ONE

Introduction

1.1 Introduction

Not to put too fine a point on it, but the study of commodity prices has long been something of an academic stepchild. Most work on the subject is in the domain of specific fields, notably agricultural economics. Only a smattering of articles on the subject has appeared in broader publications, such as the *Journal of Political Economy* or the *Journal of Finance*.

Especially in finance, this relative obscurity arguably reflects the niche role of commodities in the broader financial markets, as compared to equity and fixed-income markets. But commodities are in the process of becoming mainstream. In the 1990s, and especially the 2000s, many major banks and investment banks have entered into commodities trading. Indeed, commodity trading – especially in energy – has become an important source of profits for major financial institutions such as Goldman Sachs, Morgan Stanley, and Citibank. Simultaneously, and relatedly, many investors have entered into the commodities market. In particular, pension funds and other portfolio managers have increasingly viewed commodities as a separate asset class that, when combined with traditional stock and bond portfolios, can improve risk-return performance. Furthermore, financial innovation has eased the access of previously atypical participants into the commodity markets. Notably, commodity index products (such as the GSCI, now the S&P Commodity Index) and exchange traded funds (ETFs) have reduced the transaction costs that portfolio managers and individual investors incur to participate in the commodities markets. Thus, a confluence of forces has dramatically increased the importance of, and interest in, commodities and commodity prices.

What is more, this increase in the presence of investors and large financial intermediaries in commodity markets combined with extraordinary price movements in commodities in the mid-2000s to make commodity prices an important political issue. Most notably, the unprecedented spike in the price of oil in 2008 created a political firestorm in the United States (and elsewhere) that has led to numerous calls to regulate the markets more restrictively. Indeed, the coincidence of the entry of new financial players into the markets and rocketing prices led many market participants, politicians, and pundits to attribute the latter phenomenon to the former, and hence to call for limitations on the ability of financial institutions, portfolio investors, and individual investors to buy and sell commodities.

Thus, the 2000s have seen commodities achieve an economic and political prominence that they had lacked since a much earlier era (the nineteenth and early twentieth centuries) in which a far larger portion of the population earned a living producing or processing commodities. Unfortunately, the modeling of commodity prices has not kept pace. Most practitioners have adopted reduced-form models, such as the model that is the basis for the Black-Scholes option pricing formula, to analyze commodity prices and to price commodity derivatives. Structural models of commodity prices that explicitly account for the implications of intertemporal optimization through storage have been around since Gustafson (1958) and have been developed by Scheinkman and Schectman (1983), Williams and Wright (1991), and others. These models, however, have been in a state of relative stasis. Moreover, the empirical analysis of these models has been extremely limited, and little use has been made of them to answer questions related to the effects of speculation.

I intend this book to push the structural modeling of commodity pricing forward, to provide a better understanding of the economics of commodity pricing for the benefit of both academics and practitioners. This book builds on the rational expectations, dynamic programming–based theory of storage epitomized by Williams and Wright (1991), but goes beyond the existing literature in many ways.

First, whereas the received models typically incorporate only a single source of economic uncertainty (e.g., a single net demand shock), I (a) demonstrate that such models are incapable of explaining salient features of commodity price dynamics, and (b) introduce models with multiple shocks that can capture many of these features.

Second, unlike received work, in this book I exploit important cross-sectional differences among commodities to derive empirical implications.[1]

[1] See especially Deaton and Laroque (1995, 1996), which lump together commodities as disparate as copper and corn in a single empirical framework.

As I discuss in more detail later in this introduction, commodities can differ on a wide variety of dimensions. Some, such as copper or oil, are produced continuously and have relatively non-seasonal demands. Others, such as corn or soybeans, are produced seasonally. These fundamental differences lead to distinctive price behaviors; the ability of suitably adapted models to explain these differences sheds a bright light on the strengths and limitations of the received structural approach.

Third, quite curiously, the empirical literature structural models of commodity prices tend to focus on low-frequency (e.g., annual) data.[2] They also tend to focus on the spot prices of commodities, as well as on price levels and the first moments (means) of prices; they typically ignore the behavior of higher moments, such as variances and measures of covariation between different prices. These models particularly tend to ignore time variation in these variances and covariances and the association between these time variations and fundamentals.

These tendencies have several pernicious effects. For one thing, they obscure the potentially illuminating differences between continuously and periodically produced commodities. For another, and perhaps more important, they result in a slighting of a tremendously rich source of data: high-frequency (e.g., daily) data on a wide variety of derivatives on a similarly wide spectrum of commodities. In particular, there are abundant data on commodity futures prices for a wide variety of commodities. Moreover, data on other commodity derivatives are becoming increasingly available. Most notably, commodity options are more widely traded than ever, and hence option price data are becoming commonplace.

To exploit these data, the book focuses on the implications of structural commodity price models for the behavior of commodity spot *and* forward prices at *high frequency*; the variances of these spot and forward prices and the correlations between them; the comovements of quantity variables (e.g., inventories) and prices; and the prices of other commodity derivatives, most notably options. Moreover, I continually confront these implications with the data, to see where the models work – and where they do not.

The basic approach is to see what data are available to test the models, derive the implications of the models for the behavior of these observables, and evaluate the performance of the models when faced with the data.

This presents the models with extreme challenges and, as discussed later, they quite often fail. But that is part of the plan. After breaking the models, we can learn by examining the pieces.

[2] See again Deaton and Laroque (1995, 1996).

The transparency of fundamentals in commodity markets (in contrast to equity or currency markets, for instance) holds out the promise of allowing structural models of commodity price behavior to be devised 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.

As just noted, the cross-sectional diversity of commodities represents a potentially valuable source of variation that can be exploited to gain better understanding of the determinants of commodity prices and their behavior and to inform structural models of commodity markets. To understand this fundamental point more clearly, it is worthwhile to examine this diversity in more detail and, at the same time, introduce some modeling issues that this diversity raises and discuss the received modeling literature.

1.2 A Commodity Taxonomy

Although the catchall term "commodity" is widely applied to any relatively homogeneous good that is not a true asset, it conceals tremendous diversity, diversity that has material impacts on price behavior and modeling.[3]

[3] The distinction between a commodity and an asset proper is that an asset (such as a stock or a bond) generates a stream of consumption (à la the "trees" in a Lucas [1978] model) or a stream of cash flows that can be used to buy consumption goods (e.g., a bond), whereas a commodity is itself consumed. An asset represents a "stock" that generates a flow of benefits. There are some potential ambiguities in this distinction, especially inasmuch as this book discusses repeatedly the role of "stocks" of commodities. However, as discussed later, even though there are commodity stocks, commodity forward prices behave differently from asset forward prices. Whereas asset forward prices always reflect full carrying costs (the opportunity cost of capital net of the asset's cash flow), commodity forward prices do not. This distinction between a consumption good and a true asset has important implications for the possibility of bubbles in commodity prices, that is, self-sustaining price increases not justified by fundamentals. Williams et al. (2000) and Gjerstad (2007) show that experimental consumption good markets almost never exhibit

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. Shipping services are another non-storable commodity. Although ships are obviously durable, the services of a ship are not: space on a ship that is not used today cannot be stored for use at a later date. Shipping derivatives are also increasingly common; derivatives on bulk commodities began trading in the early 2000s, and the first container derivative trade took place in early 2010.

Most other commodities are storable (at some cost), but there is considerable heterogeneity among goods in this category. This heterogeneity occurs on the dimensions of temporal production patterns, temporal demand patterns, and the nature of the capital used to produce them.

Some commodities 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 seasonality in demand. Heating oil, natural gas, and gasoline are prime examples of this type of commodity.

Other commodities are produced periodically (e.g., seasonally) rather than continuously, 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.[4]

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.

In sum, there is considerable diversity among commodities. This presents challenges and opportunities for the economic modeler. As to challenges,

bubbles; in contrast, Smith, Suchanek, and Williams (1998) show that experimental *asset* markets are chronically prone to bubbles.

[4] Adverse weather events sometimes can have effects that span crop years. For instance, the intense drought of 2010 in central Russia devastated the 2010 crop, but also left the ground very dry. This delayed planting of the 2011 crop, which raises the risk of a smaller-than-normal 2011 harvest.

fundamentals-based theories must take these variations across commodities into account, so one-size-fits-all models are inappropriate. As to opportunities, this cross-sectional variation has empirical implications that can be exploited to test fundamental-based structural models.

1.3 Commodity Markets and Data

Fortunately, just as there are many different commodities, there are many actively traded commodity markets. These markets produce prices that are of interest and important in their own right, but which also can be used to test structural commodity models.

In particular, although with a few limited exceptions there are few liquid and transparent "spot" markets for commodities,[5] there are active, liquid, and transparent futures markets for many commodities. (But not all; some important commodities such as iron ore have no active futures market.)

A futures contract is a financial instrument obligating the buyer (seller) to purchase (sell) a specified quantity of a particular commodity of a particular quality (or qualities) at a particular location (or locations) at a date specified in the contract. For instance, the July 2010 corn futures contract traded at the Chicago Mercantile Exchange requires the buyer (seller) to take (make) delivery of 5,000 bushels of #2 corn at a location along the Illinois River chosen by the seller during the month of July 2010. The buyer and the seller agree on the price terms, but all of the other contract terms are established by formal organizations – futures exchanges.[6]

These futures exchanges operate centralized auction markets where buyers and sellers can negotiate transactions. The exchanges typically host continuous, double-sided auctions. Historically, these auctions were face-to-face affairs that took place on an exchange floor – "the pit." In recent years, most trading has migrated to electronic, computerized exchange systems (although commodities lagged behind financial futures in this regard). The prices negotiated during these auctions are broadcast around the world and represent the primary barometer for commodities prices. Buyers and

[5] A spot market is a market for immediate delivery. Practically speaking, even a spot transaction involves a separation in time between the consummation of a transaction and the delivery of the commodity, so spot trades are properly very short-term forward transactions.

[6] A futures contract is a particular kind of forward contract. A forward contract is any contract that specifies performance at a future date. A futures contract is a forward contract traded on an exchange, where the performance on the contract is guaranteed by an exchange clearinghouse. The terms "future" and "forward" and "futures price" and "forward price" are often used interchangeably.

sellers of physical commodities base the prices of their transactions on these futures exchange prices.

At present, futures markets exist for many physical commodities, including energy products (especially oil, heating oil, gasoil, natural gas, and gasoline); grains and oilseeds (including wheat, corn, and soybeans); industrial metals (such as copper, aluminum, lead, and nickel); precious metals (gold, silver, platinum, and palladium); fibers (notably cotton); meats (live hogs, live cattle, and pork bellies); and non-grain and non-meat food products (such as coffee, cocoa, and sugar). In some cases, there are futures contracts traded on different varieties of the same commodity; for instance, Brent crude oil futures and West Texas Intermediate (WTI) crude oil futures are both traded.

For most of these commodities, futures contracts with different delivery dates are traded. In the WTI crude market, for instance, contracts calling for delivery in every month for the next several years are traded simultaneously. In the corn market in Chicago, contracts are actively traded for delivery in March, May, July, September, and December.

Thus, for most active commodity futures, at any time there are multiple prices, each corresponding to a different delivery date. The locus of forward prices on the same commodity for different delivery dates is called the "forward curve."

Forward curves exhibit a variety of shapes. Some slope up (i.e., prices rise with time to expiration). Such markets are said to be in "contango" or to exhibit "carry." Others slope down; these are said to be "inverted" or "in backwardation." Others are humped, rising over some range of delivery dates, falling over others. Some exhibit seasonality, with highs and lows corresponding to different seasons.

Any good theory should be able to generate forward curves that exhibit the observed diversity in these curves in actual markets. Moreover, any good theory should demonstrate how these curves evolve with changes in fundamental conditions.

In sum, there are vast amounts of futures price data for commodities that exhibit the diverse characteristics discussed in the taxonomy. Because the markets operate continuously, these data are of very high frequency. There are at least daily data on futures contracts for a variety of delivery dates for many commodities; for some commodities, intra-day data are also available. This represents a vast repository of information that can be used to test, challenge, and potentially break structural commodity models.

But although futures markets provide the most extensive source of price data, there are other commodity markets and hence other data sources.

Many commodity forwards are traded on a bilateral basis in over-the-counter (OTC) markets; that is, they are traded off exchange. Moreover, options on commodities are traded both on exchange and in OTC markets. Some of these options are vanilla puts and calls. Others are more exotic instruments, such as spread options that have a payoff dependent on the difference between futures prices on a particular commodity, with different delivery dates, or swaptions that are effectively options on the forward curve for a particular commodity.

It is well known that the prices of options depend on the characteristics of the dynamics of the underlying price, most notably on the volatility (and perhaps higher moments) of the underlying price. Thus, a copper option's price depends on the volatility (and perhaps higher moments) of the forward price of copper. This means that options prices incorporate information about commodity price dynamics, and hence these options prices can be used to test predictions of commodity price models.

As discussed later, structural models for storable commodities also make predictions about the behavior of commodity inventories, that is, commodity stocks. Indeed, the behavior of stocks is of particular interest because it speaks to long-standing and ongoing battles over whether speculation on futures markets distorts prices. An important economic role of prices is to guide the allocation of resources, so distortions in prices will manifest themselves in distortions in the allocation of real things – such as commodity stocks.

Unfortunately, with a few exceptions, data on commodity stocks are far less abundant than data on prices. Nonetheless, there are some markets – notably the London Metal Exchange's markets for industrial metals – that produce high-quality, high-frequency data on stocks. I will use these data where available.

There is, in brief, a plethora of data that can be used to test commodity price models. This book develops models tailored to the specific features of particular commodities, generates predictions from these models, and uses some of this bounty of data to test these predictions.

1.4 An Overview of the Remainder of the Book

The objective of this book is to develop models customized to capture the salient features of particular commodities (e.g., storability, production frequency), derive the implications of these models for the behavior of commodity prices and stocks, and then examine how well these predictions stack up against the data.

The term "behavior" is meant to be very encompassing. In this book I derive and test the implications of structural models not just for the level of forward curves, but also for the variances of forward prices of different maturities, the correlations between forward prices of different maturities, and the pricing of options. Moreover, I focus on *high-frequency* commodity price behavior. That is, I derive and test implications about the day-to-day behavior of commodity prices. This focus allows me to mine the rich seams of futures data and should also make the work of particular interest to commodity market participants who trade every day.

The rest of the book is organized as follows.

1.4.1 Modeling Storable Commodity Prices

The received theory of storage, based on the stochastic dynamic programming rational expectations modeling framework, is an economically grounded approach for understanding how commodity prices behave and how fundamentals affect commodity prices. This framework can be adapted to each particular commodity. In this chapter, I review the basics of the model and the computational approach to solving it. This chapter also introduces the partial differential equation (PDE) approach to solving for forward prices that is an important part of the solution algorithm. The chapter reviews PDE methods and shows how these methods can be used to determine the prices of more complicated contingent claims in the storage economy. These claims include vanilla options, swaptions, and spread options. Because multiple sources of risk are necessary to provide a reasonable characterization of commodity price behavior, the chapter focuses on modern PDE approaches (such as splitting techniques) for high-dimension problems.

1.4.2 A Two-Factor Model of a Continuously Produced Commodity

The conventional commodity price modeling approach assumes a single source of uncertainty. This approach cannot explain the imperfect and time-varying correlation between forward prices with different maturities. This chapter explores the implications of multiple risk sources with different time series properties; that is, two demand shocks of differing persistence. The chapter focuses on the implications of this framework for the behavior of forward curves, the variances of prices, the correlations between different forward prices, and the dependence of these moments on supply and demand conditions. It also explores the pricing of options in the storage economy with multiple demand shocks, including an analysis of commodity volatility surfaces and how these volatility surfaces depend on underlying supply and demand conditions.

1.4.3 The Empirical Performance of the Two-Factor Model

Formal econometric testing of rational expectations models is challenging, so techniques that mix calibration and estimation are common in macroeconomics where rational expectations models are the most widely used theoretical tool. This chapter adapts those techniques to the study of an important commodity market – the copper market. I use the Extended Kalman Filter (EKF) and a search over the relevant parameter space to determine the persistences and volatilities of the demand shocks of the model of Chapter 3 that best capture the dynamics of copper prices. I show that the extended, multi-shock model captures salient features of copper price behavior, but that the model cannot capture well the behavior of long-term (e.g., 27-month) forward prices; moreover, even though the model accurately characterizes the dynamics of the variance of the spot price (as documented by Ng-Pirrong), it does a poor job of capturing the behavior of longer-dated variances (e.g., 3-month futures price variances) and the volatility implied from copper options prices. This motivates an analysis of what changes to the model are necessary to capture these dynamics.

1.4.4 Commodity Pricing with Stochastic Fundamental Volatility

Traditional storage models (and the model of Chapter 3) assume homoskedastic net demand disturbances; that is, the variance of the fundamental shocks remains constant over time. This chapter explores the implications of stochastic demand variability. Such stochastic variability is plausible. The chapter shows that a model that incorporates stochastic demand volatility can explain otherwise anomalous behavior in commodity markets, namely, episodes where both inventories and prices rise and fall together (whereas the traditional storage model implies that they should typically move in the opposite direction). This is an important finding, because some critics of commodity speculation have asserted that such comovements are symptomatic of speculative price distortions. Moreover, adding stochastic fundamental volatility results in a more accurate characterization of the behavior of forward price variances and option implied volatilities.