = \psi'(y + z^*) \left(1 + \frac{\partial z^*}{\partial y} \right).$$

If storage operators are not trading in the spot market (that is, $z^* = 0$), then $P(s, y) = \psi(y)$ and, even though total storage is nonzero, we get the no-storage spot price. In particular, if the harvest rises by Δy, the spot price will fall by approximately $|\psi'(y)|\Delta y$. Suppose, instead, that $z^* > 0$, so that storage operators are currently selling the commodity from storage. Now an increase in the harvest of Δy has two effects on the supply of the commodity in the spot market: the amount supplied by producers rises by Δy and the amount supplied by storage operators changes by $\frac{\partial z^*}{\partial y} < 0$. Since these two changes work in opposing directions, the sensitivity of the spot price to the level of harvest is lower than in the case when $z^* = 0$.[26] The greater sensitivity of the spot price to harvest shocks when storage operators are not trading in the spot market can be seen more clearly in the bottom panel of Figure 3, which plots the spot price volatility, $|\phi(y)\frac{\partial P}{\partial y}|$. The high-volatility region does not arise in the frictionless case except during a stock-out. Even when storage does contain a friction, spot price volatility is relatively low whenever $z^* \neq 0$. However, it can spike up dramatically when $z^* = 0$.

The behavior of the convenience yield and the extent of backwardation is shown in Figure 4. The top panel plots the convenience yield, $CY(s, y)$, using the same format as Figure 2. When

[26] In our numerical simulations, $-1 < \frac{\partial z^*}{\partial y} < 0$, so that $1 + \frac{\partial z^*}{\partial y} > 0$ and positive harvest shocks result in lower

Figure 3: Spot price behavior

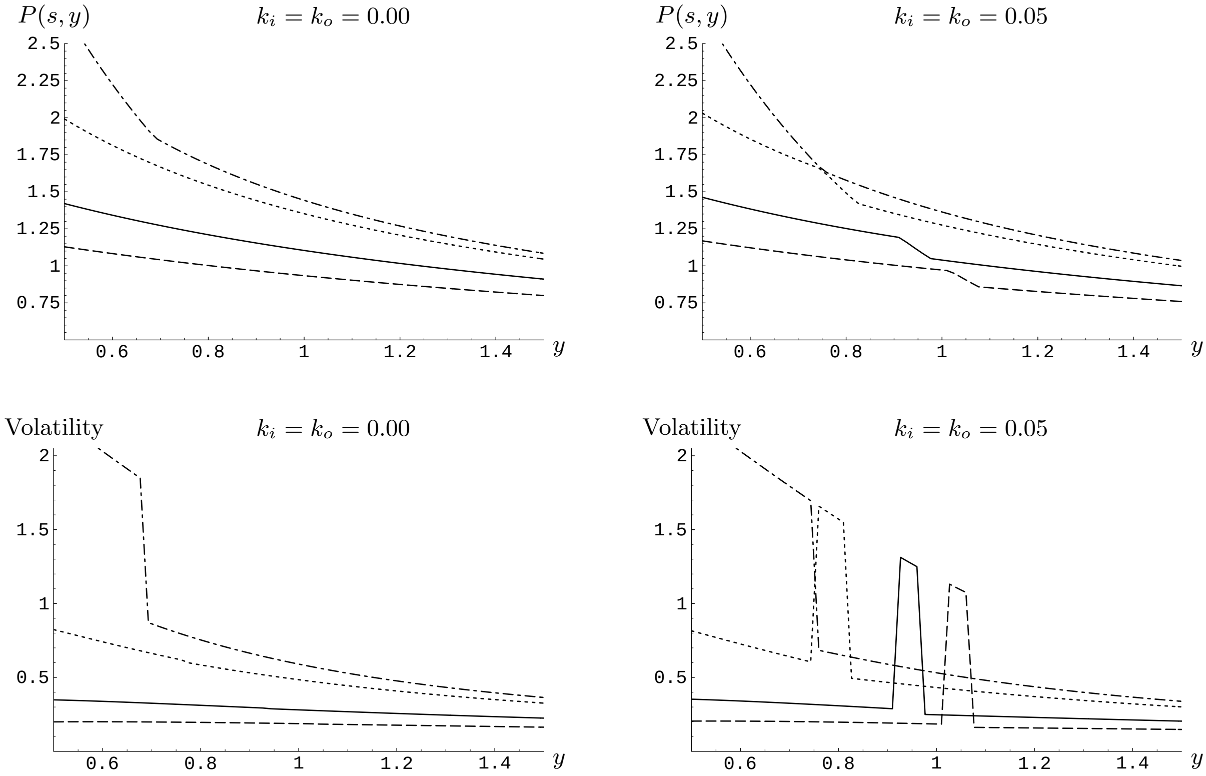

Notes. Each graph in the top panel plots the spot price, $P(s, y)$, as a function of the harvest, y, for four different levels of total storage: the dashed-dotted curves correspond to $s = 0$, the dotted curves to $s = 0.1$, the solid curves to $s = 0.5$, and the dashed curves to $s = 1$. Each graph in the bottom panel plots the instantaneous volatility of the spot price for the same four storage levels. In each panel the left graph corresponds to the case where $k_i = k_o = 0$, while the right graph corresponds to the case where $k_i = k_o = 0.05$. All other parameter values are given in Table 1.

there are no frictions, the convenience yield is only nonzero during a stock-out. In contrast, when frictions are introduced into the model both positive and negative convenience yields can arise when inventory is strictly positive. This explains the difference between the two graphs in the bottom panel of Figure 4, which plot the extent of backwardation, $B(0.25, s, y)$, as defined in equation (17). Recall that $B(\tau, s, y)$ measures the amount by which the payoff from an immediate sale exceeds the payoff from a forward sale τ years in the future and (from Proposition 7) that it equals the present value of the flow over the next τ years of the convenience yield. In the frictionless case, when a stock-out is required for the convenience yield to be nonzero, backwardation is close to zero except for very low levels of storage; it is never negative. If a friction is present, the extent of backwardation behaves similarly to the frictionless case when storage is low, but can become negative for other storage levels.

The non-monotonic relationship between B and the harvest, which only occurs when the

spot prices.

Figure 4: Convenience yield and backwardation

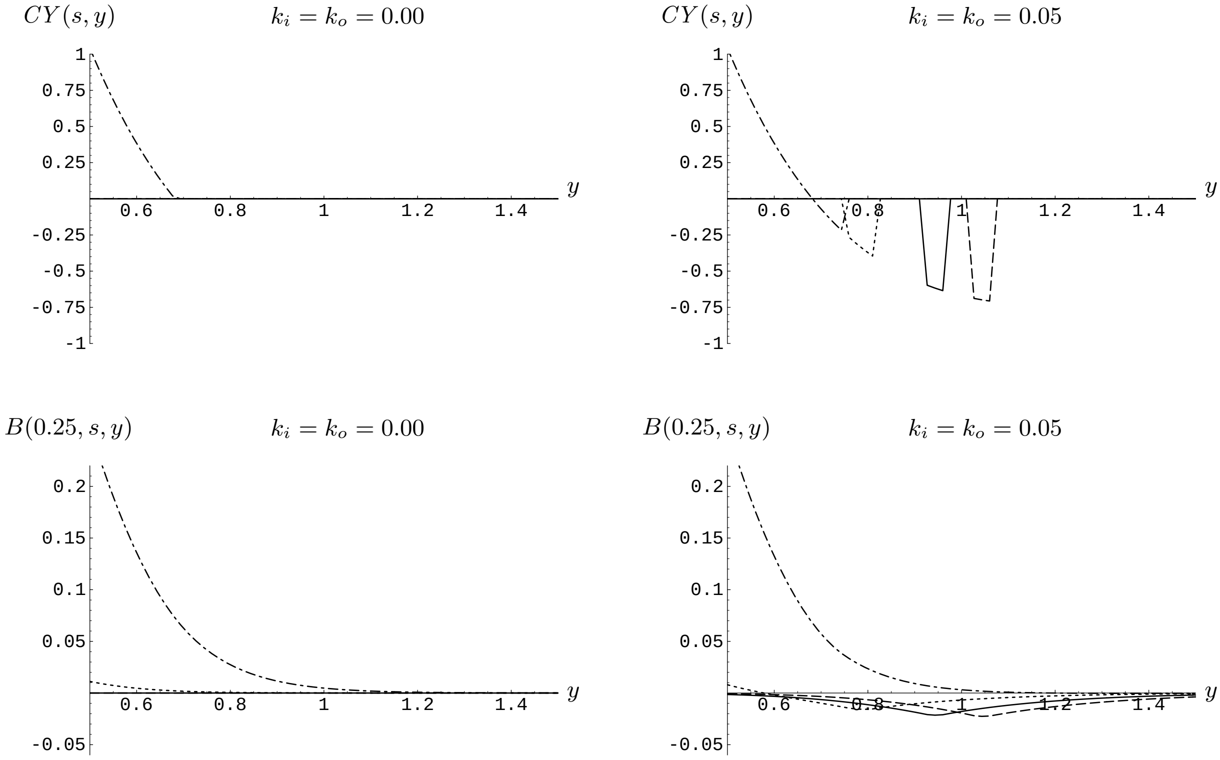

Notes. Each graph in the top panel plots the convenience yield, $CY(s,y)$, as a function of the harvest, y, for four different levels of total storage: the dashed-dotted curves correspond to $s = 0$, the dotted curves to $s = 0.1$, the solid curves to $s = 0.5$, and the dashed curves to $s = 1$. Each graph in the bottom panel plots the extent of backwardation, $B(0.25, s, y)$, for the same four storage levels. In each panel the left graph corresponds to the case where $k_i = k_o = 0$, while the right graph corresponds to the case where $k_i = k_o = 0.05$. All other parameter values are given in Table 1.

storage friction is present, impacts on the term structure of forward price volatilities. Recall from equation (17) that the forward price satisfies

$$F(\tau, s, y) = e^{(r+\varepsilon)\tau} \left(P(s,y) - B(\tau, s, y)\right). \tag{19}$$

Application of Itô's Lemma shows that $\log F$ has volatility

$$\frac{\phi(y)}{F(\tau,s,y)} \left| \frac{\partial F}{\partial y} \right| = \frac{\phi(y)}{P(s,y) - B(\tau,s,y)} \left| \frac{\partial P}{\partial y} - \frac{\partial B}{\partial y} \right|,$$

while $\log P$ has volatility

$$\frac{\phi(y)}{P(s,y)} \left| \frac{\partial P}{\partial y} \right|.$$

The two graphs in Figure 5 plot the amount by which the (log) spot price volatility exceeds the (log) forward price volatility as a function of the harvest level for the same four storage levels as in the preceding figures. The left graph corresponds to the case where $k_i = k_o = 0$, while the right graph corresponds to the case where $k_i = k_o = 0.05$. In the frictionless case, the spot

Figure 5: Samuelson's hypothesis

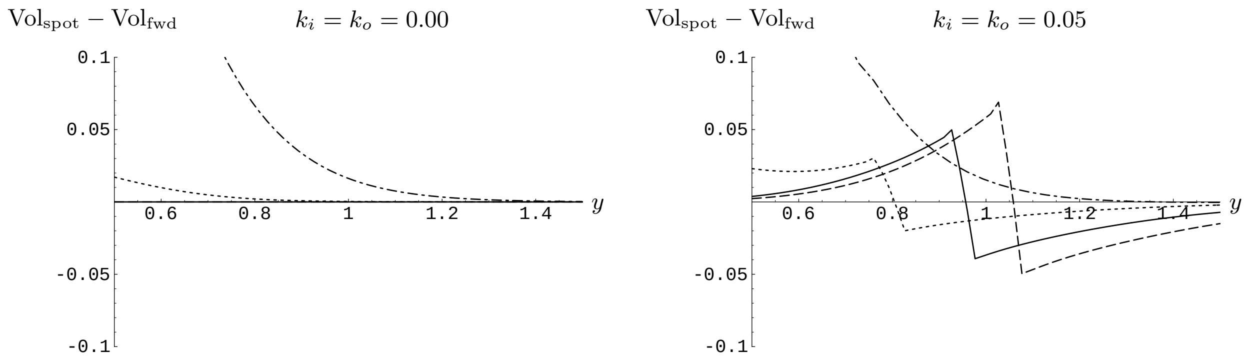

Notes. Each graph plots the amount by which spot price volatility exceeds (three-month-ahead) forward price volatility, as a function of the harvest, y, for four different levels of total storage: the dashed-dotted curves correspond to $s = 0$, the dotted curves to $s = 0.1$, the solid curves to $s = 0.5$, and the dashed curves to $s = 1$. The left graph corresponds to the case where $k_i = k_o = 0$, while the right graph corresponds to the case where $k_i = k_o = 0.05$. All other parameter values are given in Table 1.

price is significantly more volatile than the forward price during a stock-out, very slightly more volatile when inventory is very low, and the two prices have the same volatility for all other levels of inventory. That is, the Samuelson (1965) hypothesis that forward price volatility rises as the delivery date nears holds in the frictionless case, although the change in volatility only occurs when inventory is low.[27]

The right hand graph in Figure 5 shows that, when the storage friction is present, during stock-outs the behavior of spot and forward price volatility is similar to that in the frictionless case. However, two differences arise when inventory is positive. Firstly, during moderately poor harvests, the spot price is more volatile than the forward price. This occurs because, as is evident from Figures 3 and 4, harvest shocks move P and B in the same direction; equation (19) shows that the two have opposing effects on the forward price. Secondly, for moderately good harvests, the spot price is actually less volatile than the forward price. This time, Figures 3 and 4 show that harvest shocks move P and B in opposite directions, amplifying the effect on the forward price. Both of these results recall the observation by Bessembinder et al. (1996) that Samuelson's hypothesis is driven by a negative correlation between changes in the spot price and changes in the slope of the term structure of forward prices — in our case, by a positive

[27] Routledge et al. (2000) claim (pp. 1317–1318) that Samuelson's hypothesis can be violated in their (frictionless) storage model if inventory is sufficiently large. However, the apparent violations that they note only arise when the delivery date is so close, and inventory so high, that stock-outs cannot occur prior to delivery. In that case, the forward price equals the spot price scaled up by the usual cost-of-carry formula, so that forward price volatility equals spot price volatility scaled up by the same factor. When the volatilities are measured using log prices, as here, Samuelson's hypothesis continues to hold.

correlation between changes in the spot price and changes in the extent of backwardation.

The origin of this behavior is the region where the convenience yield is negative (and, therefore, where the spot price is expected to grow rapidly). For example, for moderately low harvest levels, an increase in the harvest lowers the spot price but, because it moves the system closer to the region where the convenience yield is negative, makes imminent rapid increases in the spot price more likely. The market will factor in the likelihood that the reduction in the spot price will soon be reversed and, as a result, the forward price will fall by less than the spot price. For moderately high harvest levels, an increase in the harvest also lowers the spot price, but — in contrast to the previous case — the system moves further away from the region where the convenience yield is negative. This makes imminent rapid increases in the spot price less likely than before, compounding the effect of the reduced current spot price. Thus, the forward price will fall by more than the spot price.

6.3 Simulations and the stylized facts of commodity markets

In this subsection we report the results of our analysis of storage and price data simulated using our model, with and without the storage friction. For each parameter combination we consider, we begin from a situation with storage equal to zero and the harvest at its long-run level and then simulate 600 years of harvests using an Euler approximation with 260 steps per year (that is, approximately one step per business day). We then drop the first 100 years of observations, giving us a 500 year sample of approximately daily observations that should not depend on our chosen initial state.

6.3.1 The importance of demand elasticity and storage frictions

Section 6.2 identifies α and k as key determinants of the behavior of spot and forward prices. In this section we investigate their role in more detail. We begin in Table 2, by reporting some results from repeating the simulation exercise described at the start of Section 6.3 for different combinations of the key parameters α and k. (All other parameters are as specified in Table 1.) The first two columns report the average storage level and the probability that the storage level is less than 0.01 of the mean (which we take as our approximation of the frequency with which stock-outs occur).[28] The next five columns reveal information about the distribution of the interest- and storage-adjusted basis, $i_t = I(\tau, s_t, y_t)$ for the case where $\tau = 0.25$. The first three of these columns report the probabilities that the basis is negative, zero, and positive, while the remaining columns report the conditional means.[29] The final two columns in Table 2 report the degree of (annual) autocorrelation in the spot price, both when storage is not allowed (the first

[28] Due to rounding errors introduced during our numerical solution process, we cannot simply use the frequency with which $s = 0$.

[29] Since the underlying data derives from our numerical solution of the social planner's problem, and is therefore subject to rounding errors, we regard i as zero if and only if $|i| \leq 0.0001$ (that is, the basis must be no greater than one basis point).

Table 2: The impact of demand elasticity and storage frictions

	Storage level		Interest- and storage-adjusted basis					Spot price auto.	
	Mean	Pr[stock-out]	Pr$[i<0]$	Pr$[i=0]$	Pr$[i>0]$	$E[i\|i<0]$	$E[i\|i>0]$	Without	With
$k=0.000$									
$\alpha=1.0$	0.483	0.110	0.460	0.540	0.000	−0.098	n/a	0.394	0.641
$\alpha=1.5$	0.810	0.055	0.290	0.710	0.000	−0.107	n/a	0.375	0.680
$\alpha=2.0$	1.148	0.032	0.197	0.803	0.000	−0.110	n/a	0.353	0.702
$\alpha=2.5$	1.493	0.021	0.124	0.876	0.000	−0.128	n/a	0.328	0.690
$\alpha=3.0$	1.847	0.013	0.078	0.922	0.000	−0.159	n/a	0.303	0.702
$k=0.025$									
$\alpha=1.0$	0.451	0.121	0.331	0.016	0.653	−0.138	0.015	0.394	0.619
$\alpha=1.5$	0.789	0.057	0.179	0.022	0.800	−0.173	0.018	0.375	0.664
$\alpha=2.0$	1.138	0.031	0.093	0.019	0.889	−0.227	0.021	0.353	0.693
$\alpha=2.5$	1.496	0.020	0.064	0.019	0.917	−0.237	0.022	0.328	0.685
$\alpha=3.0$	1.864	0.013	0.055	0.018	0.926	−0.212	0.030	0.303	0.696
$k=0.050$									
$\alpha=1.0$	0.419	0.136	0.310	0.010	0.679	−0.150	0.027	0.394	0.599
$\alpha=1.5$	0.765	0.059	0.164	0.013	0.823	−0.188	0.033	0.375	0.645
$\alpha=2.0$	1.126	0.031	0.085	0.011	0.904	−0.245	0.039	0.353	0.682
$\alpha=2.5$	1.497	0.020	0.049	0.010	0.942	−0.303	0.047	0.328	0.679
$\alpha=3.0$	1.878	0.012	0.031	0.009	0.959	−0.359	0.052	0.303	0.689

Notes. The first two columns report the average storage level and the probability that the storage level is less than 1% of the mean (which we take as our approximation of the frequency with which stock-outs occur), while the next five columns reveal information about the distribution of the basis i defined in (18). The final two columns report the degree of (annual) autocorrelation in the spot price, both when storage is not allowed (the first column) and when it is allowed (the second column). We use a 500 year sample of approximately daily observations. Baseline parameter values are given in Table 1.

column) and when it is allowed (the second column).[30] Each panel corresponds to a different level of frictions, while each row within each panel corresponds to a different demand elasticity.

As expected, higher levels of the parameter determining demand elasticity lead to higher average inventory and less frequent stock-outs. When storage is frictionless, the adjusted basis is occasionally negative, usually zero, and never positive; greater demand elasticity reduces the probability that the basis is negative but increases its average magnitude when it is negative. Although introducing a storage friction reduces the frequency of negative adjusted basis values, the most dramatic effect is that the adjusted basis is usually positive. On average, the magnitude of the adjusted basis is much smaller in positive states than negative ones. Greater demand elasticity reduces the probability that the basis is negative and raises the probability that it is

[30]The annual autocorrelations are calculated using a 500-year sample of non-overlapping annual observations (that is, we retain only every 260th observation).

positive; it increases the average magnitude of the adjusted basis in both negative and positive states. Finally, comparison of the various rows in each panel shows that greater demand elasticity reduces the serial correlation in the spot price when storage is not possible and increases it when storage is possible. In fact, in some cases storage doubles the degree of autocorrelation, suggesting that the competitive storage model may, after all, be able to explain high serial correlation in commodity prices — but only if the demand elasticity is so high that stock-outs are rare.

As Table 2 shows, one distinguishing feature of storage frictions is that the interest- and storage-adjusted basis can become positive. Difficulties in measuring storage costs will make it difficult to test this prediction accurately. Ng and Pirrong (1994) estimate the interest- and storage-adjusted basis for four industrial metals and find it is negative for all days in their approximately six-year sample. Other authors have investigated the properties of the interest-adjusted basis. For example, Heaney (2006) analyzes monthly data on metal spot and forward prices for the period November 1964 to December 2003 and finds that the maximum value of the interest-adjusted basis is 1.9% for copper, 3.1% for lead, and 5.7% for zinc. If the interest-adjusted basis is adjusted for storage costs, the distribution shifts left. However, provided storage costs are not too large, Heaney's results offer some support for positive values of the interest- and storage-adjusted basis, and therefore some support for the role of frictions in commodity markets.[31]

6.3.2 Commodity price dynamics

An important advantage of formulating our model in continuous time is that we are able to simulate price data of any frequency. We exploit this feature in the current section, where we analyze the implications of storage frictions for the behavior of daily spot and forward prices. We begin by plotting simulated daily data in order to informally illustrate the behavior predicted by the model. We then fit various GARCH(1,1) models to our full simulated data set of 500 years of daily commodity prices, and find that (as long as frictions are present) our model captures some key price behavior observed in actual commodity markets.

Figure 6 displays eight years of simulated daily data for the cases when there are no frictions (the left hand graphs) and when $k_i = k_o = 0.05$ (the right hand graphs); in both cases $\alpha = 2$ and $\tau = 0.25$. There are no stock-outs in this sample period.[32,33] The top pair of graphs plots

[31] Fama and French (1988) analyze daily data on spot and forward prices of industrial metals over the period 1972–1983 and find that the proportion of days on which the interest-adjusted basis (for forward contracts with three months until delivery) is positive ranges from 3.2% (tin) to 60% (aluminium). In their data set, the basis has greater magnitude (on average) on days when it is negative than on those when it is positive. For example, the respective averages are -0.0286 and 0.0026 for tin, and -0.0421 and 0.0079 for aluminium.

[32] We choose a period that does not contain a stock-out in order to highlight the differences between the cases with and without storage frictions. When stock-outs occur the two cases both feature high spot prices, high spot price volatility, and a large negative interest- and storage-adjusted basis.

[33] Note from Table 2 that the average storage levels for these parameterizations of the model are $\bar{s} = 1.148$ and

Figure 6: Eight years of simulated daily data

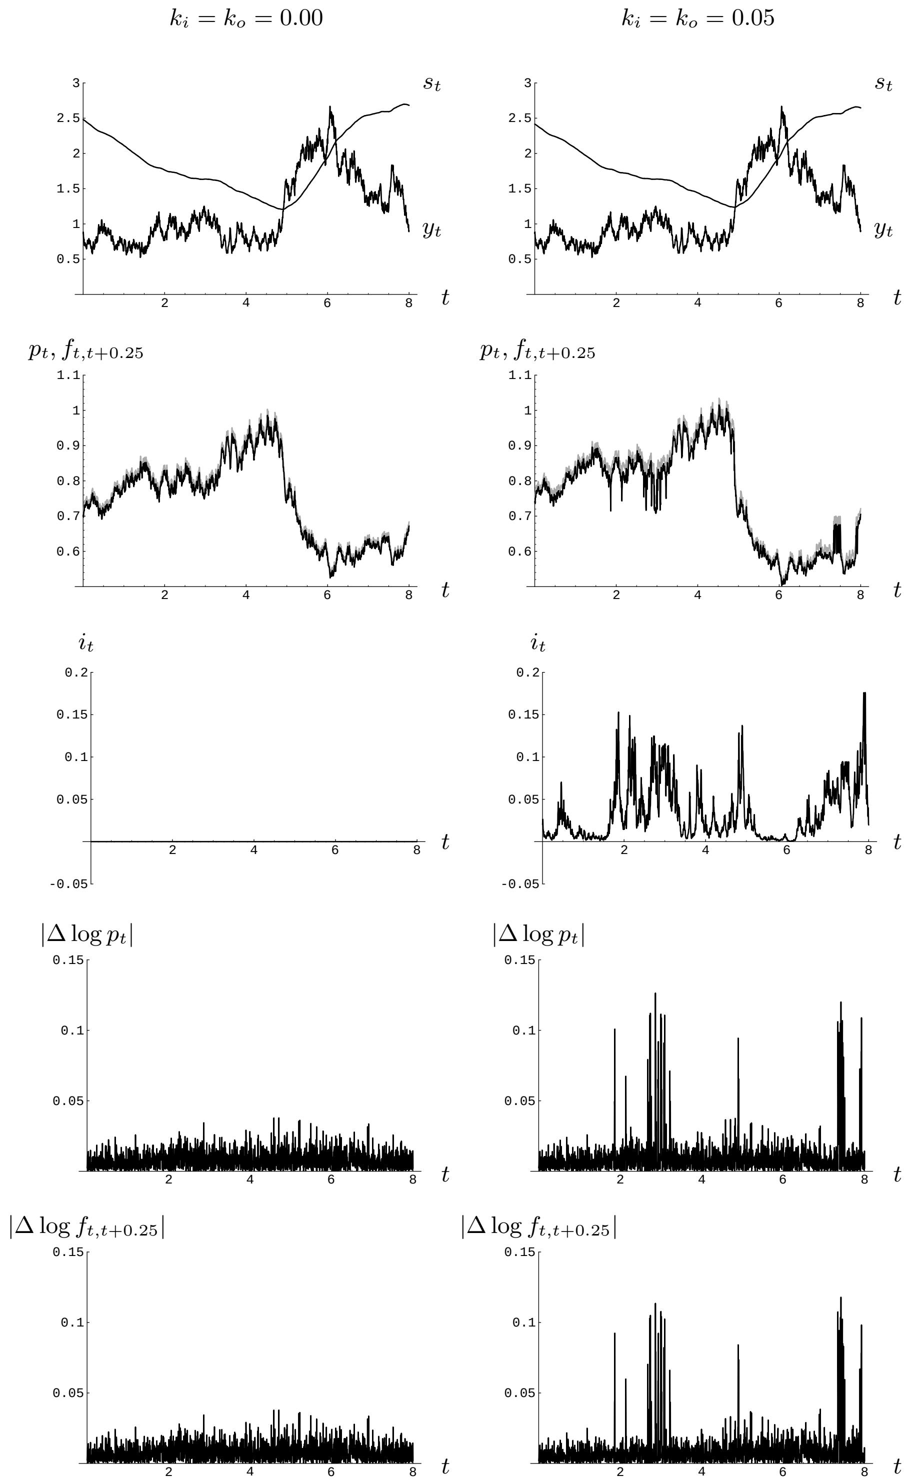

Notes. Each column of graphs plots eight years of simulated daily data for the cases when there are no frictions (the left hand graphs) and when $k_i = k_o = 0.05$ (the right hand graphs). The top pair of graphs plots the simulated paths of storage (the smooth curve) and the harvest (the jagged curve), while the second pair plots the spot price (black) and forward price (light gray). The third pair of graphs plot the interest- and storage- adjusted basis. The fourth and fifth pair of graphs plot the magnitude of spot and forward price returns, respectively.

the simulated paths of storage (the smooth curve) and the harvest (the jagged curve), while the second pair plots the spot price (black) and forward price (light gray). Inspection of the top two graphs in each column might suggest that the behavior is similar in the two cases. However, there are subtle differences that become clear when we consider the other graphs in each column.

The third pair of graphs in Figure 6 illustrates one of the most important implications of storage frictions for the issues we analyze in this paper. The left graph shows that the interest- and storage-adjusted basis is approximately zero at all times during this eight year period, while the right hand graph shows sustained periods when the basis is positive and economically significant. In the cases we have considered here, the demand curve is sufficiently steep for low levels of demand that stock-outs almost never occur. This means that unless frictions are present the instantaneous convenience yield will almost always be zero, which in turn implies that the basis will be approximately zero except when stock-outs are imminent.

The fourth pair of graphs, which plots the magnitude of spot price returns, show that when there are frictions in storage there are several brief periods during which spot price volatility explodes to levels several times greater than long-run average volatility, even though inventory is positive. During these periods speculators withdraw from the spot market, as can be seen from the top graph, which shows that total storage remains constant during these periods of high spot price volatility.[34] In contrast, in the frictionless case spot price volatility is relatively stable as long as total storage is positive. The fifth pair of graphs, which plots the magnitude of three-month-ahead forward price returns, suggests that forward price volatility behaves in much the same way as spot price volatility.

The intuition underlying this volatility clustering is straightforward. From Figure 3, the volatility spikes occur when $z^* = 0$ — that is, when speculators withdraw from the spot market and are therefore unable to buffer the economy from harvest shocks. As Figure 2 shows, for any given level of storage, the region where $z^* = 0$ spans a nontrivial interval of harvest values. The continuity of the harvest process therefore ensures that the state will remain in this region for some time. Thus, a sudden increase in volatility indicates that speculators have largely withdrawn from the spot market, and the non-infinitesimal nature of the storage frictions means that they will not immediately return; while they are inactive in the spot market, volatility will remain high.

The top pair of graphs in Figure 7 plots the spot price (black) and the market value of the stored commodity (light gray), $v_t = V(s_t, y_t)$. In the frictionless case the market value of the stored commodity is exactly equal to the spot price. When the storage friction is present, the market value is a constant proportion of the spot price whenever speculators are buying the

$\bar{s} = 1.126$ for the cases $k_i = k_o = 0$ and $k_i = k_o = 0.05$ respectively. Thus, in both cases total storage exceeds its long-run average for the entire eight year period.

[34]From the perspective of an individual speculator, the sudden increase in volatility increases the value of the speculator's delay options, encouraging it to keep existing inventory in storage, at least until some of the uncertainty surrounding future levels of the spot price is resolved.

Figure 7: Spot price and the market value of the stored commodity

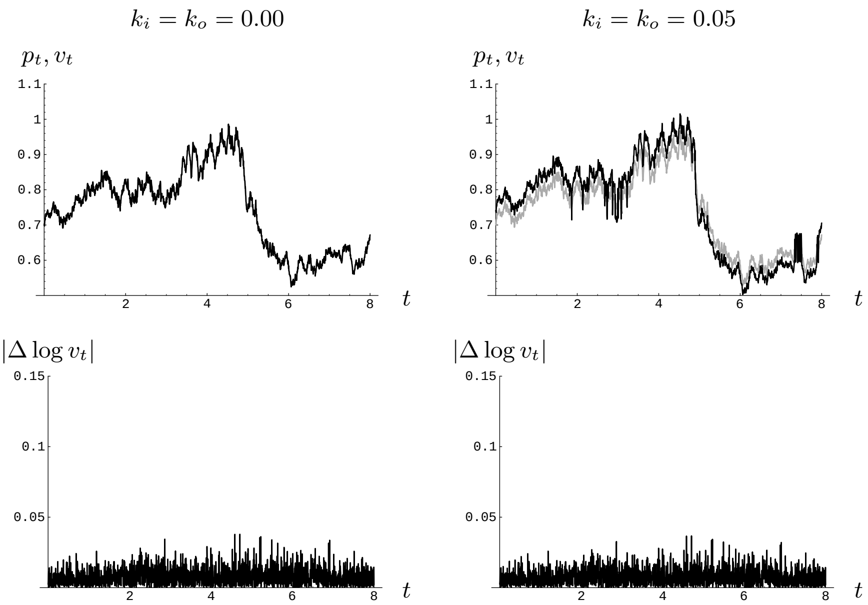

Notes. Each column of graphs plots eight years of simulated daily data for the cases when there are no frictions (the left hand graphs) and when $k_i = k_o = 0.05$ (the right hand graphs). The top pair of graphs plots the spot price (black) and the market value of the stored commodity (light gray). The bottom pair of graphs plots the magnitude of market value returns.

commodity, and a (different) constant proportion when they are selling; the ratio of market value to spot price fluctuates between the two constants when speculators are not active in the spot market. This separation between the market value and the spot price insulates the market value of the stored commodity from the volatility spikes that affect the spot market. This is clear from the bottom pair of graphs, which plots the magnitude of market value returns in the two cases: frictions have little impact on the volatility of the market value of the stored commodity.

We now explore how well the volatility of the commodity price series mimics the stylized facts suggested by commodity price studies by fitting various GARCH processes to our simulated daily spot and forward price data and comparing our results with those contained in other authors' studies of actual commodity price data. Our results are reported in Table 3.[35] The top panel fits a standard GARCH(1,1) model, with t-distributed disturbances, to the spot and three-month-ahead forward prices separately. Specifically, we fit

$$\begin{aligned} \Delta \log p_t &= \mu + e_t, \\ e_t &\sim t(0, \sigma_t^2, \nu), \\ \sigma_t^2 &= \beta_0 + \beta_1 \sigma_{t-\Delta t}^2 + \beta_2 e_{t-\Delta t}^2, \end{aligned} \tag{20}$$

[35]Because we are using 500 years of daily data, traditional measures of statistical significance convey little information. Therefore, we only report the point estimates of the parameters in Table 3.

Table 3: GARCH models for spot and forward prices

	$k = 0.000$		$k = 0.025$		$k = 0.050$	
	Spot	Forward	Spot	Forward	Spot	Forward
Model 1: Standard GARCH(1,1)						
constant	0.000000	0.000000	0.000002	0.000002	0.000011	0.000009
$\sigma^2_{t-\Delta t}$	0.982	0.984	0.930	0.925	0.802	0.801
$e^2_{t-\Delta t}$	0.018	0.016	0.066	0.071	0.190	0.204
1/df	0.015	0.014	0.162	0.161	0.223	0.224
Model 2: GARCH(1,1) with basis as an explanatory variable						
constant	0.000000	0.000000	0.000002	0.000002	0.000011	0.000009
$\sigma^2_{t-\Delta t}$	0.982	0.984	0.921	0.919	0.772	0.783
$e^2_{t-\Delta t}$	0.017	0.016	0.071	0.075	0.199	0.207
$\lvert i_{t-\Delta t}\rvert$	0.000009	0.000005	0.000040	0.000029	0.000168	0.000105
1/df	0.014	0.013	0.157	0.158	0.216	0.219
Model 3: GARCH(1,1) with an asymmetric basis effect						
constant	0.000000	0.000000	0.000002	0.000002	0.000013	0.000010
$\sigma^2_{t-\Delta t}$	0.982	0.984	0.923	0.921	0.766	0.782
$e^2_{t-\Delta t}$	0.017	0.016	0.068	0.073	0.200	0.207
$i^-_{t-\Delta t}$	0.000009	0.000005	0.000076	0.000041	0.000442	0.000196
$i^+_{t-\Delta t}$	–	–	0.000021	0.000012	0.000124	0.000064
1/df	0.014	0.013	0.160	0.160	0.220	0.222

Notes. Each column reports the estimates of the variance equation in a GARCH(1,1) model for the indicated price and friction. The simulated data is the same as that used in Table 2. Baseline parameter values are given in Table 1.

to our simulated spot prices and the analogous equation to our simulated forward prices. Baillie and Myers (1991) fit this model to daily spot and futures price data for five different agricultural commodities and gold. For the agricultural commodities they report estimates of β_1 in the range 0.67–0.87 for spot prices and 0.61–0.88 for futures prices. Their corresponding estimates for β_2 are in the ranges 0.09–0.29 and 0.10–0.21, while their estimates of $1/\nu$ lie in the ranges 0.058–0.14 and 0.15–0.17.[36] As Table 3 shows, when there are no storage frictions there is almost no role for lagged-squared residuals in explaining price volatility. Moreover, the model is unable to generate prices with similar values of $1/\nu$ to those observed in actual commodity price data. However, when a friction is present the model generates price behavior remarkably similar to that identified by Baillie and Myers (1991). For example, when $k = 0.025$ both the coefficient on the ARCH term and the estimate of $1/\nu$ are close to their empirical counterparts. This suggests that frictions play an important part in explaining commodity price dynamics.

Ng and Pirrong (1994) argue that (20) is misspecified and inconsistent with the theory of

[36] Yang and Brorsen (1992) fit a modified GARCH(1,1) model — day-of-the-week and other dummies appear in the variance equation — to daily spot and futures price data on seven agricultural commodities. Their estimates for β_1 and β_2 are in the ranges 0.47–0.91 and 0.05–0.15 respectively. (They arbitrarily set $1/\nu = 0.1$.)

storage. In particular, they claim that demand and supply conditions determine both price volatility and the interest- and storage-adjusted basis. Given that storage data is not readily available (certainly not with daily frequency), Ng and Pirrong investigate the role that fundamentals play in determining price volatility by allowing σ_t^2 to vary with the magnitude of the lagged basis. Motivated by their approach, we re-estimate the GARCH model with the variance equation in (20) replaced by[37]

$$\sigma_t^2 = \beta_0 + \beta_1 \sigma_{t-\Delta t}^2 + \beta_2 e_{t-\Delta t}^2 + \beta_3 |i_{t-\Delta t}|.$$

The results of fitting this model to our simulated data are reported in the second panel of Table 3. Ng and Pirrong find that the coefficients on the lagged basis term are positive and statistically significant for the four industrial metals they consider. Moreover, spot price volatility is more sensitive to the basis term than is forward price volatility. As Table 3 shows, storage frictions are necessary for the basis term to affect price volatility in our model. Consistent with the empirical behavior documented in Ng and Pirrong (1994), spot price volatility is more sensitive to the basis term than forward price volatility.[38]

The third panel of Table 3 reports the results of allowing the interest- and storage-adjusted basis to have an asymmetric effect on price volatility. This extension is motivated by Ng and Pirrong (1996), who test whether spot price volatility is greater when the basis is negative than when it is positive. In this case, we re-estimate the GARCH model with the variance equation in (20) replaced by

$$\sigma_t^2 = \beta_0 + \beta_1 \sigma_{t-\Delta t}^2 + \beta_2 e_{t-\Delta t}^2 + \beta_3 i_{t-\Delta t}^- + \beta_4 i_{t-\Delta t}^+,$$

where

$$i^- = \begin{cases} -i & \text{if } i \le 0, \\ 0 & \text{if } i > 0, \end{cases} \quad \text{and} \quad i^+ = \begin{cases} 0 & \text{if } i \le 0, \\ i & \text{if } i > 0. \end{cases}$$

Consistent with the evidence in Ng and Pirrong (1996), we find that price volatility is more sensitive to negative basis than to positive basis when frictions are present in the model.[39,40]

[37] Ng and Pirrong (1994) estimate a bivariate GARCH model, featuring $i_{t-\Delta t}^2$ rather than $|i_{t-\Delta t}|$ in the variance equation. They report that the regressions involving $i_{t-\Delta t}^2$ and $|i_{t-\Delta t}|$ lead to similar qualitative results but that the quadratic term leads to higher log likelihood values. We also find that the two competing formulations lead to similar results, but in our case the latter formulation is superior.

[38] Ng and Pirrong (1994) report estimates of β_1 in the range 0.65–0.82 for spot prices and 0.73–0.88 for forward prices for the four industrial metals. Their corresponding estimates for β_2 are in the ranges 0.065–0.086 and 0.063–0.078. Table 3 shows that storage frictions are necessary if the competitive storage model is to match these estimates.

[39] No asymmetric effect can be estimated in the frictionless case due to its lack of negative basis observations.

[40] Our model suggests that when frictions are present, storage will have an effect on volatility in the presence of the basis terms. In fact, an improvement in fit does result from adding storage directly to the model in the third panel of Table 3. Statistical results are available on request.

7 Concluding remarks

In this paper we have presented a competitive storage model of commodity prices with storage frictions. By introducing a degree of irreversibility into storage decisions, the friction in our model effectively embeds a valuable real option in each unit of the commodity that is held in storage. A convenience yield appears simply because the usual calculation of the expected return from storage mis-measures the value of the stored commodity by ignoring the expected capital gain in the embedded option made available by storage, as well as the opportunity cost of holding this option. If the embedded option is included when calculating the expected return from storage there is no expected return shortfall; if it is omitted, there can be a (positive or negative) expected return shortfall and therefore a nonzero convenience yield.

Our formulation of the model in continuous time makes it relatively straightforward to simulate high-frequency data. We used such data to investigate the implications of storage frictions for the behavior of commodity spot and forward prices, as well as the market value of the stored commodity, and found that frictions have a subtle, but important, direct impact on price dynamics. For example, frictions cause the spot price to experience intermittent periods of very high volatility, corresponding to situations when speculators withdraw from the spot market. When we fitted various GARCH(1,1) models to our simulated spot and forward price data we found similar behavior to that observed in actual price data by numerous authors, provided the friction is present. Similarly, although Samuelson's hypothesis — that spot prices are more volatile than forward prices — holds in the frictionless case, it no longer holds when frictions are present. In particular, during moderately good harvests, forward prices are more volatile than the spot price.

The ease with which simulated price and storage data are able to be generated using this model suggests that a rich array of predictions can be made. Our investigation of price volatility — specifically, into GARCH behavior and Samuelson's hypothesis — is representative of what can be done, but many other avenues can be explored. It may also be possible to calibrate the model, in much the same way as Chambers and Bailey (1996), Deaton and Laroque (1992, 1996), and Routledge et al. (2000) have calibrated the frictionless model, in order to estimate the magnitude of storage frictions. The calibration process would be eased if reliable high-frequency storage data were available in addition to the widely-available price data. This suggests that emerging commodity markets, such as those for electricity, might be ideal test-beds for models such as the one presented in this paper.

Our analysis shows that transaction costs materially affect decisions and market outcomes, including price processes. Although we have focussed on commodity markets, we are hopeful that the ideas presented in this paper can be applied to other markets in which frictions and speculation play a role. The real estate market is an obvious candidate, but even financial markets, with their transaction costs, may be able to be analyzed using variants of the model presented in this paper.

References

Anderson, R. and Danthine, J.-P. (1983). The time pattern of hedging and the volatility of futures prices. *Review of Economic Studies*, 50:249–266.

Baillie, R. T. and Myers, R. J. (1991). Bivariate GARCH estimation of the optimal commodity futures hedge. *Journal of Applied Econometrics*, 6(2):109–124.

Bessembinder, H., Coughenour, J. F., Seguin, P. J., and Smeller, M. M. (1996). Is there a term structure of futures volatilities? Reevaluating the Samuelson hypothesis. *Journal of Derivatives*, 4:45–58.

Brennan, D., Williams, J., and Wright, B. D. (1997). Convenience yield without the convenience: A spatial-temporal interpretation of storage under backwardation. *Economic Journal*, 107(443):1009–1022.

Brennan, M. J. (1958). The supply of storage. *American Economic Review*, 48(1):50–72.

Carlson, M., Khokker, Z., and Titman, S. (2006). Equilibrium exhaustible resource price dynamics. *Journal of Finance*, forthcoming.

Casassus, J., Collin-Dufresne, P., and Routledge, B. R. (2005). Equilibrium commodity prices with irreversible investment and non-linear technology. *NBER Working Paper 11864*.

Chambers, M. J. and Bailey, R. E. (1996). A theory of commodity price fluctuations. *Journal of Political Economy*, 104(5):924–957.

Cox, J. C., Ingersoll, J. E., and Ross, S. A. (1985). A theory of the term structure of interest rates. *Econometrica*, 53:385–407.

Deaton, A. and Laroque, G. (1992). On the behavior of commodity prices. *Review of Economic Studies*, 59(1):1–23.

Deaton, A. and Laroque, G. (1996). Competitive storage and commodity price dynamics. *Journal of Political Economy*, 104(5):826–923.

Fama, E. F. and French, K. R. (1988). Business cycles and the behavior of metals prices. *Journal of Finance*, 43(5):1075–1093.

Frechette, D. L. and Fackler, P. L. (1999). What causes commodity price backwardation? *American Journal of Agricultural Economics*, 81(4):761–771.

Heaney, R. (2006). An empirical analysis of commodity pricing. *Journal of Futures Markets*, 26(4):391–415.

Heinkel, R., Howe, M. E., and Hughes, J. S. (1990). Volatility and commodity price dynamics. *Journal of Futures Markets*, 10(5):519–533.

Hong, H. (2000). A model of returns and trading in futures markets. *Journal of Finance*, 55(2):959–988.

Hull, J. C. (2005). *Fundamentals of Futures and Options Markets*. Pearson Education, Upper Saddle River, New Jersey, 5th edition.

Kaldor, N. (1939). Speculation and economic stability. *Review of Economic Studies*, 7(1):1–27.

Kawai, M. (1983). Price volatility of storable commodities under rational expectations in spot and futures markets. *International Economic Review*, 24(2):435–459.

Kogan, L., Livdan, D., and Yaron, A. (2005). Futures prices in a production economy with investment constraints. Technical report, Wharton School, University of Pennsylvania.

Kohn, M. (1978). Competitive speculation. *Econometrica*, 46(5):1061–1076.

Litzenberger, R. H. and Rabinowitz, N. (1995). Backwardation in oil futures markets: Theory and empirical evidence. *Journal of Finance*, 50(5):1517–1545.

Michaelides, A. and Ng, S. (2000). Estimating the rational expectations model of speculative storage: A Monte Carlo comparison of three simulation techniques. *Journal of Econometrics*, 96:231–266.

Muth, J. F. (1961). Rational expectations and the theory of price movements. *Econometrica*, 29(3):315–335.

Ng, S. (1996). Looking for evidence of speculative stockholding in commodity markets. *Journal of Economic Dynamics and Control*, 20:123–143.

Ng, V. K. and Pirrong, S. C. (1994). Fundamentals and volatility: Storage, spreads, and the dynamics of metals prices. *Journal of Business*, 67(2):203–230.

Ng, V. K. and Pirrong, S. C. (1996). Price dynamics in refined petroleum spot and futures markets. *Journal of Empirical Finance*, 2:359–388.

Richard, S. F. and Sundaresan, M. (1981). A continuous time equilibrium model of forward prices and futures prices in a multigood economy. *Journal of Financial Economics*, 9:347–371.

Routledge, B. R., Seppi, D. J., and Spatt, C. S. (2000). Equilibrium forward curves for commodities. *Journal of Finance*, 55(3):1297–1338.

Samuelson, P. A. (1965). Proof that properly anticipated prices fluctuate randomly. *Industrial Management Review*, 6(2):41–49.

Samuelson, P. A. (1971). Stochastic speculative price. *Proceedings of the National Academy of Sciences*, 68(2):335–337.

Scheinkman, J. A. and Schechtman, J. (1983). A simple competitive model with production and storage. *Review of Economic Studies*, 50(3):427–441.

Telser, L. G. (1958). Futures trading and the storage of cotton and wheat. *Journal of Political Economy*, 66(3):233–255.

Trigeorgis, L. (1996). *Real Options: Managerial Flexibility and Strategy in Resource Allocation.* MIT Press, Massachusetts.

Working, H. (1948). Theory of the inverse carrying charge in futures markets. *Journal of Farm Economics*, 30(1):1–28.

Working, H. (1949). The theory of price of storage. *American Economic Review*, 39(6):1254–1262.

Wright, B. D. and Williams, J. C. (1982). The economic role of commodity storage. *Economic Journal*, 92(367):596–614.

Wright, B. D. and Williams, J. C. (1984). The welfare effects of the introduction of storage. *Quarterly Journal of Economics*, 99(1):169–192.

Wright, B. D. and Williams, J. C. (1989). A theory of negative prices for storage. *Journal of Futures Markets*, 9(1):1–13.

Yang, S.-R. and Brorsen, B. W. (1992). Nonlinear dynamics of daily cash prices. *American Journal of Agricultural Economics*, 74(3):706–715.

Appendices

A Proofs

A.1 Proof of Proposition 1

Since total consumption of the commodity equals $y + z$, consumers' surplus equals

$$CS(y, z) = \int_0^{y+z} (\psi(q) - \psi(y + z))\, dq,$$

producers' surplus equals $y\psi(y + z)$ and speculators' surplus equals $z\psi(y + z)$.[41] The flow of total surplus is therefore

$$TS(z; s, y) = \int_0^{y+z} (\psi(q) - \psi(y + z))\, dq + y\psi(y + z) + z\psi(y + z) = \int_0^{y+z} \psi(q)\, dq.$$

[41] We ignore production costs because the harvest is exogenous.

A.2 Proof of Proposition 2

The marginal benefit of sales from storage equals

$$\psi(y+z) - (1-k_i)\frac{\partial W}{\partial s}$$

when $z < 0$ and

$$\psi(y+z) - (1+k_o)\frac{\partial W}{\partial s}$$

when $z \geq 0$.

The simpler of the two cases in the proposition is where $s = 0$, since then we only need to consider the behavior of the marginal benefit of sales from storage over the interval $z \in [-y, 0]$. If $\psi(0) < (1-k_i)\frac{\partial W}{\partial s}$ then the marginal benefit of sales from storage is negative for all $z \in [-y, 0)$, so that the social planner should choose $z^* = -y$. If $\psi(y) > (1-k_i)\frac{\partial W}{\partial s}$ then the marginal benefit of sales from storage is positive for all $z \in [-y, 0]$, so that the social planner should choose $z^* = 0$. In all other cases, the social planner should choose $z^* \in [-y, 0]$ defined implicitly by

$$\psi(y+z^*) = (1-k_i)\frac{\partial W}{\partial s}.$$

Now we turn to the case where $s > 0$. This is more complicated, as we need to consider the behavior of the marginal benefit of sales from storage over both the intervals $z \in [-y, 0]$ and $z \in [0, \infty)$. If $\psi(0) < (1-k_i)\frac{\partial W}{\partial s}$ then the marginal benefit of sales from storage is negative for all $z \in [-y, \infty)$, so that the social planner should choose $z^* = -y$. If

$$\psi(y) \leq (1-k_i)\frac{\partial W}{\partial s} \leq \psi(0)$$

then the marginal benefit of sales from storage is zero for some $z^* \in [-y, 0]$, so that the social planner should choose this z^*, which is defined implicitly by

$$\psi(y+z^*) = (1-k_i)\frac{\partial W}{\partial s}.$$

If

$$(1-k_i)\frac{\partial W}{\partial s} < \psi(y) < (1+k_o)\frac{\partial W}{\partial s}$$

then the marginal benefit of sales from storage is positive for any allowable $z < 0$ and negative for all $z > 0$, implying that $z^* = 0$ is socially optimal. Finally, if $\psi(y) \geq (1+k_o)\frac{\partial W}{\partial s}$ then the marginal benefit of sales from storage is zero for some $z^* \geq 0$, so that the social planner should choose this z^*, which is defined implicitly by

$$\psi(y+z^*) = (1+k_o)\frac{\partial W}{\partial s}.$$

A.3 Proof of Proposition 3

We begin by proving that if the market value of a storage operator equals $x\frac{\partial W}{\partial s}$ then $w = xz^*/s$ is an optimal storage policy. We then show that if a storage operator adopts this policy then its market value equals $x\frac{\partial W}{\partial s}$.

Suppose that the market value of a storage operator equals $G(x; s, y) = x\frac{\partial W}{\partial s}$ whenever it holds x units of the commodity in storage. If the firm sells $w\,dt$ units of the commodity over the next interval of time lasting dt years, then it has current market value

$$wp\,dt + e^{-r\,dt}E^*[G(x + dx; s + ds, y + dy)],$$

which reduces to

$$G(x; s, y) + \left(wp - (\pi(w) + \varepsilon x)\frac{\partial G}{\partial x} - (\pi(z^*) + \varepsilon s)\frac{\partial G}{\partial s} + (\nu - \lambda\phi)\frac{\partial G}{\partial y} + \frac{1}{2}\phi^2\frac{\partial^2 G}{\partial y^2} - rG\right) dt.$$

It therefore chooses w in order to maximize

$$wp - \pi(w)\frac{\partial G}{\partial x} = wp - \pi(w)\frac{\partial W}{\partial s}.$$

From Proposition 2, if $z^*(s, y) < 0$ then $p = (1 - k_i)\frac{\partial W}{\partial s}$ and the firm chooses w in order to maximize

$$(w(1 - k_i) - \pi(w))\frac{\partial W}{\partial s} = \begin{cases} 0 & \text{if } w < 0, \\ -(k_i + k_o)w\frac{\partial W}{\partial s} & \text{if } w \geq 0. \end{cases}$$

It follows that any $w^* < 0$ is optimal when $z^*(s, y) < 0$. Similarly, if $z^*(s, y) > 0$ then $p = (1 + k_o)\frac{\partial W}{\partial s}$ and the firm chooses w in order to maximize

$$(w(1 + k_o) - \pi(w))\frac{\partial W}{\partial s} = \begin{cases} (k_i + k_o)w\frac{\partial W}{\partial s} & \text{if } w < 0, \\ 0 & \text{if } w \geq 0. \end{cases}$$

It follows that any $w^* > 0$ is optimal when $z^*(s, y) > 0$. Finally, if $z^*(s, y) = 0$ then $(1-k_i)\frac{\partial W}{\partial s} \leq p \leq (1 + k_o)\frac{\partial W}{\partial s}$ and the firm chooses w in order to maximize

$$wp - \pi(w)\frac{\partial W}{\partial s} = \begin{cases} (p - (1 - k_i)\frac{\partial W}{\partial s})w & \text{if } w < 0, \\ (p - (1 + k_o)\frac{\partial W}{\partial s})w & \text{if } w \geq 0. \end{cases}$$

Since this function is increasing in w for $w < 0$ and decreasing in w for $w \geq 0$, it follows that $w^* = 0$ is optimal when $z^*(s, y) = 0$. Optimality of the proposed storage policy $w = xz^*/s$ follows immediately.

Now we value a firm that manages its storage facility according to the policy described by $w^* = xz^*/s$. Since it sells $w^*\,dt$ units of the commodity over the next interval of time lasting dt years, its market value, G, satisfies

$$G(x; s, y) = w^*p\,dt + e^{-r\,dt}E^*[G(x + dx; s + ds, y + dy)],$$

which reduces to

$$0 = \frac{x}{s}z^*p - \left(\frac{x}{s}\pi(z^*) + \varepsilon x\right)\frac{\partial G}{\partial x} - (\pi(z^*) + \varepsilon s)\frac{\partial G}{\partial s} + (\nu - \lambda\phi)\frac{\partial G}{\partial y} + \frac{1}{2}\phi^2\frac{\partial^2 G}{\partial y^2} - rG. \quad \text{(A-1)}$$

The form of the storage policy, together with the constant returns to scale of the storage technology, implies that the market value of the firm will equal

$$G(x; s, y) = xV(s, y)$$

for some function V to be determined. Substituting this expression for G into (A-1) shows that V must satisfy

$$0 = -(\pi(z^*) + \varepsilon s)\frac{\partial V}{\partial s} + (\nu - \lambda\phi)\frac{\partial V}{\partial y} + \frac{1}{2}\phi^2\frac{\partial^2 V}{\partial y^2} - (r+\varepsilon)V + \frac{z^*}{s}\left(p - \frac{\pi(z^*)}{z^*}V\right).$$

The form of π allows us to rewrite this as

$$0 = -(\pi(z^*)+\varepsilon s)\frac{\partial V}{\partial s}+(\nu-\lambda\phi)\frac{\partial V}{\partial y}+\frac{1}{2}\phi^2\frac{\partial^2 V}{\partial y^2}-(r+\varepsilon)V+\frac{z^*\pi'(z^*)}{s}\left(\frac{\partial W}{\partial s} - V\right)+\frac{z^*}{s}\left(p - \pi'(z^*)\frac{\partial W}{\partial s}\right).$$

The last term on the right hand side is identically equal to zero.[42] Therefore V is determined by

$$0 = -(\pi(z^*) + \varepsilon s)\frac{\partial V}{\partial s} + (\nu - \lambda\phi)\frac{\partial V}{\partial y} + \frac{1}{2}\phi^2\frac{\partial^2 V}{\partial y^2} - (r+\varepsilon)V + \frac{z^*}{s}\pi'(z^*)\left(\frac{\partial W}{\partial s} - V\right).$$

It is straightforward to show that $V = \frac{\partial W}{\partial s}$ is a solution to this equation.

A.4 Proof of Proposition 4

The discussion prior to Proposition 4 in the text confirms that $P(s,y)$ is a competitive equilibrium spot price and that the market value of each unit of the commodity held in storage equals $\frac{\partial W}{\partial s}$. Proposition 2 shows that when $z^*(s,y) < 0$

$$V(s,y) = \frac{\partial W}{\partial s} = \frac{\psi(y + z^*(s,y))}{1 - k_i} = \frac{P(s,y)}{1-k_i},$$

and when $z^*(s,y) > 0$

$$V(s,y) = \frac{\partial W}{\partial s} = \frac{\psi(y + z^*(s,y))}{1 + k_o} = \frac{P(s,y)}{1+k_o}.$$

A.5 Proof of Proposition 5

Differentiating equation (5) with respect to s shows that $V = \frac{\partial W}{\partial s}$ satisfies

$$0 = -(\pi(z^*(s,y)) + \varepsilon s)\frac{\partial V}{\partial s} + (\nu - \lambda\phi)\frac{\partial V}{\partial y} + \frac{1}{2}\phi^2\frac{\partial^2 V}{\partial y^2} - (r+\varepsilon)V + \theta,$$

where

$$\theta = \left(\psi(y + z^*(s,y)) - \pi'(z^*(s,y))\frac{\partial W}{\partial s}\right)\frac{\partial z^*}{\partial s}.$$

Note that (i) $\frac{\partial z^*}{\partial s} = 0$ in the region where $z^* = 0$ and (ii) the term in large brackets equals zero in the region where $z^* \neq 0$. Thus $\theta = 0$ for all (s,y), implying that

$$\frac{E^*[dV]}{dt} - (r+\varepsilon)V = -(\pi(z^*) + \varepsilon s)\frac{\partial V}{\partial s} + (\nu - \lambda\phi)\frac{\partial V}{\partial y} + \frac{1}{2}\phi^2\frac{\partial^2 V}{\partial y^2} - (r+\varepsilon)V = 0.$$

[42] This is obviously the case when $z^* = 0$. Moreover, Proposition 2 shows that $p = \pi'(z^*)\frac{\partial W}{\partial s}$ whenever $z^* \neq 0$, so that the last term vanishes whenever $z^* \neq 0$ as well.

A.6 Proof of Proposition 6

The convenience yield equals

$$CY = (r+\varepsilon)P - \frac{E^*[dP]}{dt}.$$

When $z^*(s,y) = 0$, the spot price is $P(s,y) = \psi(y)$ and Itô's Lemma implies that

$$CY = (r+\varepsilon)\psi(y) - (\nu - \lambda\phi)\psi'(y) - \frac{1}{2}\phi^2\psi''(y).$$

When $z^*(s,y) > 0$, Proposition 2 shows that the spot price is

$$P(s,y) = \psi(y+z^*) = (1+k_o)\frac{\partial W}{\partial s}.$$

Itô's Lemma implies that

$$\frac{E^*[dP]}{dt} = (1+k_o)\left(-(\pi(z^*)+\varepsilon s)\frac{\partial^2 W}{\partial s^2} + (\nu - \lambda\phi)\frac{\partial^2 W}{\partial s \partial y} + \frac{1}{2}\phi^2\frac{\partial^3 W}{\partial s \partial y^2}\right).$$

The Hamilton-Jacobi-Bellman equation, (5), implies that

$$0 = -(\pi(z^*)+\varepsilon s)\frac{\partial W}{\partial s} + (\nu - \lambda\phi)\frac{\partial W}{\partial y} + \frac{1}{2}\phi^2\frac{\partial^2 W}{\partial y^2} - rW + TS(z^*; y).$$

Differentiating with respect to s shows that

$$0 = -(\pi(z^*)+\varepsilon s)\frac{\partial^2 W}{\partial s^2} + (\nu - \lambda\phi)\frac{\partial^2 W}{\partial s \partial y} + \frac{1}{2}\phi^2\frac{\partial^3 W}{\partial s \partial y^2} - (r+\varepsilon)\frac{\partial W}{\partial s}.$$

where we have used the Envelope Theorem to remove the terms involving $\frac{\partial z^*}{\partial s}$. It follows that

$$\begin{aligned}
\frac{E^*[dP]}{dt} &= (1+k_o)\left(-(\pi(z^*)+\varepsilon s)\frac{\partial^2 W}{\partial s^2} + (\nu - \lambda\phi)\frac{\partial^2 W}{\partial s \partial y} + \frac{1}{2}\phi^2\frac{\partial^3 W}{\partial s \partial y^2}\right) \\
&= (1+k_o)(r+\varepsilon)\frac{\partial W}{\partial s} \\
&= (r+\varepsilon)P,
\end{aligned}$$

whence $CY = 0$. The case when $z^*(s,y) < 0$ can be treated in the same way.

A.7 Proof of Proposition 7

Fix the delivery date of the forward contract at T, and allow the current date t to vary. Using the definition of B in equation (17), it is straightforward to show that

$$\begin{aligned}
\frac{E^*[dB]}{dt} - (r+\varepsilon)B &= \frac{E^*[dP]}{dt} - e^{-(r+\varepsilon)(T-t)}\frac{E^*[dF]}{dt} \\
&\quad - (r+\varepsilon)e^{-(r+\varepsilon)(T-t)}F - (r+\varepsilon)\left(P - e^{-(r+\varepsilon)(T-t)}F\right) \\
&= \frac{E^*[dP]}{dt} - (r+\varepsilon)P - e^{-(r+\varepsilon)(T-t)}\frac{E^*[dF]}{dt}.
\end{aligned}$$

Since $E^*[dF]/dt = 0$, it follows from equation (13) that

$$0 = \frac{E^*[dB]}{dt} - (r+\varepsilon)B + CY.$$

When combined with the initial condition

$$B(0,s,y) = P(s,y) - F(0,s,y) = 0$$

this implies the result in Proposition 7.

B Numerical solution method

We use policy iteration to solve for the social planner's optimal storage policy on a discrete grid in (s, y)-space. The starting point is an initial guess for the social planner's storage policy, denoted $z^{(0)}(s, y)$, which we calculate using Proposition 2 for the case where $\frac{\partial W}{\partial s} = 1$ (that is, we set $\frac{\partial W}{\partial s}$ equal to the price when consumption equals the long-run average harvest). Given the storage policy $z^{(n)}(s, y)$, we solve the finite difference approximation of the partial differential equation (9) after making the substitution $z^* = z^{(n)}$. We use the central difference to approximate $\frac{\partial W}{\partial y}$, while the approximation for $\frac{\partial W}{\partial s}$ depends on the sign of $\pi(z^{(n)}(s, y)) + \varepsilon s$. We use the backward difference if $\pi(z^{(n)}(s, y)) + \varepsilon s > 0$ and the forward difference otherwise. We impose numerical boundary conditions along the boundaries of the grid. Once we have found the finite difference approximation to the social planner's objective function, which we denote $W^{(n)}(s, y)$, we use Proposition 2 to calculate the new storage policy $z^{(n+1)}(s, y)$. We repeat this sequence of steps until convergence occurs.

Once we have found the social planner's storage policy, we calculate the spot price directly using $P(s, y) = \psi(y + z^*(s, y))$. The extent of backwardation, $B(\tau, s, y)$, is calculated by solving the partial differential equation corresponding to the present value of the flow of the convenience yield (see Proposition 7) using a finite difference method. The forward price is then calculated directly using

$$F(\tau, s, y) = e^{(r+\varepsilon)\tau}(P(s, y) - B(\tau, s, y)).$$