1

Structural Models of Commodity Prices

Craig Pirrong

1.1 INTRODUCTION

The transparency of fundamentals in commodity markets (in contrast to equity or currency markets, for instance) holds out the promise of devising structural models of commodity price behavior that can illuminate the underlying factors that drive these prices, and which perhaps can be used to value contingent claims on commodities. There has been much progress on these models in recent years, but the empirical data show that real-world commodity price behavior is far richer than that predicted by the current generation of models, and that except for non-storable commodities, structural models currently cannot be used to price derivatives. The models and empirical evidence do, however, point out the deficiencies in reduced form commodity derivative pricing models, and suggest how reduced form models must be modified to represent commodity price dynamics more realistically. They also suggest additional factors that may be added to the models (at substantial computational cost) to improve their realism.

This chapter sketches out the current state of fundamental models of commodity markets. It starts with a taxonomy of commodities, and then proceeds to discuss models for storable and non-storable commodities, and structural models for each.

1.2 A COMMODITY TAXONOMY

Although the catchall term "commodity" is widely applied to anything that is not a true asset, it conceals tremendous diversity, diversity that has material impacts on price behavior and modeling.

The most basic divide among commodities is between those that are storable, and those that are not. The most important non-storable commodity is electricity (although hydro generation does add an element of storability in some electricity markets). Weather is obviously not storable – and it is increasingly becoming an important underlying in commodity derivatives trading.

Most other commodities are storable (at some cost), but there is considerable heterogeneity among goods in this category. Some are continuously produced and consumed, and are not subject to significant seasonality in demand; industrial metals such as copper or aluminum fall into this category. Some are continuously produced and consumed, but exhibit substantial

---

Craig Pirrong, University of Houston, Houston, TX 77204. Tel: 713-743-4466. Email: cpirrong@uh.edu

Risk Management in Commodity Markets: From Shipping to Agriculturals and Energy Edited by Hélyette Geman

seasonality in demand. Heating oil, natural gas, and gasoline are prime examples of this type of commodity. Other commodities are produced seasonally, but there is also variation within the category of seasonally produced commodities. Grains and oilseeds are produced seasonally, but their production is relatively flexible because a major input – land – is quite flexible; there is a possibility of growing corn on a piece of land one year and soybeans the next, and an adverse natural event (such as a freeze) may damage one crop, but does not impair the future productivity of land. In contrast, tree crops such as cocoa or coffee or oranges are seasonally produced, but utilize specialized, durable, and inflexible inputs (the trees) and damage to these inputs can have consequences for productivity that last beyond a single crop year.

Fundamentals-based models must take these variations across commodities into account. Moreover, this cross sectional variation has empirical implications that can be exploited to test fundamental-based structural models.

1.3 FUNDAMENTAL MODELS FOR STORABLE COMMODITIES

In a nutshell, a fundamental model derives commodity prices as the equilibrium result of basic supply and demand factors. In contrast, a reduced form model merely specifies the dynamics of a commodity price (or a forward curve of commodity prices), usually in the form of a stochastic differential equation.

The Theory of Storage is the canonical fundamental commodity price model. Early versions of the theory of storage (due to Kaldor (1939) and Working (1949)) posit that commodity inventories generate a stream of benefits – a convenience yield – and that marginal convenience yield varies inversely with the level of inventories. This theory was devised to explain the fact that the forward prices of storable commodities are routinely below the spot price plus the costs of holding inventory until contract expiration. However, it is *ad hoc* and does not provide an equilibrium model of the determinants of the marginal benefit of inventory holding.

A more solidly grounded Theory of Storage is embodied in the rational expectations model of Scheinkman and Schectman (1983). In this model, a random amount of a commodity is produced every period, and competitive agents allocate production between current consumption and storage. The stored commodity can be consumed in the future. In a competitive market, the equilibrium storage decision maximizes the discounted expected utility of the representative agent. This decision depends on two state variables: current output and current inventories. In brief, agents add to inventories when production is higher than average (especially when current stocks are low) and consume inventory when production is lower than average (especially when stocks are high).

The constraint that storage cannot be negative has important pricing implications. When demand is low (and/or stocks are high) it is optimal for agents to hold inventory. In a competitive market, forward prices must cover the costs of storage in order to induce the optimal decision in these conditions, and the forward price will equal the spot price plus the cost of holding inventory to the next production date. When demand is very high and stocks are low, however, it is optimal to consume all inventories. Since storage links spot and futures prices, such a “stockout” disconnects these prices. Moreover, when it is not efficient to store, forward prices should punish storage by failing to cover by the costs of holding inventory. Indeed, if demand is sufficiently high, the equilibrium forward price during a stockout is less than the spot price. This is sometimes referred to as a “backwardation” or an inversion. Thus, in contrast to the *ad hoc* convenience yield theory, this version of the theory of storage provides a fundamentals-based, structural model of forward price structures.

As is common with rational expectations models, Scheinkman-Schectman requires a numerical solution of a dynamic programming problem. That is, it cannot be solved in closed form. However, since the problem is typically a contraction mapping, it is readily amenable to solution using standard recursive techniques. In the original Scheinkman-Schectman model, the commodity is produced every period, and the production shocks are IID. Solutions to the dynamic program are computationally cheap when there is a single independent and identically distributed (IID) demand shock. See Williams and Wright (1991) for detailed descriptions of the relevant numerical techniques.

Such a model is appropriate for a continuously produced commodity with IID demand shocks. The model implies that prices should be autocorrelated even when demand shocks are IID because storage links prices over time. When demand is low today, for instance, more of the commodity is stored, increasing future supply and thereby depressing prices in the future.

Deaton and Laroque (1992) perform empirical tests on the storage model with IID demand shocks using annual data on a variety of commodities encompassing all types of storables, including tree crops, grains, and continuously produced goods such as copper. Deaton-Laroque find that real-world commodity prices exhibit far more persistence than the storage model can generate. Storage in the presence of IID demand shocks can produce autocorrelations on the order of 20 %, far below the 90 % that Deaton-Laroque find. In later emprical work, these authors (Deaton-Laroque, 1996) attribute virtually all of the persistence in commodity prices to autocorrelation in demand shocks.

The use of annual data and the homogeneous treatment of these commodities are problematic given that the frequency of the storage decision is less than a year, and varies across these disparate commodity types. Moreover, this approach disregards an important source of valuable price information: daily data on futures prices for various maturities that are available for a wide variety of commodities.

Exploitation of high frequency futures price data, and of the cross-sectional variation in commodity characteristics, requires use of more sophisticated models. Pirrong (2006, 2007) extends the basic Scheinkman-Schectman framework to include multiple, autocorrelated demand shocks for a commodity in which the frequency of the storage decision is the same as the frequency of production.[1] This is appropriate for continuously produced commodities such as copper. Solution of such a model is substantially more computationally demanding (due to the curse of dimensionality common to dynamic programming models), but provides more realistic characterizations of storage economics, and its implications for the behavior of commodity price structures. Moreover, given the solution of the storage problem, it is possible to solve partial differential equations to determine the price of any forward contract with a maturity greater than the frequency of production.

This more complicated model can generate price behavior that mimics some of the features documented for industrial metals by Ng-Pirrong (1994). Specifically, in this model, spot prices are more volatile when the market is in backwardation (a signal of tight supply and

---

[1] Routledge, Seppi, and Spatt (2000) and Deaton-Laroque (1996) also model commodity prices when demand is autocorrelated. Pirrong permits multiple demand shocks with differing persistence.

demand conditions), and the correlation between spot prices and forward prices (e.g., a three-month forward price) is near 1 when stocks are high and the market is at full carry, and is below 1 and decreasing in the amount of backwardation (and in inventories) when the market is in backwardation. These basic behaviors are documented in Ng-Pirrong, but the more complex model calibrated to match the behavior of prices in the copper market cannot duplicate other aspects of the Ng-Pirrong empirical dynamics. For instance, in Ng-Pirrong three-month forward prices are substantially more volatile (though less volatile than spot prices) when the market is in backwardation than when it is not; in the augmented storage model, in contrast, three-month forward price volatilities vary much more weakly with backwardation (and stocks). Moreover, the augmented storage model does a poor job at explaining the dynamics of more distant forward prices, such as 15 or 27-month copper prices. Even if one of the demand shocks is highly persistent (and nearly integrated), in the storage model the long maturity forward prices exhibit virtually no variability when the market is in backwardation, whereas real-world 15 and 27-month copper prices exhibit substantial volatility. Over such a long period, a current demand shock (even if highly persistent) has little power to forecast demand in the distant future, so in the model distant forward prices do not vary in response to demand shocks.

Other extensions of the model can capture market price behaviors that are otherwise puzzling. For instance, in 2005–2006 many market commentators, and even a committee report of the United States Senate, declared that the simultaneous increase in energy prices and inventories observed during that period was symptomatic of a disconnection between market fundamentals and prices, driven by speculative excess. Pirrong (2008) modifies the basic storage model to include stochastic volatility in the net demand shock to explain this seemingly anomalous behavior.

The intuition is quite straightforward. Inventory is largely held to smooth the impact of fundamental shocks. If these shocks become more volatile, it is optimal to hold larger inventories. When market participants perceive that fundamental volatility has increased, they rationally increase inventories. This requires a reduction in consumption and, concomittantly, an increase in prices. Hence, the model predicts simultaneous increases in inventories and prices during periods of heightened risks; since the risks of hurricanes and geopolitically driven disruptions in energy production quite plausibly increased beginning in late-2005 (think Hurricanes Rita and Katrina; the Lebanon War; Iraq; turmoil in Nigeria, Venezuela, and other energy-producing regions), the model can explain the inventory and price movements that baffled so many market analysts.

Despite the modest empirical successes of the augmented storage model, empirical work that exploits the diversity of commodities points out difficulties with the received rational expectations version of the theory of storage. Specifically, Pirrong (1999), Osborne (2004), and Chambers and Bailey (1996) model seasonal commodities. In these models, storage decisions occur more frequently than production (as is realistic). Moreover (again realistically), agents receive information about the size of the future crop prior to harvest. In this model, the state variables are the current demand shock, current inventories, and information about the size of the next harvest.

This model predicts that (a) well prior to the harvest, spot prices (“old crop” prices) should exhibit little correlation with “new crop” futures prices (i.e., futures with delivery dates immediately following the harvest), and (b) information about the size of the harvest should have little impact on spot prices but a big impact on new crop prices. These predictions obtain regardless of whether demand shocks are highly persistent. The intuition behind these results is straightforward. Except under highly unusual circumstances (e.g., a large crop at the previous harvest and low demand, leading to high current inventories, combined with a forecast of an extremely short upcoming crop), agents seldom find it efficient to carry positive inventory into the new harvest: why carry supplies from when they are relatively scarce (right before the harvest), to when they are relatively abundant (immediately following the harvest)? Thus, storage cannot link "old crop" spot prices and new crop futures prices, and information that relates to the size of the new crop is largely immaterial to the price of the old crop, as it affects neither the demand for the old crop (which is driven by demand up to the time of the harvest) nor its supply (which was established at the last harvest and subsequent storage decisions).

This model predicts the differential behavior of new crop and old crop prices even if, as Deaton-Laroque posit, demand is highly autocorrelated. Thus, examining seasonals prices at weekly (rather than annual) frequency can help determine whether high demand autocorrelation is indeed the key factor in explaining the persistence of commodity price shocks.

In reality, however, there is a high correlation (typically between 90 % or higher) between old crop and new crop corn, wheat, cotton, and soybean futures prices, and both old crop and new crop prices respond by about the same amount in the same direction to official forecasts of crop size. Thus, neither storage nor high demand autocorrelation can explain the behavior of seasonal commodity futures prices. This raises doubts about the reliability of the storage model. Pirrong (1999) discusses some possible factors that can explain the evident intertemporal connections between old crop and new crop prices, including intertemporal substitution (ruled out in the basic storage model) and final goods production (e.g., soybeans are used to produce oil and meal, rather than consumed directly as the basic storage model assumes). The former explanation is somewhat *ad hoc* and difficult to test. The latter is conceptually rigorous, but increases the dimensionality of the dynamic programming problem, because it is necessary to add a state variable (final goods inventory) and solve additional equilibrium conditions (one each for the raw and final good markets). At present, the curse of dimensionality precludes sufficiently timely solution of the problem to permit rigorous empirical testing.

In sum, fundamentals-based structural rational expectations models of storable commodities shed some light on the behavior of commodity prices, commodity price forward curves, and commodity price dynamics, but it is clear that these models are missing important features. The non-negativity constraint on storage that plays a central role in this type of model can shed light in a rigorous, equilibrium-based way on the reasons that (a) futures curves sometimes are in backwardation, (b) commodity price volatilities and correlations are time varying, and (c) volatilities and correlations covary with inventories and the slope of the forward curve. However, this type of model fails miserably in explaining why old crop and new crop futures prices behave so similarly. Moreover, although it can closely mimic the behavior of spot prices (namely the evolution of spot price volatility over time), its ability to capture the dynamics of forward prices degrades rapidly with time to maturity. This last feature may reflect the fact that the model takes certain factors (namely productive capacity) as fixed, whereas in reality agents can invest in new capacity. The curse of dimensionality again sharply constrains our current ability to investigate this possibility.

The basic storage model is clearly not ready for derivatives pricing prime time, due both to its empirical deficiencies, and the curse of dimensionality. However, it does shed serious doubts on the reasonableness of received reduced form models used to price commodity derivatives. These models typically assume constant volatilities, and for curve sensitive products (such as spread options or swaptions), constant correlations. The storage model, which at least captures some important aspects of commodity price determination, shows clearly that these assumptions are dubious (as does much empirical evidence).

1.4 NON-STORABLE COMMODITIES

Life is far easier when studying non-storable commodities, such as electricity, because the lack of storability makes it unnecessary to solve recursively a dynamic programming problem to determine the efficient (and competitive equilibrium) allocation of resources, and hence to determine the equilibrium evolution of prices. For a true non-storable, every instant of time is distinct from every other instant, and intertemporal connections only arise due to persistence in demand or supply shocks. If these shocks are Markovian, an assumption that does not do too much violence to reality, it is a straightforward exercise to determine the non-storables' spot price as a function of current supply and demand fundamentals, and given specification of the dynamics of these fundamentals, to characterize the dynamics of the spot price.

This approach has been applied most frequently and successfully to the study of electricity markets by Eydeland and Geman (1998), Pirrong and Jermakyan (1999, 2008), and Eydeland and Wolyniec (2002). The basic approach in this research is to posit that the spot price of electricity depends on a small number of drivers, notably load (the demand for electricity), available capacity, and a fuel price (or a set of fuel prices). Each of these drivers evolves in a Markovian way. Moreover, especially when one considers load, an abundance of data makes it straightforward to determine empirically these dynamics. More specifically, the relevant supply curve of electricity (the relation between the spot price and load, conditional on the fuel price) is flat for low levels of load, but increases steeply as load nears available capacity. This supply curve implies that the dynamics of prices are time varying. When load is low, prices are low and exhibit relatively little variability, but when load is near capacity, prices can "spike" and exhibit extreme variability.

Since (a) a relatively small set of well-behaved, observable factors explains a substantial fraction of the variability in power prices, and (b) no solution of a dynamic programming problem is necessary to determine the relations between the fundamental state variables and spot prices, for non-storables it is possible to use a fundamentals-based model to price derivatives. Both Eydeland-Wolyniec and Pirrong-Jermakyan do just that, although in slightly different ways. The key nettle that must be grasped in doing so is that the underlying state variables are not traded assets, and hence the market is incomplete.[2] Thus, any derivatives price depends on a market price of risk function that must be inferred from the prices of traded claims. This fact is emphasized explicitly in Pirrong-Jermakyan, who use inverse techniques to estimate the market price of risk, but it is implicit in the Eydeland-Wolyniec approach as well, and as a result their calibration techniques effectively determine model parameters in the equivalent pricing measure.

The assumption of non-storability for electricity is quite apt for some markets, such as Texas, that are almost strictly fossil-fueled. The assumption is less realistic for other markets, such as Scandinavia, where hydro power is central; although electricity cannot be stored, water can be, and hence optimization in a hydro market requires solution of a dynamic

---

[2] Since electricity is not storable, and hence cannot be a proper asset, this problem cannot be avoided by going to a reduced form model that specifies the dynamics of the spot price.

programming problem. Non-storability is clearly apposite for other "commodities", most notably weather, that are definitely not storable.

The fundamentals-based models of non-storable prices are far preferable to reduced form models, particularly when pricing derivatives. Non-storables markets are inherently incomplete, so both structural and reduced form models must confront the problem of determining a market price of risk. Moreover, non-storability results in extreme non-linearities – such as large spikes in power prices – that are very difficult to capture in reduced form models, but which are a natural feature of well-specified fundamental models because these non-linearities are a direct consequence of fundamental supply and demand factors. Reduced form models also face difficulties in addressing the seasonality in many non-storables, whereas this is not a problem for the structural models. Finally, fundamental models can more readily handle the pricing of contingent claims with payoffs that depend on multiple factors (e.g., spark spread options, whose payoffs depend on load and fuel prices) because these factors are built into the pricing model, whereas multiple reduced form models must be bolted together to price these claims.

1.5 SUMMARY

Structural models of commodity price behavior have improved our understanding of commodity price dynamics, but for storable commodities there is still a yawning gap between theory and evidence. The modern Theory of Storage has shown how inventory decisions in competitive markets subject to random demand shocks can influence the shape of commodity forward curves. However, this theory cannot mimic the richness of commodity price behavior, especially for seasonal commodities and for long-dated forward contracts. These empirical deficiencies and the curse of dimensionality hamper the utility of these models as derivative pricing tools. In contrast, structural models of non-storables' prices can capture salient features of non-storables' prices, and with sufficiently sophisticated techniques for extracting information about market risk prices from the prices of traded claims, can be used to price and hedge non-storable commodity contingent claims.

1.6 REFERENCES

Chambers, M. and R. Bailey. (1996). A Theory of Commodity Price Fluctuations. *Journal of Political Economy* **104**, 924–957.

Deaton, A. and G. Laroque. (1992). On the Behavior of Commodity Prices. *Review of Economic Studies* **59**, 1–23.

Deaton, A. and G. Laroque. (1996). Competitive Storage and Commodity Price Dynamics. *Journal of Political Economy* **104**, 896–923.

Eydeland, A. and H. Geman. (1998). Fundamentals of Electricity Derivatives, in R. Jameson, ed., *Energy Modelling and the Management of Uncertainty*. London: Risk Publication.

Eydeland, A. and K. Wolyniec. (2002). *Energy and Power Risk Management: New Developments in Modeling, Pricing, and Hedging*. New York: John Wiley & Sons Inc.

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

Ng, V. and C. Pirrong. (1994). Fundamentals and Volatility: Storage, Spreads, and the Dynamics of Metals Prices. *Journal of Business* **67**, 203–230.

Osborne, T. (2004). Market News in Commodity Price Theory. *Review of Economic Studies* **7**, 133–164.

Pirrong, C. (1999). Searching for the Missing Link. Working paper, Washington University.

Pirrong, C. (2006). Price Dynamics and Derivatives Prices for Continuously Produced, Storable Commodities. Working paper, University of Houston.

Pirrong, C. (2007). Testing the Theory of Storage. Working paper, University of Houston.

Pirrong, C. (2008). Stochastic Volatility of Fundamentals, Commodity Prices, and Inventories. Working paper, University of Houston.

Pirrong, C., and M. Jermakyan. (2008). The Price of Power. Forthcoming, *Journal of Banking and Finance*.

Pirrong, C., and M. Jermakyan. (1999). Pricing Forwards and Options Using the Mesh-Based Partial Differential Equation Approach, in R. Jameson, ed., *Energy Modelling and the Management of Uncertainty*. London: Risk Publications.

Routledge, B., D. Seppi, and C. Spatt. (2000). Equilibrium Forward Curves for Commodities. *Journal of Finance* **55**, 1297–1338.

Scheinkman, J., and J. Schectman. (1983). A Simple Competitive Model of Production with Storage. *Review of Economic Studies* **50**, 427–441.

Williams, J., and B. Wright. (1991). *Storage and Commodity Markets*. Cambridge: Cambridge University Press.

Working, H. (1949). The Theory of the Price of Storage. *American Economic Review* **39**, 1254–1262.