Commodity Prices and the Option Value of Storage

Lewis Evans

Victoria University of Wellington

Graeme Guthrie*

Victoria University of Wellington

April 26, 2010

Abstract

We incorporate a friction into the standard competitive storage model of commodity prices and derive equilibrium storage policies and spot prices. The friction introduces an element of irreversibility to storage decisions, which leads to periods when storage operators do not trade in the spot market and spot-price volatility is substantially greater than normal. It also drives a wedge between the spot price and the market value of the stored commodity. When the return on storage is correctly measured using the market value of the stored commodity, instead of the spot price, the convenience yield is zero except during stock-outs. However, when the return from storage is incorrectly calculated using the spot price, the convenience yield is nonzero in the no-trade region.

JEL Classification code: D51, G13, Q1, Q4

Keywords: commodity prices, convenience yield, real options, frictions

1 Introduction

The theoretical framework generally used to analyze commodity price behavior is a competitive storage model in which speculators buy a commodity from producers and store it for subsequent sale to consumers.[1] The key driver of the behavior in this model is the inability of inventory to be negative. The standard model is frictionless in all other respects. In particular, there are no costs associated with moving the commodity into or out of storage (although there are usually ongoing costs of holding the commodity in storage for a period of time). However, commodities

---

*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

[1] See Muth (1961), Samuelson (1971), Kohn (1978), Wright and Williams (1982, 1984), and Scheinkman and Schechtman (1983) on the existence and uniqueness of rational expectations equilibria and Deaton and Laroque (1992, 1996), Chambers and Bailey (1996), and Routledge et al. (2000) on the empirical performance of competitive equilibrium storage models.

are physical (not financial) assets, and transporting them between spot markets and storage facilities is costly. This paper incorporates these costs into the standard competitive storage model and identifies some of the implications for spot prices and storage policies.

In our model a cost is incurred each time a unit of the commodity is moved into storage.[2] This friction introduces an element of irreversibility into storage decisions—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). As a result, the real option to delay selling the stored commodity can be valuable, and this option drives a wedge between the spot price and the market value of the stored commodity.[3] A storage operator that wishes to increase its inventory by one unit must buy that unit in the spot market, as well as incur the cost of moving it into storage. In return, the operator receives one unit of the commodity ready for immediate sale, as well as the real option to wait and sell it at some future date. It is optimal for the operator to buy the commodity in the spot market only when the value of this delay option is at least as great as the cost of moving the commodity into storage. In contrast, a storage operator that wishes to reduce its inventory by one unit must sell that unit in the spot market. Because such a sale destroys the real option to wait and sell at a later date, it is optimal for the operator to sell the commodity in the spot market only when this delay option is worthless. Thus, storage operators will be buyers in the spot market when the delay option's value is above some strictly positive threshold, and they will be sellers when its value is zero. When the value of the delay option lies anywhere between these two thresholds, storage operators do not trade in the spot market.[4]

There is a single explicit market in our model, the spot market for the commodity. However, the model features an additional market, the financial market for ownership of the storage operators themselves. It follows that there are actually two goods in our model—the commodity being traded in the spot market and the commodity being held in storage—with the market value of the stored commodity being determined in the market for ownership of storage operators. This ensures that the expected growth rate in the market value of the stored commodity, adjusted for ongoing costs of holding the commodity in storage, equals the appropriate risk-adjusted discount rate.[5] As long as storage operators are trading in the spot market, the spot price will inherit this property. However, whenever storage operators do not participate in the spot market, the

---

[2] Our theory also applies when it is costly to remove the commodity from storage.

[3] Heinkel et al. (1990) analyze the call option associated with holding a commodity in storage, but the decision to release the commodity from storage is completely irreversible. In Routledge et al. (2000), ownership of the stored commodity confers a valuable timing option that can be exercised during stock-outs.

[4] Constantinides (1986), Dumas and Luciano (1991), and Chavas et al. (2000) derive optimal trading strategies for price-taking individuals facing transaction costs, and show that no-trade situations result. However, they do not derive equilibrium prices.

[5] We will assume that all individuals are risk neutral, so that the expected growth rate in the market value of the stored commodity, adjusted for the ongoing cost of storage, equals the risk-free interest rate. However, relaxing this assumption is straightforward, requiring only that we replace the actual process for the harvest in our model with its "risk-neutral" counterpart.

items traded there are created and consumed instantaneously, so that there are no intertemporal restrictions on the spot price: the expected growth rate in the spot price, adjusted for ongoing storage costs, almost always deviates from the risk-adjusted discount rate.

Researchers have devoted a great deal of effort to understanding the expected return from storage. The return is typically calculated as the sum of the expected growth rate in the spot price and the so-called convenience yield, which is a flow of nonpecuniary benefits that firms receive as long as they store the commodity (Kaldor, 1939; Working, 1948, 1949). The convenience yield is believed to compensate investors for any expected rate of return shortfalls. 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. One flaw with these motivations of the convenience yield, which we address in this paper, is that they involve benefits that are received only 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." However, storing the commodity involves having the cake and *not* eating it.

In our model, the owner of the stored commodity receives no benefits—pecuniary or otherwise—other than the capital gain. We show that, provided the rate of return is calculated using the market value of the stored commodity, the expected rate of return from storage equals the risk-adjusted discount rate. However, if the spot price is used instead, then the expected rate of return from storage will deviate from the risk-adjusted discount rate when storage operators withdraw from the spot market. The difference between the expected and required rates of return would normally be attributed to a convenience yield. Thus, our model generates an endogenous convenience yield that is the result of mis-measuring the return from storage and is nonzero only when storage operators are not active in the spot market.[6]

At all times, the market value of the stored commodity equals the sum of the spot price and the market value of the real option to delay selling the stored commodity in the spot market. The expected rate of return on storage is therefore a weighted average of the expected growth rates in the spot price and the option value, adjusted for ongoing storage costs. Equilibrium in the market for ownership of the stored commodity ensures that the (adjusted) weighted average

---

[6] Wright and Williams (1989) offer another explanation for positive inventory in the presence of low expected returns to storage that is based on mis-measurement. They 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 higher 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 appear in the data as positive industry-wide storage with a negative expected return to storage.

always equals the risk-adjusted discount rate, so if the (adjusted) expected growth rate in the spot price is lower than this, the (adjusted) expected growth rate in the option value must be higher. Because the convenience yield equals the amount by which the (adjusted) expected growth rate in the spot price falls short of the risk-adjusted discount rate, it must also equal the amount by which the (adjusted) expected growth rate in the option value exceeds this level. That is, the convenience yield can be interpreted as the expected excess return on the real option to delay selling the stored commodity.

We believe that this interpretation of the convenience yield can be extended beyond our particular model set-up. 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. Rather than equalling a flow of benefits received during the period over which the return from storage is being calculated—which is the standard interpretation—we argue that it actually represents changes in the present value of benefits that will be received only some time after the measurement period, when the commodity is released from storage. Thus, storage operators are not having their cake and eating it too. Instead, they are having their cake in order that it may be eaten in the future.

Although much of our discussion centers on the expected value of changes in the spot price, the no-trade region also leads to interesting behavior in the standard deviation of spot-price changes. As long as storage operators are trading in the spot market, they will help dampen shocks (in our model, supply shocks). However, when storage operators 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 storage operators are trading in the spot market, interrupted by relatively brief periods of much higher volatility when they do not participate. When the friction is omitted from the model, the no-trade region does not arise and simulated prices exhibit little heteroscedasticity.

Our paper adds to a growing literature that uses continuous-time equilibrium models to analyze the behavior of commodity prices. Carlson et al. (2007) 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. (2009) 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. (2009) 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.

The model is introduced, and its general properties described, in Section 2. This involves first deriving a storage policy that a welfare-maximizing social planner would choose and then proving that the resulting spot price can be achieved in an equilibrium setting where individual storage operators take the spot price as given and adopt storage policies that maximize the value of their business. Section 3 analyzes price behavior, starting by comparing the spot price and the market value of the stored commodity, and then focussing on the qualitative properties of spot prices. One particular case of the model is solved numerically, allowing us to document some of the quantitative price behavior. Section 4 offers some concluding remarks. All proofs can be found in Appendix A.

2 Solving for the equilibrium commodity spot price

2.1 Model set-up

There are three types of agents: consumers, producers, and storage operators. The latter 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 storage operators can store the commodity). All agents are risk neutral and storage capacity is infinite. Let s_t denote the total quantity of the commodity held in storage at date t and suppose that inventory decays at the constant rate $\varepsilon > 0$. Each unit of the commodity purchased by storage operators increases inventory by only $1 - \kappa$ units, but the commodity can be removed from storage for sale in the spot market without any additional cost.[7] 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 storage operators sell the commodity to consumers[8] and

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

The market-clearing spot price at date t is $p_t = \psi(y_t + z_t)$, where the consumers' inverse demand function ψ satisfies $\psi' < 0$ 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. The risk-free interest rate, r, is constant.

2.2 Solving the social planner's problem

The main focus of this paper is the commodity price that results from the competitive interaction between storage operators. We will see that a relatively simple way to determine an equilibrium

---

[7] 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.

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

storage policy is to first consider the problem facing a hypothetical social planner who maximizes the present value of the incremental flow of total surplus attributable to storage, which equals[9]

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

If $z_t > 0$ (so that consumption of the commodity at date t is raised due to storage activities) then the sum of consumers' and storage operators' incremental surpluses is the increased area under the demand curve, which is the part lying between consumption levels y_t and $y_t + z_t$. In contrast, if $z_t < 0$ (so that consumption of the commodity at date t is reduced due to storage activities) then the sum of consumers' and storage operators' surpluses falls by an amount equal to the area under the demand curve between consumption levels $y_t + z_t$ and y_t.

The planner takes s and y as given and chooses the level of sales from storage $z_t = z(s_t, y_t)$ at all future dates t in order to maximize the present value of total surplus,

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

The corresponding Hamilton-Jacobi-Bellman equation is

$$0 = \max_z \; -(\pi(z) + \varepsilon s)\frac{\partial W}{\partial s} + \nu \frac{\partial W}{\partial y} + \frac{1}{2}\phi^2 \frac{\partial^2 W}{\partial y^2} - rW + TS(z; y). \tag{3}$$

The choice of z in (3) is constrained by the requirements that consumption and storage cannot be negative; that is, $y + z \geq 0$ for all (s, y) and $z \leq 0$ whenever $s = 0$. If $z^*(s, y)$ maximizes the argument on the right hand side of (3) then W satisfies

$$0 = -(\pi(z^*(s, y)) + \varepsilon s)\frac{\partial W}{\partial s} + \nu \frac{\partial W}{\partial y} + \frac{1}{2}\phi^2 \frac{\partial^2 W}{\partial y^2} - rW + \int_y^{y+z^*(s,y)} \psi(q)\, dq. \tag{4}$$

The policy function z^* and the value function W must be solved simultaneously. However, we are able to express the optimal inventory management policy, $z^*(s, y)$, in terms of the marginal social value of the stored commodity, $\frac{\partial W}{\partial s}$, by solving the maximization problem in (3).

Proposition 1 *The socially-optimal storage policy depends on the marginal social value of the stored commodity,* $\frac{\partial W}{\partial s}$*, as follows. When* $s > 0$*:*

- *if* $\frac{\partial W}{\partial s} \leq \psi(y)$ *then the planner reduces inventory, choosing* $z^* \geq 0$ *defined implicitly by* $\psi(y + z^*) = \frac{\partial W}{\partial s}$*;*
- *if* $\psi(y) < \frac{\partial W}{\partial s} \leq \frac{\psi(y)}{1-\kappa}$ *then the planner holds inventory constant, so that* $z^* = 0$*;*
- *if* $\frac{\psi(y)}{1-\kappa} < \frac{\partial W}{\partial s} \leq \frac{\psi(0)}{1-\kappa}$ *then the planner raises inventory, choosing* $z^* \in [-y, 0]$ *defined implicitly by* $\psi(y + z^*) = (1 - \kappa)\frac{\partial W}{\partial s}$*;*
- *if* $\frac{\partial W}{\partial s} > \frac{\psi(0)}{1-\kappa}$ *then the planner stores the entire harvest, so that* $z^* = -y$*.*

---

[9] 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).

When $s = 0$:

- *if* $\frac{\partial W}{\partial s} \leq \frac{\psi(y)}{1-\kappa}$ *then the planner sells the entire harvest to consumers, so that* $z^* = 0$*;*
- *if* $\frac{\psi(y)}{1-\kappa} < \frac{\partial W}{\partial s} \leq \frac{\psi(0)}{1-\kappa}$ *then the planner stores some of the harvest, choosing* $z^* \in [-y, 0]$ *defined implicitly by* $\psi(y + z^*) = (1-\kappa)\frac{\partial W}{\partial s}$*;*
- *if* $\frac{\partial W}{\partial s} > \frac{\psi(0)}{1-\kappa}$ *then the planner stores the entire harvest, so that* $z^* = -y$*.*

Proposition 1 shows that for any particular level of the harvest, there will be three distinct ranges for the marginal value of storage. If it is sufficiently low then the planner will sell the commodity out of storage at a rate such that consumers' willingness to pay for one additional unit of consumption, $\psi(y + z)$, equals the increase in overall welfare from storing one additional unit, $\frac{\partial W}{\partial s}$. In contrast, when the marginal value of storage is sufficiently high, the planner will purchase the commodity and store it at a rate such that consumers' willingness to pay for one additional unit of consumption, $\psi(y+z)$, equals the increase in overall welfare from storing $1-\kappa$ additional units, $(1-\kappa)\frac{\partial W}{\partial s}$. Between these two ranges, the marginal value of storage is not high enough to justify diverting any of the harvest from the spot market and it is not low enough to justify boosting the harvest with sales from storage.[10]

2.3 Competitive equilibrium

We now show that if $z^*(s, y)$ denotes a socially optimal storage policy then there exists a competitive equilibrium with spot price equal to

$$p_t = P(s_t, y_t) \equiv \psi(y_t + z^*(s_t, y_t)) \tag{5}$$

at date t. Our first step is to consider a firm that currently holds x units of the commodity in storage and takes the spot price process in (5) 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 at each date t, $w_t = w(x_t; s_t, y_t)$, in order to maximize

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

where its total inventory evolves according to

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

The choice of w is constrained by the requirement that $w \leq 0$ whenever $x = 0$; that is, the firm's storage cannot be negative.[11] Solutions to this equation, one of which is described in the following proposition, correspond to optimal storage policies for a price-taking storage operator.

---

[10] In the frictionless case (that is, when $\kappa = 0$), it is socially optimal for storage operators to not trade in the spot market only when either (i) inventory is positive and $\frac{\partial W}{\partial s} = \psi(y)$, or (ii) inventory is zero and $\frac{\partial W}{\partial s} \leq \psi(y)$.

[11] 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 2 *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*

$$w(x; s, y) = \frac{x}{s} z^*(s, y). \tag{6}$$

(Negative values of w correspond to spot market purchases.) The market value of such a storage operator equals $x\frac{\partial W}{\partial s}$.

If each storage operator i, who has x_t^i units of the commodity in storage, follows the policy described in Proposition 2, then total inventory is $\sum_i x_t^i = s_t$ and total sales from storage equals

$$\sum_i w_t^i = \sum_i \frac{x_t^i}{s_t} z^*(s_t, y_t) = z^*(s_t, y_t).$$

The spot price at date t would equal $\psi(y_t + z^*(s_t, y_t))$, which is the expression in (5). That is, if each firm takes the price in (5) as given and maximizes its own market value using the policy in Proposition 2, then the market-clearing spot price equals the one in (5). That is, the price in (5) can arise in a competitive equilibrium.

Proposition 3 *There exists a competitive equilibrium in which the aggregate storage policy is given in Proposition 1, the market-clearing spot price is given by equation (5), and the market value of each unit of the stored commodity is $V(s, y) \equiv \frac{\partial W}{\partial s}$.*

The form of the policy in Proposition 1 determines the behavior of aggregate storage in this equilibrium. In particular, the presence of frictions means that the periods when storage operators appear on the demand side of the spot market are separated from those when they appear on the supply side by discrete intervals. In contrast, when there are no frictions, storage operators switch from one side of the market to the other rapidly and repeatedly.

The equilibrium behavior of individual storage operators can be understood in terms of the real options associated with storage, as we explain in the next section. A storage operator that wishes to increase its inventory by one unit must buy that unit in the spot market, as well as incur the cost of moving it into storage (in our model, an amount equal to the product of $\kappa/(1-\kappa)$ and the spot price). In return, the operator receives one unit of the commodity ready for immediate sale, as well as the real option to wait and sell it at some future date. It is optimal for the operator to buy the commodity in the spot market only when the value of this delay option is at least as great as the cost of moving the commodity into storage. In contrast, a storage operator that wishes to reduce its inventory by one unit must sell that unit in the spot market. Because such a sale destroys the real option to wait and sell at a later date, it is optimal for the operator to sell the commodity in the spot market only when this delay option is worthless. Thus, storage operators will be buyers in the spot market when the delay option's value is above some strictly positive threshold, and they will be sellers when its value is zero. When the value of the delay option lies anywhere between these two thresholds, storage operators do not trade in the spot market.

3 Price behavior

3.1 The real option to delay selling the stored commodity

There is a single explicit market in our model, the spot market for the commodity. However, our model has an additional market, the financial market for ownership of the storage operators themselves. Equilibrium in the latter market requires that the expected rate of return earned by storage operators equals the risk-free interest rate, which determines the market value of storage firms, as well as their behavior and ultimately the spot price. Thus there are actually two 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 stored commodity, where (from Proposition 3) the latter is equal to the marginal value that the planner attributes to inventory. These prices are not equal, but storage operators' ability to transform one good into the other means that they are closely related.

The spot price will never fall below $(1-\kappa)V(s,y)$, because if it did storage operators would demand an infinite amount of the commodity in the spot market, making an immediate gain of $(1-\kappa)V(s,y)-P(s,y)$ on each unit purchased. Except during stock-outs, it will never climb above $V(s,y)$, because if it did storage operators would instantaneously sell their total inventory in the spot market, making an immediate gain of $P(s,y)-V(s,y)$ on each unit sold. The spot price is thus constrained to a band bounded below by $(1-\kappa)V(s,y)$ and above by $V(s,y)$.[12] When storage operators are selling the commodity in the spot market, the spot price lies at the top of this band; when they are buying the commodity from producers, it lies at the bottom. However, if the spot price lies above the lower bound and below the upper one then storage operators do not trade in the spot market at all: the spot price is too high for purchasing the commodity in the spot market to be profitable, and too low for selling the stored commodity to be profitable. These results are summarized in the following proposition.

Proposition 4 *As long as inventory is positive, the market-clearing spot price satisfies*

$$(1-\kappa)V(s,y) \leq P(s,y) \leq V(s,y),$$

with $P(s,y)$ equalling the lower bound if $z^*(s,y)<0$ and the upper bound if $z^*(s,y)>0$, and otherwise equalling $P(s,y)=\psi(y)$.

It is useful to think of each unit of the stored commodity being a bundle comprising a unit of the commodity committed to immediate sale in the spot market and the real option to delay this sale until a future date of the owner's choosing. The value of the first component is the spot price $P(s,y)$ and the value of the second component, which we denote by $U(s,y)$, satisfies

$$V(s,y)=P(s,y)+U(s,y). \tag{7}$$

---

[12] The upper and lower bounds are identical if there are no transaction-cost frictions, in which case the spot price always equals $V(s,y)$.

Proposition 4 shows that $U(s, y)$ equals zero when storage operators are selling in the spot market, equals $(\kappa/(1-\kappa))P(s, y)$ when they are buying in the spot market, and is otherwise somewhere between these two bounds.

3.2 Convenience yield and the expected return to storage

The expected return from storage is typically calculated as the sum of the expected growth rate in the spot price and the so-called convenience yield, which is a flow of nonpecuniary benefits that firms receive as long as they store the commodity. The convenience yield is believed to compensate investors for any expected rate of return shortfall. In this section we analyze the convenience yield that arises endogenously in our model and we interpret it in terms of the decomposition of the stored commodity into the commodity committed for immediate sale and the real option to delay that sale.

Although the rate of return from storage is usually calculated using the spot price to value the stored commodity, Section 3.1 shows that the correct approach is to use V. Moreover, because the owner of the stored commodity receives no cash flows or benefits other than the capital gain, the total expected rate of return from storage equals the expected growth rate in V, adjusted for the decay associated with storage. Since V is the price of a traded asset—ownership of one unit of the stored commodity—equilibrium in the market for ownership of this asset will ensure that the expected rate of return from buying the stored commodity and holding it for later sale equals the risk-free interest rate. Since inventory decays at rate ε, the expected growth rate in the price must exceed the risk-free interest rate by this amount.

Proposition 5 *The expected growth rate in the price of the stored commodity equals*

$$\frac{1}{V}\frac{E[dV]}{dt} = r + \varepsilon. \tag{8}$$

Proposition 5 shows that, if the stored commodity is valued correctly, a convenience yield is not needed to equate the expected rate of return from storage to the risk-free interest rate.

As long as storage operators are active in the spot market, the spot price is proportional to V and so will share the same expected growth rate. However, in the region where $z^* = 0$, the lack of storage-operator involvement means that the spot price is simply equal to the level at which demand for consumption of the commodity equals the level of the harvest. In this case, there is no intertemporal smoothing on either the demand or supply sides of the market, and so nothing to prevent the expected growth rate in the spot price from deviating from $r + \varepsilon$. Trends in the harvest therefore flow through directly into trends in the spot price. The next proposition gives the precise form of the expected growth rate in the spot price.

Proposition 6 (Spot-price drift) *The expected growth rate in the spot price equals*

$$\frac{1}{P}\frac{E[dP]}{dt} = \begin{cases} r + \varepsilon, & \text{if } z^*(s, y) \neq 0, \\ (\nu\psi'(y) + \frac{1}{2}\phi^2\psi''(y))/\psi(y), & \text{if } z^*(s, y) = 0. \end{cases} \tag{9}$$

We will see that in the frictionless case the region where $z^*(s, y) = 0$ has measure zero during periods when inventory is positive, so that the expected growth rate in the spot price is almost always equal to $r + \varepsilon$ except during stock-outs.

In the region where storage operators are active in the spot market, the expected growth rate in the spot price is independent of the current level of the spot price: there is no short-run mean reversion in this case. However, when the spot market separates from the market for ownership of the stored commodity, the expected growth rate varies with the current level of the spot price. Given this complex behavior, tests for mean reversion using high-frequency spot-price data that are based on simple autoregressive processes, for example, are inappropriate. Alternative approaches, such as using the dynamics of the term structure of futures prices as in Bessembinder et al. (1995), may be more effective.

Proposition 6 shows that there will be situations in which the expected growth rate in the spot price differs from $r + \varepsilon$. The usual approach to dealing with this *apparent* violation of equilibrium behavior is to posit a convenience yield. It follows immediately from Proposition 6 that in our model this convenience yield equals

$$CY \equiv (r+\varepsilon)P - \frac{E[dP]}{dt} = \begin{cases} 0, & \text{if } z^*(s,y) \neq 0, \\ (r+\varepsilon)\psi(y) - \nu\psi'(y) - \frac{1}{2}\phi^2\psi''(y), & \text{if } z^*(s,y) = 0. \end{cases}$$

It is nonzero only when storage operators are not trading in the spot market. In the frictionless case, the convenience yield is almost always equal to zero as long as inventory is positive; that is, it is a stock-out phenomenon.

There are no nonpecuniary benefits from storage in our model. The only gain from holding the stored commodity is the capital gain, and when the stored commodity is valued correctly the expected capital gain, adjusted for storage-related decay, is the risk-free interest rate. Thus, the convenience yield in our model is the result of the stored commodity being misvalued when the “return” from storage is being calculated.

Nevertheless, it is possible to interpret the convenience yield as more than simply the consequence of a valuation error. Recall that the value of the stored commodity is equal to the sum of the spot price and the value of the real option to delay selling the commodity: $V = P + U$. Using this decomposition to eliminate V from equation (8) shows that

$$\frac{E[dP]}{dt} + \frac{E[dU]}{dt} = (r+\varepsilon)\left(P+U\right),$$

which can be rearranged to give

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

That is, the convenience yield is actually the expected excess return on the real option to delay selling the stored commodity. If the spot price is expected to increase at a rate less than $r + \varepsilon$ then the option will be expected to grow in value at a greater rate: the value of the bundle that is the stored commodity will grow at rate $r + \varepsilon$ on average.

3.3 Spot-price volatility

While the previous subsection focuses on the drift in the spot price, this one concentrates on spot-price volatility. Recall from Itô's Lemma that the spot price has volatility

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

If storage operators are currently active in the spot market ($z^* \neq 0$), then the effect of a positive harvest shock on the spot price will be partly offset by storage operators either buying more (if $z^* < 0$) or selling less (if $z^* > 0$). Thus, spot-price volatility will be relatively low when $z^* \neq 0$. In contrast, if storage operators are not trading in the spot market ($z^* = 0$), then a positive harvest shock will flow directly through to the spot market. Thus, spot-price volatility will be relatively high when $z^* = 0$. Moreover, although z^* is continuous at the boundaries between the regions where $z^* = 0$ and $z^* \neq 0$, $\frac{\partial z^*}{\partial y}$ will not be, so that there will be a (discontinuous) jump in volatility when the system moves between these regions.

It follows that the spot price will experience distinct periods of especially high volatility, which begin when storage operators withdraw from the spot market and end when they return. Whenever storage operators are trading in the spot market, spot price shocks will be relatively small and—as the result of the smoothing effects of storage—more persistent than harvest shocks. In contrast, whenever storage operators are absent, spot price shocks will be relatively large and less persistent than harvest shocks. In the special case where the friction is zero, the region where $z^*(s,y) = 0$ has measure zero, so that the high-volatility/low-persistence periods will be absent (except during stock-outs).

3.4 Numerical analysis

In this section we present the details of a particular implementation of our model. 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 $\psi(x) = (\mu/x)^\alpha$, where η, μ, σ, and α are positive constants. The harvest flow generated by this process is mean-reverting with unconditional mean and variance of μ and $\sigma^2\mu/(2\eta)$, respectively. Deviations from the long-run level are expected to decline at rate η, the harvest can never become negative and, provided $2\eta\mu \geq \sigma^2$, the harvest will always be positive.[13] We set $\mu = 1$, $\eta = \log 2$, and $\sigma = (\log(2^{1/2}))^{1/2}$, so that the long-run average harvest level is unity, the half-life of harvest shocks is one year, and the unconditional standard deviation of the harvest is half the unconditional mean. The units in which the price is measured are chosen so that the price is unity when consumption equals the long-run average harvest. We set $\alpha = 2$ (so that the price-elasticity of demand equals $-1/2$),

---

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

Figure 1: Equilibrium storage policy

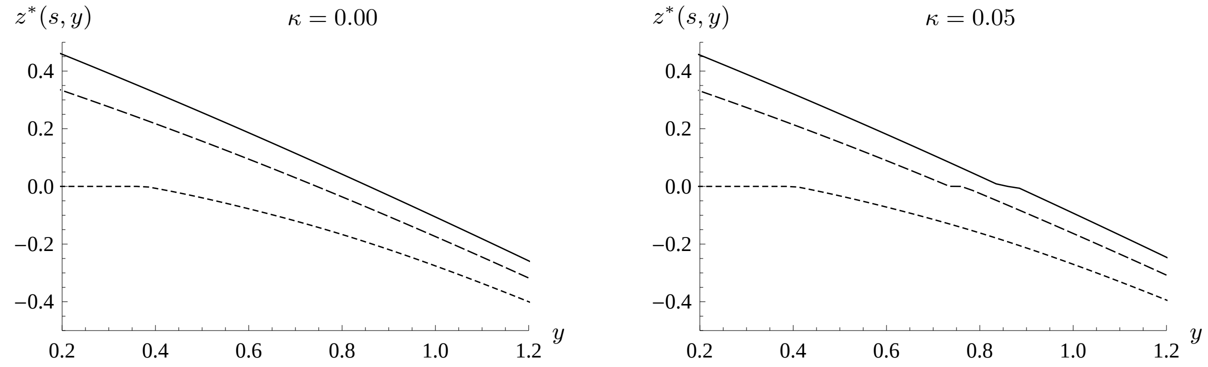

Notes. Each graph plots $z^*(s, y)$ as a function of the harvest (y) for three different levels of total storage: the dotted curves correspond to $s = 0$, the dashed curves to $s = 0.25$, and the solid curves to $s = 0.50$.

$\varepsilon = 0.03$, and $r = 0.04$. We consider two different values for the costs of moving the commodity into storage: $\kappa = 0$ and $\kappa = 0.05$.

We present the model's main outputs by plotting them as functions of the harvest level for various levels of storage. As time evolves, the harvest level changes; as long as $z^* \neq 0$, then the level of storage changes as well. However, changes in storage are of order $O(dt)$, whereas harvest changes are of order $O(\sqrt{dt})$. Thus, for the purposes of interpreting the behavior in the various graphs, it is reasonable to treat the storage level as constant for small changes in the harvest.

We begin by solving for the equilibrium storage policy using the approach described in Appendix B. Figure 1 illustrates the resulting policy for two different levels of the friction, $\kappa = 0$ in the left-hand graph and $\kappa = 0.05$ in the right-hand one. The three curves plot $z^*(s, y)$ as a function of the harvest (y) for different levels of total storage: the dotted curves correspond to $s = 0$, the dashed curves to $s = 0.25$, and the solid curves to $s = 0.50$. When the harvest is poor, storage operators sell the commodity from storage (that is, $z^* > 0$) as long as storage is positive; when the harvest is good, storage operators buy the commodity and store it (that is, $z^* < 0$); 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 right-hand graph shows that when a friction is introduced, a range of harvest values appears for which storage operators temporarily withdraw from the spot market—they neither buy nor sell the commodity until the harvest moves out of this region.

Figure 2 uses the same format as Figure 1, but plots the market-clearing spot price. 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 $z^* = 0$, corresponding to the kinks in each of the curves in Figure 2, which are aligned with the "flat spots" in Figure 1.

Figure 2: Spot price behavior

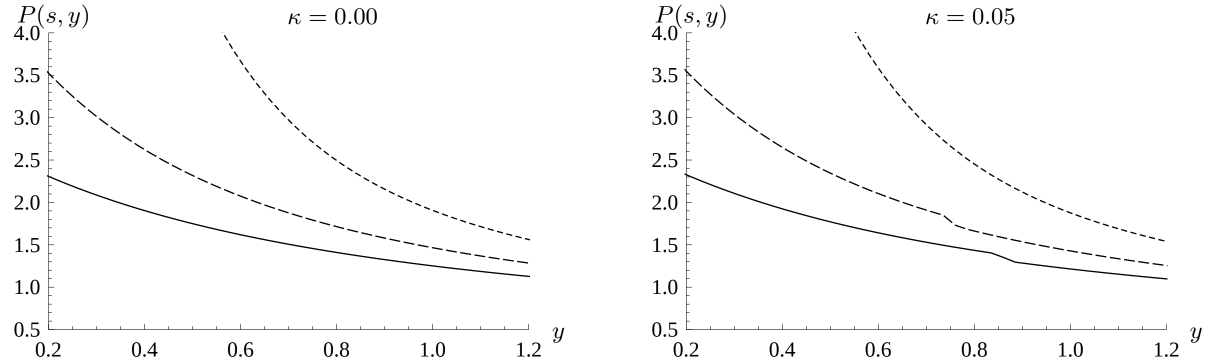

Notes. Each graph plots the spot price, $P(s,y)$, as a function of the harvest, y, for three different levels of total storage: the dotted curves correspond to $s=0$, the dashed curves to $s=0.25$, and the solid curves to $s=0.50$.

Finally, Figure 3 illustrates the behavior of the expected growth rate in the spot price (the top two graphs) and the volatility (the bottom two graphs), where the latter is expressed as a proportion of the spot price. The graphs have the same format as Figure 1. Consistent with the discussion in Sections 3.2 and 3.3, in the frictionless case, as long as inventory is positive the expected growth rate in the spot price equals $r+\varepsilon$ (so that the convenience yield is zero) and volatility is approximately constant. During a stock-out, if the harvest is sufficiently low then the inability of storage operators to offset poor harvests results in high spot-price volatility (storage operators cannot buffer harvest shocks) and expected growth rates that deviate from $r+\varepsilon$. Including the friction in the model has little impact on spot-price behavior during stock-outs. However, differences are clearly evident when inventory is positive. Although the expected growth rate still equals $r+\varepsilon$ and spot-price volatility is still low and approximately constant when storage operators are trading in the spot market, neither of these properties hold when storage operators withdraw from the spot market: the top right-hand graph shows that the expected growth rate differs from $r+\varepsilon$ and the bottom right-hand one shows that volatility increases dramatically.

4 Concluding remarks

This paper extends the standard competitive-storage model of commodity prices by incorporating a transaction cost reflecting the costs of moving a commodity into storage. We derive the conditions for a competitive equilibrium, showing that the real option to delay selling the stored commodity can be valuable and that this leads to a range of spot prices in which storage operators do not trade in the spot market. They will be buyers in the spot market when the delay option's value is at least as great as the cost of moving the commodity into storage, and

Figure 3: Spot-drift drift and volatility

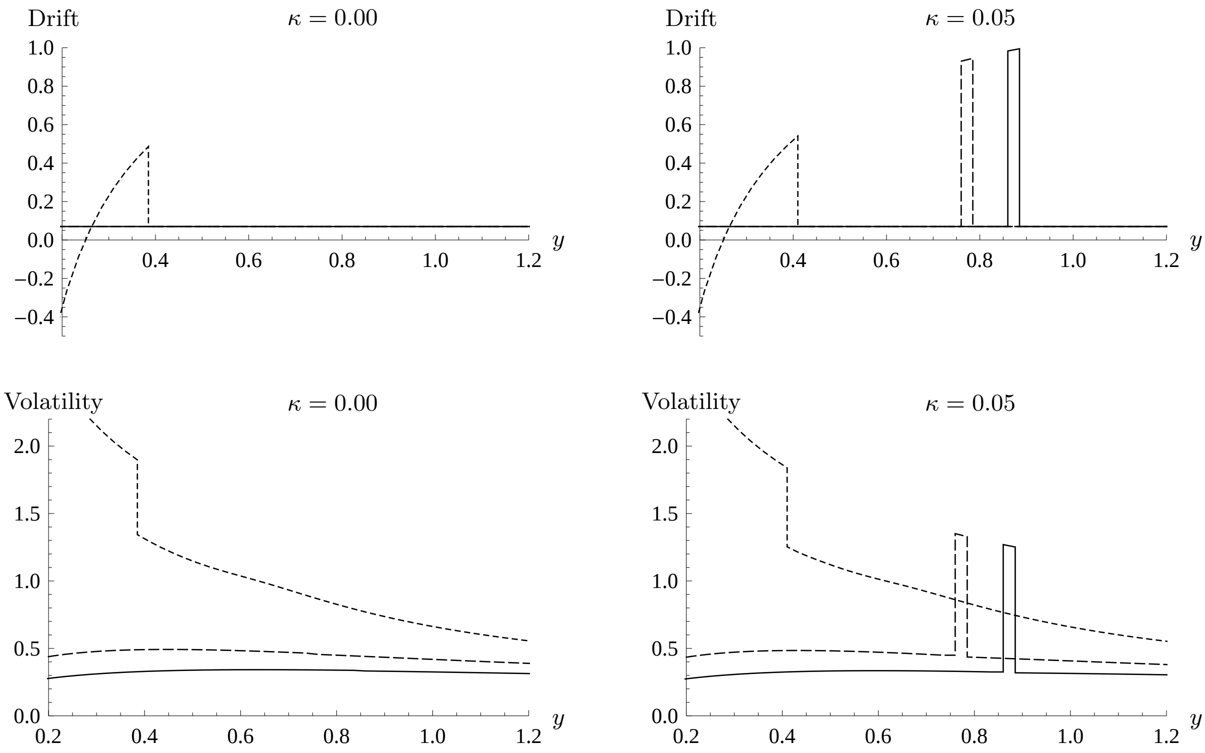

Notes. Each graph in the top panel plots the expected growth rate in the spot price as a function of the harvest, y, for three different levels of total storage: the dotted curves correspond to $s = 0$, the dashed curves to $s = 0.25$, and the solid curves to $s = 0.50$. Each graph in the bottom panel plots the instantaneous volatility of the spot price for the same three storage levels.

they will be sellers when its value is zero. When the value of the delay option lies anywhere between these two thresholds, storage operators do not trade in the spot market. It follows that there are actually two goods in our model—the commodity being traded in the spot market and the commodity being held in storage—with the market value of the stored commodity being determined in the market for ownership of storage operators. The spot price inherits the properties of the market value of the stored commodity when storage operators are active in the spot market: small and relatively persistent price shocks, and an expected growth rate (adjusted for ongoing storage costs) that equals operators' risk-adjusted discount rate. However, when storage operators are not trading in the spot market, the spot price is decoupled from the market value of the stored commodity, resulting in a period of large and relatively transient shocks to the spot price.

Although our model is framed in terms of three distinct types of agents—consumers, producers, and storage operators—our results apply to a variety of situations that better match some real-world arrangements, where the distinctions between these agents may be blurred. For example, producers might operate their own storage facilities, as is common in markets for agricultural commodities. Producer-storers will send their produce directly to the spot market when the spot price is in the no-trade region, and hold onto whatever inventory they have. Where there is no market for ownership of stand-alone storage operators, as envisaged in our model, the market value of the stored commodity will be determined in the market for ownership of producer-storers.

The results in this paper have implications for the way we think about and model commodity spot price behavior. First, in contrast to the traditional story that the convenience yield represents a flow of nonpecuniary benefits received during the period over which the return from storage is being calculated, we argue that it actually represents changes in the present value of benefits that will be received only some time after the measurement period, when the commodity is released from storage: the convenience yield is the expected excess return on the real option to delay selling the stored commodity. Moreover, if the expected rate of return on storage was measured correctly—using the market value of the stored commodity rather than the spot price—the convenience yield would be zero whenever inventory is positive.

Second, our model makes several predictions about spot price behavior that are not captured by the stochastic processes typically used when valuing commodity derivatives and other contingent claims. The simplest models assume that the commodity price follows geometric Brownian motion with a constant convenience yield, while other models use a mean-reverting stochastic process to model the behavior of the convenience yield (Schwartz, 1997; Casassus and Collin-Dufresne, 2005). Our results suggest that neither approach is consistent with a theory of equilibrium spot-price behavior incorporating frictions in the storage and trading process. For example, our results indicate that the spot price will experience periods of unusually low volatility (when the convenience yield is zero), interrupted by periods of unusually high volatility (during which the convenience yield is nonzero). This behavior is not consistent with the approaches described above, although it could be captured in a regime-switching framework. Investigation of such possibilities is left as a topic for future research.

References

Bessembinder, H., Coughenour, J. F., Seguin, P. J., and Smeller, M. M. (1995). Mean reversion in equilibrium asset prices: Evidence from the futures term structure. *Journal of Finance*, 50:361–375.

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

Carlson, M., Khokker, Z., and Titman, S. (2007). Equilibrium exhaustible resource price dynamics. *Journal of Finance*, 62(4):1663–1703.

Casassus, J. and Collin-Dufresne, P. (2005). Stochastic convenience yield implied from commodity futures and interest rates. *Journal of Finance*, 50(5):2283–2331.

Casassus, J., Collin-Dufresne, P., and Routledge, B. R. (2009). Equilibrium commodity prices with irreversible investment and non-linear technologies. Technical report, Columbia University.

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

Chavas, J.-P., Despins, P. M., and Fortenbery, T. R. (2000). Inventory dynamics under transaction costs. *American Journal of Agricultural Economics*, 82(2):260–273.

Constantinides, G. M. (1986). Capital market equilibrium with transaction costs. *Journal of Political Economy*, 94(4):842–862.

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.

Dumas, B. and Luciano, E. (1991). An exact solution to a dynamic portfolio choice problem under transaction costs. *Journal of Finance*, 46(2):577–595.

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

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

Kogan, L., Livdan, D., and Yaron, A. (2009). Oil futures prices in a production economy with investment constraints. *Journal of Finance*, 54(3):1345–1375.

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

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

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. (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.

Schwartz, E. S. (1997). The stochastic behavior of commodity prices: Implications for valuation and hedging. *Journal of Finance*, 52(3):923–973.

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

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.

Appendices

A Proofs

A.1 Proof of Proposition 1

The marginal benefit of sales from storage equals $\psi(y+z)-(1-\kappa)\frac{\partial W}{\partial s}$ when $z<0$ and $\psi(y+z)-\frac{\partial W}{\partial s}$ when $z\geq 0$.

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

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 $\psi(0)<(1-\kappa)\frac{\partial W}{\partial s}$ then the marginal benefit of sales from storage is negative for all $z\in[-y,\infty)$, so that the planner should choose $z^*=-y$. If $\psi(y)\leq(1-\kappa)\frac{\partial W}{\partial s}\leq\psi(0)$ then the marginal benefit of sales from storage is zero for some $z^*\in[-y,0]$, so that the planner should choose this z^*, which is defined implicitly by $\psi(y+z^*)=(1-\kappa)\frac{\partial W}{\partial s}$. If $(1-\kappa)\frac{\partial W}{\partial s}<\psi(y)<\frac{\partial W}{\partial s}$ 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 $\psi(y) \geq \frac{\partial W}{\partial s}$ then the marginal benefit of sales from storage is zero for some $z^* \geq 0$, so that the planner should choose this z^*, which is defined implicitly by $\psi(y + z^*) = \frac{\partial W}{\partial s}$.

A.2 Proof of Proposition 2

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, 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

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

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

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

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

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

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

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

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

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

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

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

The form of π allows us to rewrite this as

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

The last term on the right hand side is identically equal to zero,[14] so that V is determined by

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

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

A.3 Proof of Proposition 4

Proposition 1 shows that when $z^*(s, y) < 0$,

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

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

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

A.4 Proof of Proposition 5

Differentiating equation (3) with respect to s shows that $V = \frac{\partial W}{\partial s}$ satisfies

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

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

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

---

[14] This is obviously the case when $z^* = 0$. Moreover, Proposition 1 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.5 Proof of Proposition 6

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

$$\frac{E[dP]}{dt} = \nu\psi'(y) + \frac{1}{2}\phi^2\psi''(y).$$

When $z^*(s,y) > 0$, Proposition 1 shows that the spot price is $P(s,y) = \psi(y + z^*) = \frac{\partial W}{\partial s}$. Itô's Lemma implies that

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

The Hamilton-Jacobi-Bellman equation, (3), implies that

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

Differentiating with respect to s shows that

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

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

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

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

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 1 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 (4) after making the substitution $z^* = z^{(n)}$, using central differences to approximate $\frac{\partial W}{\partial y}$ and $\frac{\partial W}{\partial s}$. 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 1 to calculate the new storage policy $z^{(n+1)}(s,y)$. We repeat this sequence of steps until convergence occurs.