THE UNIVERSITY OF CALGARY

The Behaviour of Natural Gas Prices

by

Todd Allan Kemp

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES

IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE

DEGREE OF MASTER OF ARTS

DEPARTMENT OF ECONOMICS

CALGARY, ALBERTA

NOVEMBER, 1995

THE UNIVERSITY OF CALGARY

FACULTY OF GRADUATE STUDIES

The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies for acceptance, a thesis entitled "The Behaviour of Natural Gas Prices" submitted by Todd Allan Kemp in partial fulfilment of the requirements for the degree of Master of Arts.

Supervisor, Dr. A. Serletis, Department of Economics

Dr. A.J. MacFadyen, Department of Economics

Dr. G. Sick, Faculty of Management

Dec. 5, 1995
Date

Abstract

This thesis examines several different properties of natural gas prices and markets. Specifically, I address the form and functions of futures markets, the cyclical behaviour of natural gas prices (utilizing Prescott's (1986) methodology), tests on the theory of storage (utilizing Fama and French's (1988) methodology), and tests for market efficiency (utilizing Fama's (1984) methodology).

My analysis concludes that natural gas prices are weakly procyclical with output and lag the cycle by two periods. Moreover, natural gas prices, in general, tend to move in the same fashion as other energy commodity prices. Natural gas passes one of the three tests of the theory of storage and this result is inconsistent with Serletis and Hulleman's work (1994) on energy prices and storage. Further, it appears that natural gas markets have operated in an efficient manner as the current futures price appears to have the power to predict the future spot price. However, I have uncovered evidence of a time varying risk premium which is consistent with Serletis (1991) results with respect to heating oil, Fama and French's (1988) result with respect to metal prices, and finally, Cho and McDougall's (1990) results with respect to other energy prices.

Acknowledgements

I would first like to thank my supervisor, Dr. Apostolos Serletis, who set clearly defined goals and provided the necessary guidance for me to complete this thesis in such a timely manner.

Thanks must also go out to David Krause who provided his computer expertise and help with 'Shazam'.

Finally, thanks are in order to Jason Sandmaier for his proof reading of the text.

Dedication

To Elmer

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

The goal of this thesis is to examine several different properties of natural gas prices and markets. To begin this discussion, it is of some use to get a better understanding of natural gas itself.

Natural gas is used primarily as a heating and cooling tool, and is utilized in both the commercial and residential sectors of the economy. Natural gas is playing an increasingly large role in today's economy, due to several advantages it possesses over other fossil fuels; some of which are listed below.

Comprised of mainly methane, with small amounts of propane, carbon dioxide and other gases, natural gas is a homogeneous compound that does not vary very much in its composition no matter where it is found. In fact, once it enters a pipeline, it must meet specific standards, which ensures its homogeneity. This gives it a distinct advantage over fuel substitutes such as crude oil (of which many varieties are found) since its delivery through pipelines is "instantaneous." One simply has to inject gas at one end of the pipeline and simultaneously withdraw gas at the other end. Comparatively, crude oil is delivered in "batches" through a pipeline. One type of crude oil follows another down the pipeline, and one can have numerous different types of crude flowing at once. Thus, it is not possible for "instantaneous" delivery to take place. This homogeneous characteristic that natural gas possesses not only eases delivery problems associated with other fuels, but also facilitates more flexible trading volumes, and thus, easier trading.

Another advantage natural gas possesses over other fuels is that it is a comparatively cleaner burning fuel, that emits few pollutants into the environment. In an increasingly energy conscious society, this is a strong attribute to possess. Dubbed the "fuel of the future" by some, the use of natural gas is becoming more widespread and innovative. For example, automobiles utilizing natural gas as their source of fuel as opposed to unleaded gasoline have emerged on the marketplace.

With these point in mind, I believe that it is of interest to know how this growing market behaves, what some of its economic characteristics are, and thus, this is the thrust of this thesis. To reach this end, I examine four different areas relating to natural gas markets. Specifically, I look at:

(i) The form and function of natural gas futures markets.
(ii) The cyclical properties of natural gas prices.
(iii) The theory of storage.
(iv) Tests for market efficiency.

Chapter two of this thesis begins my look at natural gas markets by describing different market structures in general, and then moves on to discuss futures markets in general. Futures are an integral part of many markets whose commodity exhibits somewhat volatile price movements, and natural gas is no exception. While futures markets serve many purposes, one of their greatest attributes is that they reduce the price risk associated with this volatility. Arguably, their second most valuable attribute is that they provide an avenue for price discovery. Thus, the goal of this chapter is tô familiarize the reader with the form and function of futures markets. This is important as all of the following chapters deal with, in some form, futures prices and futures markets.

Chapter three investigates the cyclical properties of natural gas prices. Utilizing the Hodrick-Prescott filter (using Prescott's (1986) methodology) I remove the trend component of a stream of monthly natural gas futures prices to obtain its cyclical component. Using the same approach, I then do the same for a United States production index, price index, and unemployment rate. After doing so, several different avenues are taken to get an idea of how natural gas prices behave.

First, cross correlations of natural gas prices are calculated with each of production, prices and the unemployment rate. This tells us if natural gas prices are procyclical, countercyclical, or acyclical with the variable in question. Then, by shifting the natural gas price series forward and backwards several periods relative to the variable in question, we once again calculate a correlation coefficient. Engaging in this exercise, we are able to determine if, for example, production leads natural gas prices, the opposite is true, or neither leads the other.

Finally, using the same methodology as above, I compare the cyclical component of crude oil prices with those of natural gas, heating oil, and unleaded gasoline. Doing this gives us some idea about how energy prices move with respect to each other. After completion of this exercise, I will have formed quite a complete picture of how natural gas prices tend to move. This has important implications for those who make their living in the market, namely, hedgers and speculators.

Chapter four investigates the theory of storage originally depicted by Working (1949) and expanded on by Brennan (1958) and Tesler (1958). The original motivation behind this theory was to try to explain why it was observed that commodities were being stored when the price of (or return on) storage was negative. The authors conclude that this behavior is exhibited since there is some sort of convenience yield present on stored commodities. These authors develop three tests on the theory of storage, all of which I reproduce with respect to natural gas.

The first test of the theory of storage examines the prediction that supply and demand shocks cause more volatility in spot and futures prices when inventory levels are low. The second test examines the prediction that supply and demand shocks cause approximately equal changes in spot and futures prices when inventory levels are high, but shocks cause spot prices to change more than futures prices when inventory levels are low. Lastly, the third test of the theory of storage examines the prediction that supply and demand shocks produce larger changes in near term futures as opposed to longer term futures.

To conclude my analysis on natural gas markets, chapter 5 looks at tests of market efficiency and the possibility that futures markets can be used as a forecasting tool. The background to this chapter is based on the foundation of futures markets themselves. While it is almost universally accepted that futures markets reduce price exposure, the notion that current futures prices can predict future spot prices is questionable. If it is the case that current futures prices can reliably predict future spot prices, we say that this is an efficient market. To test for market efficiency, I utilize the approach of Fama (1984).

First, I will test to see if the natural gas price series is stationary using the methodology of Dickey and Fuller (1979). Then, following Fama (1984) I test various hypotheses related to the variability of risk premiums and expected spot price changes. Doing so enables us to see, if indeed, current futures prices can predict future spot prices.

All of my results and conclusions will be summarized in the concluding chapter, chapter 6.

Chapter 2

Futures Markets Operations

2.1 Introduction.

This chapter will describe the form and function of futures markets. Specifically, I will look at how futures markets originated, how they are sustained, and how a futures contract is different from other types of commodity contracts. To begin, I will look at commodity markets, then move on to the evolution of futures markets. All discussion of futures markets will be in a general tone and I will discuss the specific futures and futures market I deal with in more detail near the end of this chapter.

2.2 Commodity Markets.

A good place to begin is a short discussion of the simplest market form, a commodity market. Commodity markets facilitate trading by providing a central location for buyers and sellers to trade. These markets can take on various forms and sizes, from a small town's "farmers market" to a large market such as a stock market. Transactions on commodity markets usually occur daily, with payment and delivery made almost immediately. For example, when you buy a loaf of bread at the local grocery store, you pay the clerk for it at the counter and he immediately gives you the bread to take home.

There is some risk that one of the parties may not fulfill his end of the deal. For example, you may have given the clerk a bad cheque. While these markets satisfy the basic trading needs of providing a trading partner of both buyers and sellers they are not "complete" enough to satisfy all trading needs.

A common feature of many traded commodities is price volatility. Associated with price volatility is risk a buyer (seller) must face when contemplating future purchases (sales). A person planning to buy (sell) a commodity in one month's time may be quite disturbed to find that in one months time the commodity that he wished to purchase (sell) rose (fell) in price quite dramatically compared to the price on the market today. Conversely, he may be pleased to see that the price of the commodity has dropped (rose) relative to the price on the market today. However, people who are risk averse would like to avoid such "surprises" and be certain of which price they will pay (receive) when they buy (sell) the commodity in one months time. Thus, the desire for forward contracting arose.

2.3 Forward and Futures Markets

Forward contracting eliminates the risk associated with price volatility. A forward contract is a private transaction made now to purchase a specific amount of a cash asset at a specific price in the future. Payment and exchange of the asset occur at a date agreed upon by both parties. This way, both parties know exactly the terms of the agreement and all risk is eliminated.

Futures contracts are essentially a forward contract with a few variations. The differences between the two are as follows:

(i) Futures are traded on organized exchanges.

Natural Gas is traded on the New York Merchantile Exchange (NYMEX) which ensures that buyers and sellers have a centralized place to conduct business. In forward contracting, one must privately seek out a trading partner, and thus, futures contracting eliminates this transaction cost. I will talk about the specifics of the NYMEX in the next section of this chapter.

(ii) Futures contracts are standardized for all transactions and cannot vary.

Specifically, they outline terms for delivery volume, delivery date and quality. A stringent structure such as this helps facilitate transfers between traders as all market participants know the terms of trade. Conversely, a forward contract is highly personalized and will vary depending upon the parties involved.

(iii) Futures contracts are more liquid than forward contracts, facilitating larger trading volumes and making them easy to close.

Since forward contracts tend to be individual specific, they are difficult to trade. The trading option on a futures contract makes them more liquid than forward contracts and promotes an active trading environment and a high volume of trade since there is no need to actively search out trading partners.

While the standardization of the contract makes trading easy, it is often too rigid for the day to day operations of the commercial participants. In this light, it comes as no surprise that a typical futures contract is used solely as a financial instrument and is generally used for hedging or speculating. For example, a hedger will hold the futures contract to shield himself from spot price volatility until he wishes to make a cash transaction for a commodity. A speculator will hold a futures contract hoping that spot price movements will allow him to make a windfall gain.

The rigidity of the contract results in only about two percent of all futures resulting in actual delivery. While the rules of the exchange state that all those holding a position in a futures contract on the last day of trading of that specific contract must either make or take delivery of the specified commodity, exchanges have allowed several variations of this rule, allowing a futures position to be easily closed. This is in stark contrast to forward contracting which almost always calls for delivery for settlement.

Besides making or taking delivery of a commodity, futures market participants may choose to engage in an Exchange For Physicals (EFP) or an Alternative Delivery Procedure (ADP).

(iv) Futures contracts are safeguarded against defaults.

A common feature of futures markets are daily settlements and margin requirements. Every futures trading exchange collects margins from all of the interested brokers. Brokers are the intermediaries between the exchange and the actual customer, so the broker puts the margin in on the customers behalf. The actual amount of the margin will fluctuate between five and ten percent of the contract's total face value and is designed to be equal to the average daily fluctuation in the value of the contract that is being traded.

The margin is adjusted daily after the exchange views the closing or settlement price. Accordingly, the exchange will then request additional margin, or conversely, pay margin to the broker in question. For example, if the closing price goes up from one day to the next, the exchange will then credit the buyer's margin and debit the seller's margin.

Interestingly, if for some reason, a broker cannot fulfill the contract as specified, the clearinghouse within the exchange holds member brokers responsible for the performance of the contract. This protects the customer and ensures the credibility of the exchange.

We can now move on to discuss the specifics of the exchange on which natural gas is traded.

2.4 The New York Mercantile Exchange.

Natural gas futures are traded on the New York Mercantile Exchange (NYMEX), and what follows is a brief history of its operation. NYMEX was founded over 100 years ago. The exchange started as a produce exchange and remained relatively small until it diversified into energy futures. Beginning in 1978, they created the New York heating oil contract. It added a leaded gasoline contract in 1981 (later changed to an unleaded contract), and a West Texas Intermediate crude oil contact in 1983. Following this, several other energy contracts were added, including natural gas in 1990. Today, NYMEX is the leading energy futures exchange in the world and the third largest futures exchange in the world, with energy contracts accounting for 90 percent of its turnover.

2.4.1 The NYMEX Natural Gas Futures Contract.

As mentioned earlier, a futures contract is highly standardized. The NYMEX natural gas futures contract is no exception. The specifications are as follows:

The volume specified is 10,000 million British thermal units (MMBtu) to be delivered approximately evenly over a period of one month. This amounts to about 320 mcf/day.

The physical deliveries will take place at Sabine Pipeline's Henry Hub, which is an active spot gas trading point. It is located in Southern Louisiana on the Gulf Coast. This location is ideal as it connects most major gas pipelines across the United States. With the system of pipelines being highly integrated, access to virtually all North American markets is assured. The gas must also meet standard pipeline specifications.

2.5 Criteria for Successful Futures Trading.

Just because a commodity is traded on a spot market, it does not mean that the commodity will be successful on a futures market or even that futures trading is desirable. Several criteria must be met before a futures market will be a success for any given commodity. What follows is brief discussion of the criteria necessary for a successful futures market. Some of these criterion are more important than others and not all of them may be present in any given futures market. However, all can be deemed as desirable and the more of them a futures market possesses, the more likely it is that it will be successful.

As mentioned earlier, forward and futures markets arose as a result of price volatility. Thus, price volatility is essential for futures trading. If price fluctuations were not broad enough to cause some degree of risk to the parties involved, there would be no need for futures trading. Without price uncertainty, hedgers would not need to offset their physical holdings with a futures contract and speculators would not find a wide enough profit margin in the market and would concentrate on a different market possessing greater price fluctuations. It is generally accepted that price fluctuations from the mean price of plus or minus twenty percent per year are necessary to maintain futures trading. The higher the degree of price variability, the more likely it is that a futures market will be successful as more participants will find it desirable to hedge. Natural gas markets satisfy this criterion.

Hand in hand with price volatility is the presence of uncertian supply and demand conditions. In the energy sector, fluctuating supply and demand conditions can cause severe price volatility. For example, unseasonably cold weather will increase natural gas demand immensely. If this is unexpected, we may find that natural gas supplies are less than are needed to satisfy the demand and prices may rise dramatically as local distribution companies scramble to find supplies. Obviously, price volatility is linked to uncertain supply demand conditions and this is a necessary condition for futures market survival. Natural gas demand is highly seasonal and weather dependent, ensuring price fluctuations.

A sufficient amount of deliverable supplies must be available to ensure the proper operation of a futures market. Also, these supplies must meet the quality standards set out in the futures contract. This may seem like a paradox when we compare this criterion to the last one, as we must have both uncertain supply demand conditions but at the same time we also need to have an adequate supply in order to ensure delivery. However, there is a middle ground between the two. Although one rarely witnesses delivery in a futures market, potential delivery is essential to their operation. The simple possibility that delivery may occur will help to keep futures and spot prices in line. In fact, if the two did not converge on their own, arbitrage opportunities would force them into line. Delivery and deliverable supplies are essential for futures market operation. A storage capacity of about 30 days average demand is usually needed to meet this criterion. Natural gas is commonly stored for periods up to one year, so this criterion is satisfied.

A high degree of competition in the commodity market is also necessary for a successfully operating futures market. No one market participant should be able to exert their power and affect market prices. If one large firm could completely control supply (and thus, price), the need for a futures market would be negligible as prices would be predictable and there would be no need for either hedging or speculating. With respect to natural gas markets, the large number of producers ensures that no one producer has significant market power. Although market concentration is difficult to quantify, it is generally desirable to have the top five firm's market share at less then 50 percent and the top ten firm's market share at less than 80 percent. With a large number of producers in the natural gas industry, natural gas markets pass this criterion and can be deemed competitive.

Product homogeneity will help facilitate futures trading. Due to the standardized nature of any futures contract, it is desirable for any futures commodity to be as homogeneous as possible. Natural gas fits this criterion to a key. Composed mainly of methane, with small amounts of propane, carbon dioxide, and other gases, it is essentially a generic commodity by the time it enters a pipeline.

For the futures market to operate successfully, the futures commodity should have the characteristic of being able to be stored for long periods of time (from six to twelve months is preferred). Natural gas is certainly storable for long periods of time and satisfies this criterion. In fact, due to the nature of the natural gas industry itself, storage is essential. Often operating on long term contracts with pipelines which specify a certain amount of gas to be delivered daily, gas companies often find it useful to store their gas. This is due to the fact that natural gas demand is highly seasonal and all of the gas delivered cannot be used at all times. For example, residential gas demand tends to be low in the summer so excess gas not sold will be put in storage until the winter. Basically, contracts are designed to ship a constant amount for the duration of the contract and this usually results in over supply in the warm months and under supply un the cold months. However, by using storage facilities, excess gas from the summer months can supplement supply in the winter months.

From the above criteria, we can say with some assurance that natural gas qualifies as a product a futures market will support.

2.6 Opportunities in Futures Markets.

In this section, I will describe what functions futures markets provide and some of the activity which takes place in them.

Futures markets have three basic functions: to protect market participants from risk, to provide an avenue for price discovery, and finally, to provide an investment opportunity. Since one of the goals of this thesis is to determine if futures prices actually do predict spot prices, I will deal with that particular point in chapter 5.

The primary function of these markets is to provide an avenue for market participants to reduce risk. This is done through hedging. There are two basic types of hedges, the short hedge and the long hedge. When we say someone is "short" in a commodity, the trader is not in possession of the good. When we say someone is "long" in a commodity, the trader is in possession if the good.

If a trader is engaging in a short hedge, this means that the party possesses an inventory of a commodity and plans to sell it sometime in the future. This is usually engaged in by a trader who wishes to protect himself from the risk of a price fall between the current time period and the period in which he wishes to sell his commodity. While it is best to hedge in a falling market, it does not really matter which way a price may move, the trader will end up equally well off no matter which way prices may move. The short hedge works like this: Suppose Company A owns a volume of natural gas reserves, but for some reason, they think that natural gas prices may drop in the future. To protect the value of the projected sales on the reserves they hold, Company A will engage in a short hedge. As an example, suppose the futures price today, say July first, for delivery in December is \$1.50/MMBtu but Company A believes that the spot price will be less than \$1.50/MMBtu in December. To avoid a loss in the value of their reserves, Company A will sell into a futures market at a price of \$1.50/MMBtu. By doing so, they have guaranteed the value of their reserve at \$1.50/MMBtu. At the end of December, suppose that gas futures have fallen to a price of \$1.00/MMBtu (along with the spot price). Company A can then buy back its futures contract at a lower price than it sold it for. Thus, the profit of \$0.50/MMBtu will completely offset the loss they took on the physical value of the gas reserve.

A similar scenario would hold if futures and spot prices had increased by December. If the futures price had increased to \$2.00/MMBtu in December, Company A's loss of \$0.50/MMBtu in futures would be offset by it's gain in the value of their physical gas reserves. Notice that it does not matter if spot prices rise or fall, Company A ends up in the same position regardless since they engaged in futures trading. However, a "smart" hedger would only buy futures to cover part of his physical holdings is he thought prices might rise as he stand to gain more if all of his holdings were not locked in a at the previous price. In this light, one can see that there is a fine line between hedging and speculating.

When we say that someone is engaging in a long hedge, we mean that the trader is short in a commodity and wishes to purchase it sometime in the future. This is usually done by a trader who wants to protect himself from the risk of a price increase between the current time period and the period in which he will need to buy a commodity. The long hedge works like this: Suppose that Company B is a large industrial gas user which has to deliver it's product at a fixed price. However, for some reason, they fear that gas prices may rise in the near future and severely increase their operating costs. To protect itself from this risk, Company B will buy a futures contract to ensure that the price of gas when delivered will be the same as it is today say, for example, \$2.00/MMBtu. If prices actually rise, say to 2.50/MMBtu, then the increase in Company B's operating costs (\$0.50/MMBtu) will be offset by the profit made (\$0.50/MMBtu) in selling the gas futures contract back, just prior to maturity, at the new higher price. Thus, all gas price risk has been eliminated.

The third function of futures markets is to provide an investment opportunity. This is why speculators are involved in futures markets. Speculators have no interest in buying or selling the physical good that they are trading. They are involved in futures markets solely for financial gain. They take on the risk associated with these markets to achieve that end.

If a speculator believes that prices will rise, he will go long in futures (buy), planning to sell the contract before its maturity date at the higher price. Conversely, if he believes that prices will fall, he will open short in futures (sell) and plan to buy the contract back later at the (expected) lower price. While these seem like simple strategies, many speculators will engage in more complicated processes and open and short in contracts with different maturity dates. This is known as trading in spreads. By doing so, the speculator is actually hedging, as losses in one contract can be offset by gains in another. The only thing that matters to the speculator is the relative spread between the contracts and when the optimal time to close is. Obviously, when the trader believes that the spread is in his favor to the maximum extent possible, he will close his position in both contracts and reap the rewards. In essence, speculators supply liquidity to futures markets by assuming the risk that hedgers wish to shed.

In conclusion, this chapter has explained the need for futures markets, their operation, the specifics of the natural gas futures market and contract and the role of hedgers and speculators. Having done so, I have established the base for the chapters that follow.

Chapter 3

Cyclical Properties of Natural Gas Prices

3.1 Introduction.

At this point, an interesting application is to look at the cyclical behavior of natural gas prices. Specifically, we are interested in "detrending" our time series data. By doing so we can look for cross correlations between natural gas prices and each of output, the consumer price index and the unemployment rate. This will allow us to determine if natural gas prices are procyclical, countercyclical or acyclical with these specific economic variables. If natural gas prices are procyclical, this means that they move in the same direction as the cyclical variable we are interested in. If natural gas prices are countercyclical, they move in the opposite direction to the cyclical variable we are interested in. Finally, if they are acyclical, natural gas prices have no discernible relationship with the cyclical variable we are looking at. Also, we will be able to determine whether natural gas prices tend to lead or lag the cycle of the cyclical variable we are interested in.

The reason we have an interest in taking up such an exercise is that the cyclical behavior of natural gas prices has important implications for the hedging and speculative fields. If, for example, natural gas prices seem to be procyclical and lagging output, any mention of a change in government fiscal policy will cause hedgers and speculators to act differently than they might have been planning as they will have a better idea about which way prices will tend to move.

A specific example would work along these lines. Suppose the Clinton administration announced an increase in government spending, effective immediately. When a speculator or hedger hears this, he will undertake the actions in the futures market that he believes will benefit him the most. To choose the appropriate action, he will have to know how natural gas prices are related to output. If history tells us that natural gas prices are procyclical with output and lag output by three periods, market participants will know that natural gas prices will likely rise as a result of this government spending increase and peak three periods from now.

A speculator will buy gas futures immediately as he knows that there is high probability that their price will rise. He will likely then sell them three periods from now as he realizes that this is when they generally tend to peak. A hedger short in natural gas faced with the same scenario would want to buy futures to cover the amount of natural gas he wants to buy in the future to protect himself from a price increase. A hedger long in natural gas may now wish to hedge a smaller portion of his physical holdings, in the case of an unlikely price drop. He will subject the other portion of his holdings to price risk as it is most likely that prices will rise.

As a further application, we can run a similar exercise but compare natural gas prices with crude oil prices, which will give us a tool in which we can determine the relationship between the two and what the implications are. The reason I use crude oil as a means of comparison is that it is a commodity traded with high volume and is generally considered the "staple" of energy products.

3.2 Methodology.

To begin such an analysis, we must first detrend our time series data in order to obtain the cyclical component and render the series stationary. To do so, we follow Prescott's (1986) methodology and use the Hodrick - Prescott (HP) filter. This is by no means the only method of detrending time series. For example, taking the first differences of a series will render it stationary if it is integrated of order one. Also, one may wish to use a linear time trend if there is no unit root present in a particular series. With this in mind, I chose to use the HP filter and also first difference the series then compare the results of the two methods. What follows is a description of the HP filter and how it works.

We start by realizing that the term "trend" is rather ambiguous and needs more structure. Kydland and Prescott (1990) believe that "the trend component of a time series should be approximately the curve that a student of business cycles and growth would draw through a time plot of this time series." With this in mind, they construct the trend component of the series in the following manner:

For the logarithm of any time series X_t (where t = 1, 2, ... ,T), we can denote the trend of the particular time series as τ_t (t = 1, 2, ... , T)

They then denote the mean square deviation of the actual time series from the trend of the series as:

$$\text{(1)} \qquad MSE = \sum_{t=1}^{T} (X_t - \tau_t)^2$$

To obtain a least squares time trend from this, we want to minimize (1) such that the change in the trend values between periods $(\tau_t - \tau_{t-1})$ is constant for all periods. Prescott chooses to relax this somewhat, constructing his constraint to look like:

$$\text{(2)} \qquad \sum_{t=2}^{T-1} ((\tau_{t+1} - \tau_t) - (\tau_t - \tau_{t-1}))^2 \leq \lambda$$

We then choose λ to reflect how "tight" a fit, relative to the series, that we desire for our trend values. A large value for λ will result in close tracking of the data. Simply put, equation (2) is the sum of squares of second differences of the trend. Thus, variations in the trend growth component will be penalized. It follows that if λ is relatively large, the penalty will be reduced since the trend line will exhibit more variation, and will more closely resemble the series we are tracking. If we allow λ to approach zero, Prescott's methodlogy reduces to a least squares time trend.

Combining (1) and (2) we arrive at the following minimization problem:

$$\text{(3)} \qquad MIN \sum_{t=1}^{T} (X_t - \tau_t)^2 - \mu[\sum_{t=2}^{T-1} ((\tau_{t+1} - \tau_t) - (\tau_t - \tau_{t-1}))^2 - \lambda]$$

Where μ is the Lagrange multiplier of the constraint. This makes the second term of equation (3) the multiple, μ, of the sum of the squares of the trend component's second differences. Thus, if the value of μ is relatively large, the penalty for variations in the growth rate of the trend component will also be large. Using simple mathematics one could solve for the values of μ and the trend, τ, if we were given the time series X_t and a value for λ. However, it is much easier to simply specify the value of μ to ease in computation. For monthly data, Prescott (1986) suggests that the value of 14,400 be used. λ is then solved for when the trend is known. One can see that small values of λ give us a tight constraint and a large value of μ.

Once this approach is implemented, we can then solve for the cyclical component of the series by simply subtracting the trend from the value of the actual time series. We can then calculate the degree of contemporaneous comovement between the cyclical component of energy commodity prices and the cyclical variable that we are interested in by looking at the value of the correlation coefficient between the two cyclical components. I will denote the correlation coefficient as $\rho(j)$, $j \in \{0,\pm1,\pm2,...\}$. The "j" term simply refers to a lead or lag or the series. We derive the coefficient by first taking the covariance between two time series, here denoted as X and Y, by the following formula:

(4) $$\mathrm{COV}\,(X,Y) = \frac{1}{N-1} \sum_{t=1}^{N} (X_t-\overline{X})(Y_t-\overline{Y})$$

Next, we calculate the degree of correlation between the two series utilizing the following formula:

(5) $$\rho(0) = \mathrm{COR}(X,Y) = \frac{\mathrm{COV}(X,Y)}{(\sigma_X \sigma_Y)^{1/2}}$$

The value of the coefficient at $\rho(0)$, (The column X_t in tables 3.1 to 3.4), tells us the degree of contemporaneous comovement between the cyclical component of the natural gas price series and the cyclical component of the variable we are interested in. If a given value of $\rho(0)$ is positive, we say that the particular series is procyclical. Conversely, if $\rho(0)$ is negative, we say that the series is countercyclical. Finally, if $\rho(0)$ is zero, the series is said to be acyclical.

Next, we can further define the relationship between the series, following Fiorito and Kollintzas (1994). They suggest that for $0.5 < |\rho(0)| \leq 1$, $0.2 < |\rho(0)| \leq 0.5$, and $0 \leq |\rho(0)| \leq 0.2$, we say that the series is strongly contemporaneously correlated, weakly contemporaneously correlated, and contemporaneously uncorrelated with the cycle.

When we look at the leads (X_{t+i}) and lags (X_{t-i}) of the series, the correlation coefficient will give us information about whether a phase shift is present in the movement of a time series relative to the cycle. For example, if a given series has a positive number in the X_t column, but has a larger positive number in the X_{t-3} column, we can say that the series is procyclical but peaks three periods before the cyclical variable. In such a case, we can say that the series leads the cycle by three periods.

Conversely, if a given series has a negative number in the X_t column but has a larger (in absolute value) negative number in the X_{t+3} column, we can say that the series is countercyclical but peaks three periods after the cyclical variable we are interested in. Here, we say that the series lags the cycle by three periods.

Before going further, it is of some importance to defend the use of the HP filter. Recent allegations, such as those brought forth by King and Rebelo (1993) and Cogley and Nason (1995) seriously question the effect the HP filter has on the time series when transforming the data. Specifically, they argue that the filter may seriously affect the measurement of the variability and comovement between variables and that the interpretation of results may rely heavily on original assumptions of the properties of the time series in question. However, recent work by Baxter and King (1995) tells us that the use of the HP filter is non distortionary and will not affect our results in an adverse manner.

Having established what the HP filter is and how it is used, we can now look at the data and results..

3.3 Data and Results.

3.3.1 Cyclical Correlation with Output, Prices and Unemployment.

This section of the thesis begins by discussing the data. This section will relate only to the cyclical correlations of the log of natural gas prices with output, prices and the unemployment rate. Discussion of the results comparing the natural gas price series with crude oil prices will appear in the next section.

Although daily observations were available for the natural gas price series, only monthly observations were available for the output, prices and unemployment series. Thus, I had no choice but to use monthly observations in my analysis. Looking at figures 3.1 and 3.2, one can see the daily and monthly natural gas observations respectively. There are 62 monthly observations and 1287 daily observations for natural gas prices, starting at April 3, 1990 and ending May 15, 1995.

Taking monthly observations of all the variables, I matched the log of the natural gas price level with the United States production index (Y), the United States consumer price index (P), and the United States unemployment rate (U). Unfortunately, the data I obtained on the Y, P, and U variables were only available to April, 1993. Natural gas began trading on the NYMEX in April 1990. Cross referencing the two samples gives us 37 observations. After arranging the data, I used the HP filter to detrend the data, and calculated the correlation coefficients between the cyclical components of the series and the pertinent cyclical variable. Looking at figure 3.3, one can see a graph depicting the actual monthly natural gas time series, the calculated trend component of the series, and the derived cyclical component of the series.

Results are reported in table 3.1. Under the column labeled "volatility" I have listed the standard deviations for output, the consumer price index and the unemployment rate. The columns labeled X_{t-6} to X_{t+6} show the correlation coefficients of output, prices and the unemployment rate as they compare to the leads (+) and lags (-) of the natural gas price series.

Referring to table 3.1, one can see the results for the cyclical correlations of output, prices and unemployment with the natural gas price series. Looking at the X_t column, we can see that natural gas is weakly procyclical with output (0.34) and acyclical with prices (0.04). Since the cyclical component of output and unemployment are negatively correlated, a negative sign in the unemployment row indicates procyclical movement and a positive sign indicates countercyclical movement. In this case, the value of -0.29 indicates weak procyclical movement.

These results are not terribly surprising as one would expect output and natural gas prices to move together. If output increases, this is a signal of industrial expansion. Associated with this is an increased demand for natural gas as more of it is needed for use in factories, to heat new buildings, et cetera. One will notice that when we look in the X_{t+2} column, we see a value of 0.53 which is larger than the value of 0.34 in the X_t column. This means that natural gas prices lag the cycle by two periods, ie. this means that output tends to peak two periods before gas prices do. This should come as no surprise since natural gas is used primarily as a heating tool. Take the case of a new home being built. Obviously, it will not be heated until after its construction is complete. In case such as this, the increase in output which is caused by the actual construction, will show up earlier than the increased demand for gas (and gas price increases) since natural gas is not needed to heat the home until the final phase of the construction is complete or almost complete.

This has important implications for hedgers and speculators. From a speculator's standpoint, if he witnesses an increase in an economic indicator such as real GDP (a proxy for output), he may wish to buy natural gas futures now as their price will likely rise a short time after the increase in GDP took place. He can then reverse his position when prices do rise and collect on a windfall gain. From the view of a hedger short in natural gas, he may wish to sell into the futures market to cover the entire physical quantity of his physical stock of natural gas. He would do this if GDP had fallen in order to protect the value of his asset because it is likely that natural gas prices will also fall. Conversely, if GDP has risen, he may wish to sell futures to hedge only part of his physical holdings, as prices will likely rise.

The acyclicality with prices is also not too surprising as there is absolutely no reason to expect that if prices of other commodities move that natural gas prices will too. Natural gas, by its very nature, tends to be linked to very few other commodities and is used primarily as a heating and cooling tool. However, I would not expect the acyclicality of natural gas prices with other prices to hold for too many more years as natural gas is being used more creatively every year. For example, a growing number of people are starting to use natural gas to fuel their cars. While the numbers doing so are certainly small, they are certain to grow in an increasingly energy conscious society. Besides being a cleaner burning fuel than unleaded gasoline, its most attractive quality is that it is also comparatively cheaper to fuel your car with natural gas (although there is certainly a trade off in the power your engine will exhibit).

Next, we can test just exactly how robust our results are by first differencing the series. This is simply another method of trend elimination. By doing so, we are assuming that our data series is I(1), or stationary in the first differences of the series. These results are reported in table 3.2.

From table 3.2, we can see that natural gas prices are shown as acyclical with each of output, prices, and the unemployment rate. Notice that the volatility of the first differenced series is less than in the series that were run through the HP filter. When first differencing a time series, we are essentially taking the "long run" component out of the series. Realizing this fact and combining it with the results from the data, we can conclude that our price series is quite likely I(1).

Notice that the results of the two different methods of trend elimination are quite different. Since the two different methods give quite different results, neither can be deemed reliable. Taken in this context, my results are ambiguous with respect to this particular application. I will now move on to discuss the results of the other testing that was done.

3.3.2 Cross Price Correlations.

For this section, I examine how natural gas prices correlate with other energy commodity prices. To reach this end, I used crude oil prices as a benchmark and ran crude oil prices with natural gas prices separately through the HP filter. I then lagged and led the natural gas price by six periods. After doing so, I calculated the correlation coefficients of these leads and lags with crude oil prices from the current reference period. Matching the sample of monthly natural gas prices with the crude oil sample gives us a sample size of 62. To put these results in context and provide a means of comparison, I added two other energy commodities, heating oil and unleaded gasoline. I have 136 observations for heating oil, and 116 for unleaded gasoline. These results are reported in table 3.3.

Not surprisingly, heating oil and unleaded gasoline prices showed a high correlation with crude oil prices, with values of 0.88 and 0.86 respectively in the X_t column. Since both use crude oil in their very own formation (both heating oil and unleaded gasoline are derived through the refining of crude oil), one would expect these large values. Notice also that both these commodity prices neither lead nor lag crude oil prices. This is because a change in crude oil prices has an almost immediate effect on the prices of heating oil and unleaded gasoline. A change in crude oil prices will be immediately reflected in the futures prices of these commodities as the refining costs to make them will have risen or fallen depending on the crude oil price movement.

Natural gas prices seem to have little relation to oil prices, with a correlation coefficient of only 0.33. This weak positive (procyclical) correlation may be an indicator telling us that energy prices in general tend to move in the same direction. Natural gas prices also tend to peak one period after crude oil prices.

This has important implications for speculators. If one were to witness oil prices rising, it is likely that natural gas prices will rise shortly after. It would be best for a speculator witnessing a rise in crude oil prices to buy natural gas futures immediately and then sell when natural gas prices go up.

When first differencing the price series, one finds that the results do not change much. Referring to table 3.4, the correlations for heating oil and unleaded gasoline fall slightly but still remain high. The natural gas correlation coefficient rises slightly and we see that gas prices no longer lag oil prices. Notice that both methods of trend elimination give similar results. This would indicate that our results are quite reliable.

3.4 Conclusion.

The goal of this chapter was to examine the cyclical behavior of natural gas prices. The reason for doing so was to shed some light on how hedgers and speculators could use this information to aid in their activities in futures markets. To obtain the cyclical component of the time series, I used the HP filter to detrend the data. The results using the filter were as follows:

Comparing the cyclical component of natural gas prices to that of output, prices, and the unemployment rate, I found that natural gas prices were procyclical with output and the unemployment rate but acyclical with prices. Natural gas prices were also found to lag output by two periods. With natural gas used primarily as a heating fuel, these results were consistent with what was expected.

When comparing crude oil and natural gas prices, a weak procyclical correlation was found. I interpreted this as an indicator that, in general, energy prices tend to move together. Heating oil and unleaded gasoline prices showed high correlations with crude oil prices. These high correlations can be attributed to the fact that both commodities use crude oil in their formation. Thus, any move in crude oil prices will be directly reflected in heating oil and unleaded gasoline prices as their refining costs have changed.

To test how robust the results actually are, I used first differencing as an alternative method of trend elimination and compared the results of the two approaches. The relation of natural gas prices to output, prices and the unemployment rate turned out to be acyclical in all three cases. This is quite different than the results derived using the HP filter. Since different methods of trend elimination give quite different results, one must question the reliability of these results. If I were to put my faith in one method or the other, I would likely side with the HP filter as it is the more advanced statistical technique.

When looking at the cyclical relationship between crude oil and other commodities, the results are quite similar using either first differencing or the HP filter. This would indicate that, in this case, our results are quite reliable.

Table 3.1

Hodrick Prescott Cyclical Correlations of Output, Prices
and the Unemployment Rate with Natural Gas Prices

Correlation Coefficients

Variable	Volatility	X_{t-6}	X_{t-5}	X_{t-4}	X_{t-3}	X_{t-2}	X_{t-1}	X_t	X_{t+1}	X_{t+2}	X_{t+3}	X_{t+4}	X_{t+5}	X_{t+6}
Y	0.013	0.04	-0.02	-0.06	-0.02	0.06	0.19	0.34	0.46	0.53	0.51	0.43	0.25	0.01
P	0.004	0.23	0.24	0.24	0.22	0.18	0.10	0.04	-0.04	-0.17	-0.36	-0.51	-0.58	-0.48
U	0.048	0.22	-0.25	-0.19	-0.12	-0.17	-0.21	-0.29	-0.28	-0.28	-0.31	-0.37	-0.29	-0.11

Table 3.2

Correlations of First Differences of Output, Prices
and the Unemployment Rate with Natural Gas Prices

Correlation Coefficients

Variable	Volatility	X_{t-6}	X_{t-5}	X_{t-4}	X_{t-3}	X_{t-2}	X_{t-1}	X_t	X_{t+1}	X_{t+2}	X_{t+3}	X_{t+4}	X_{t+5}	X_{t+6}
Y	0.006	0.09	-0.20	-0.23	-0.07	-0.09	-0.02	0.07	0.15	0.30	0.22	0.25	0.09	-0.15
P	0.002	-0.03	0.05	-0.02	0.03	0.13	0.00	0.09	0.20	0.06	-0.15	-0.21	-0.37	-0.13
U	0.023	-0.17	-0.07	0.01	0.25	-0.14	0.02	-0.16	0.02	0.01	0.00	-0.35	-0.18	-0.08

Table 3.3

Hodrick Prescott Cyclical Correlations of Crude Oil Prices with Natural Gas
Heating Oil and Unleaded Gasoline Prices

		Correlation Coefficients												
Commodity	Volatility	X_{t-6}	X_{t-5}	X_{t-4}	X_{t-3}	X_{t-2}	X_{t-1}	X_t	X_{t+1}	X_{t+2}	X_{t+3}	X_{t+4}	X_{t+5}	X_{t+6}
Natural Gas	0.185	-0.18	-0.28	-0.38	-0.35	-0.18	0.09	0.33	0.38	0.27	0.03	-0.22	-0.33	-0.37
Heating Oil	0.142	-0.14	-0.14	-0.06	0.09	0.37	0.66	0.88	0.76	0.60	0.42	0.27	0.09	0.01
Unleaded Gasoline	0.157	-0.10	0.00	0.13	0.23	0.47	0.72	0.86	0.66	0.42	0.29	0.25	0.16	0.09

Table 3.4

Hodrick Prescott Cyclical Correlations of First Differences of Crude Oil Prices
with Natuaral Gas,Heating Oil and Unleaded Gasoline Prices

		Correlation Coefficients												
Commodity	Volatility	X_{t-6}	X_{t-5}	X_{t-4}	X_{t-3}	X_{t-2}	X_{t-1}	X_t	X_{t+1}	X_{t+2}	X_{t+3}	X_{t+4}	X_{t+5}	X_{t+6}
Natural Gas	0.141	0.18	-0.06	-0.27	-0.26	-0.22	0.04	0.39	0.33	0.24	-0.02	-0.26	-0.08	-0.13
Heating Oil	0.099	0.04	-0.20	-0.13	-0.29	-0.02	0.17	0.79	0.11	0.04	-0.07	-0.09	-0.22	-0.06
Unleaded Gasoline	0.114	0.01	0.02	0.06	-0.27	-0.03	0.26	0.74	0.08	-0.19	-0.20	0.12	-0.03	0.13

Figure 3.1. Natural Log of Daily Natural Gas Prices

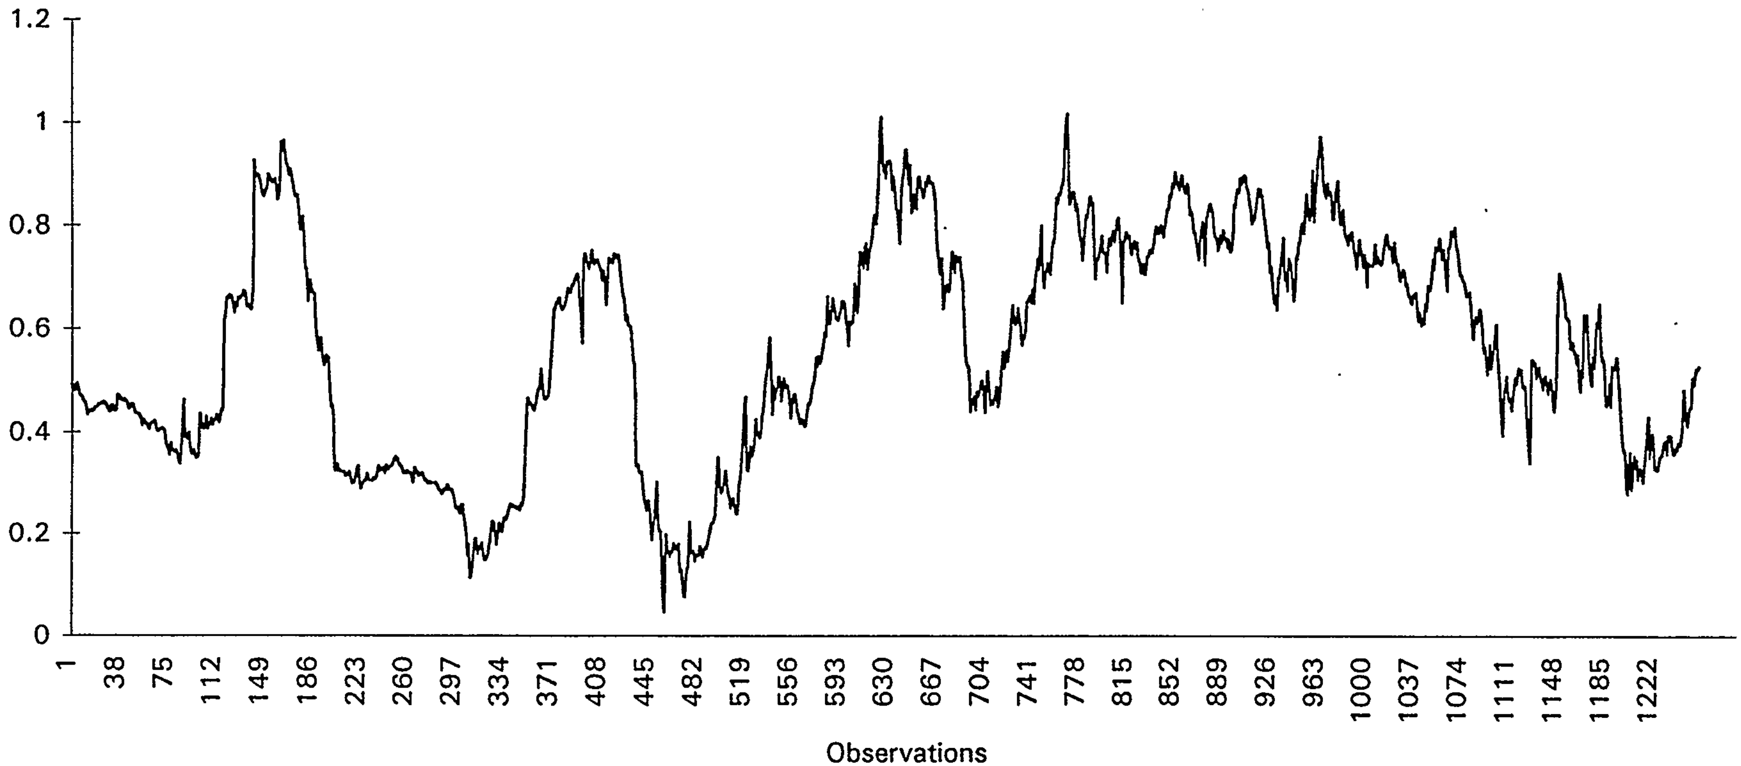

Figure 3.2. Natural Log of Monthly Natural Gas Prices

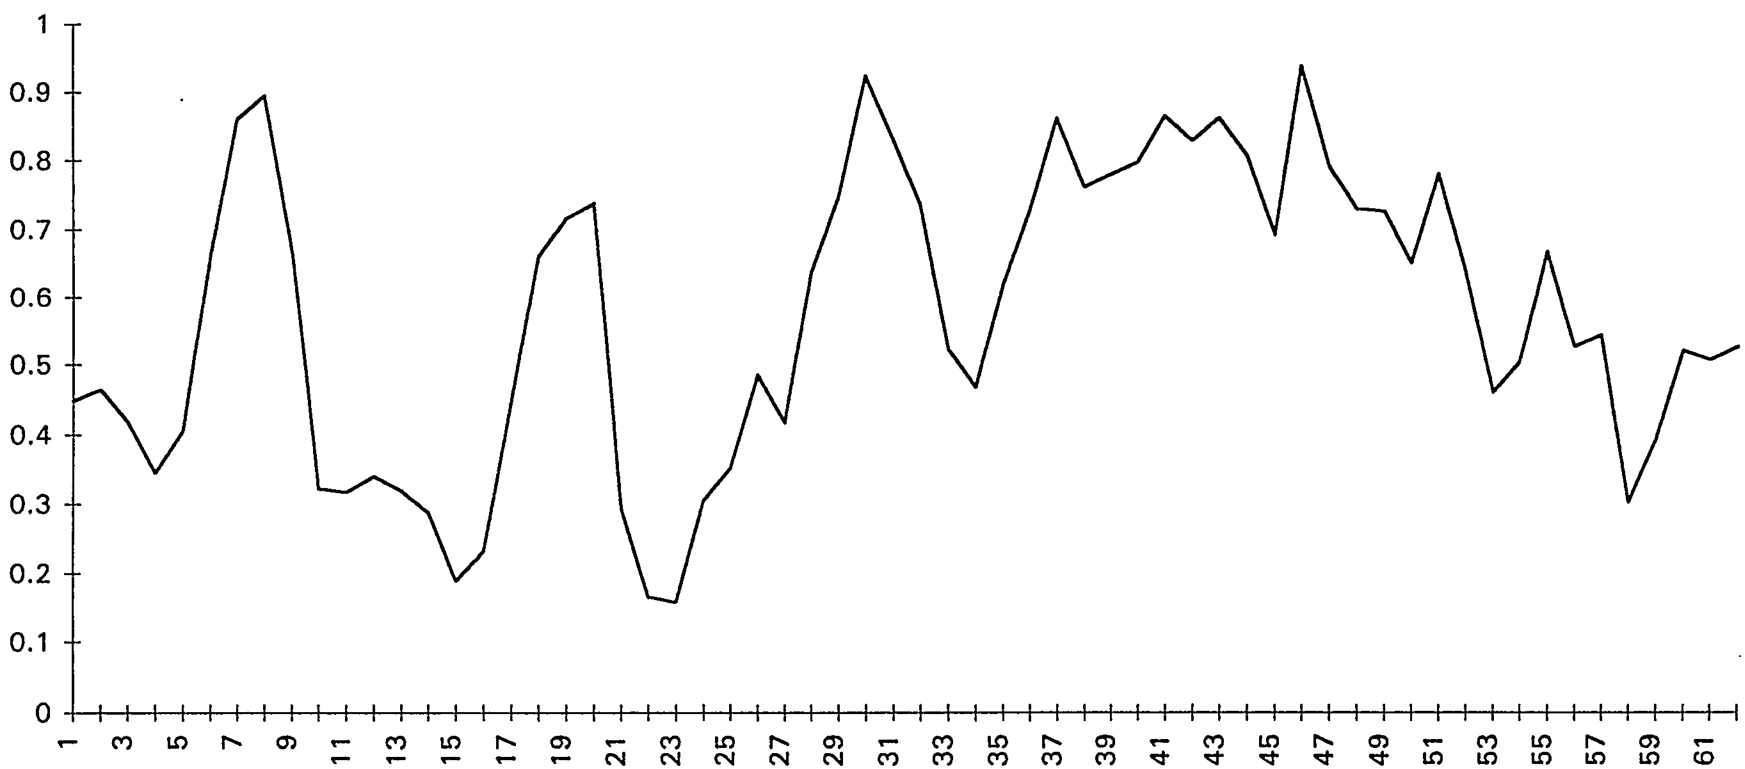

Figure 3.3 Trend and Cyclical Components of the Natural Log of Monthly Natural Gas Prices

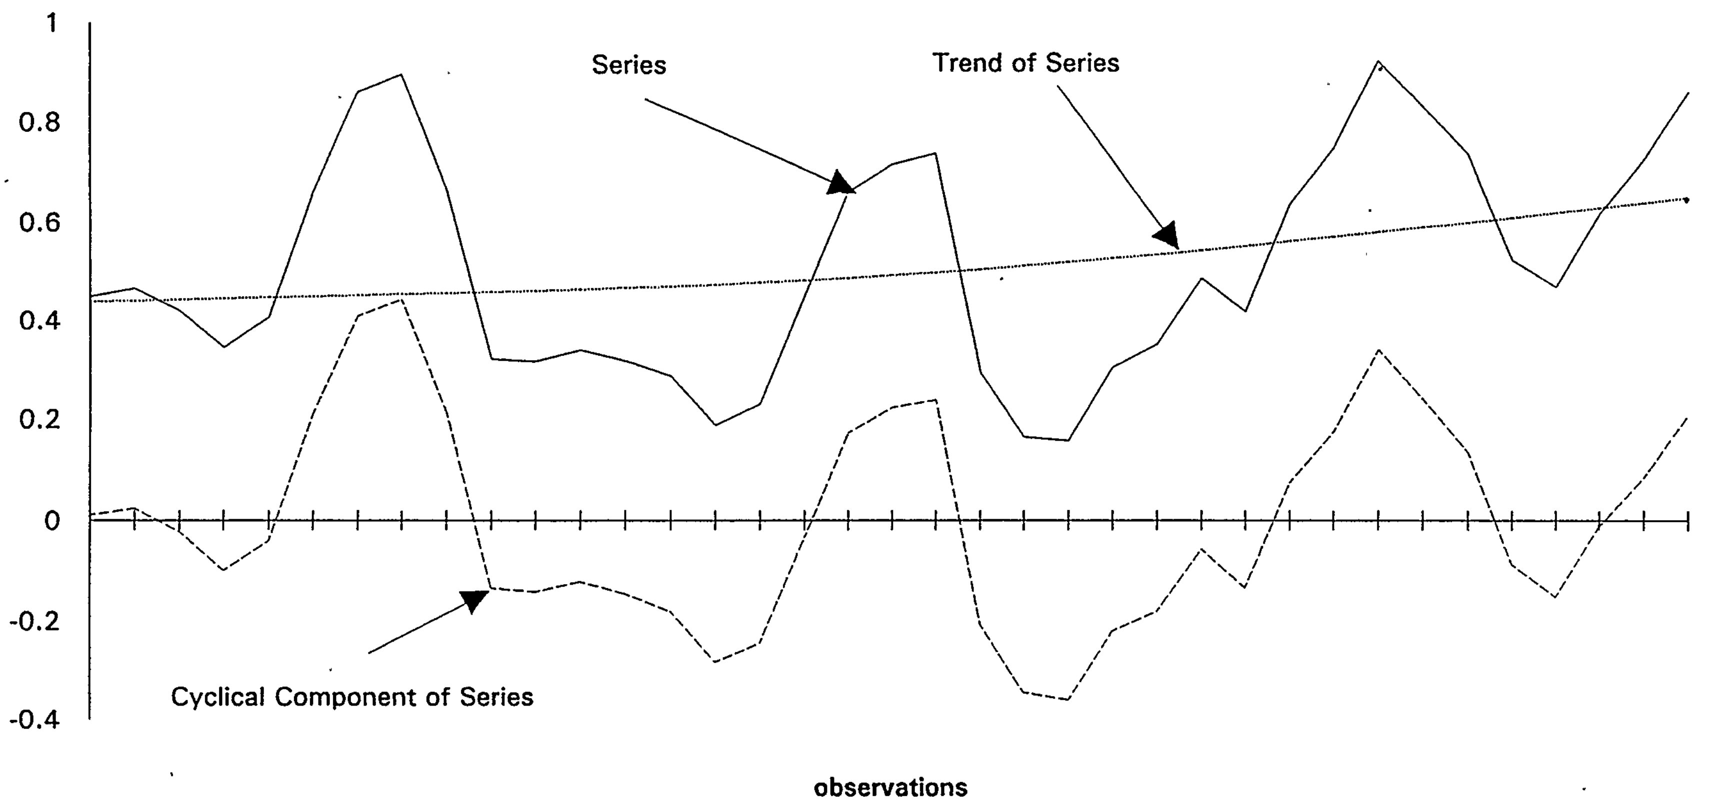

Chapter 4

The Theory of Storage

4.1 Introduction.

The theory of storage was originally developed by Working (1949) and then expanded by Brennan (1958) and Tesler (1958). The theory was developed to explain why commodities were being stored when the observed price of (or return on) storage was negative. The authors concluded that the reason firms engaged in storage when a negative return on storage was present was that the stored commodity provided a non monetary benefit to the party storing it. They called this benefit a "convenience yield." Their theory goes on to explain why futures prices are below spot prices before harvests when inventories are low and the marginal convenience yield on inventory is high. From their discussion, three predictions evolve, which try to predict the behaviour of spot and futures prices when a supply or demand shock hits the market.

The most important reason that the theory of storage was advanced was to analyze commodities whose supply fluctuated in seasonal patterns. Mainly, this analysis was focused on agricultural commodities subject to a harvest, but it has recently been extended to metals (see Fama and French (1988)) and energy products (see, for example, Serletis (1994) and Cho and McDougal (1990)). Since natural gas exhibits seasonal variations in supply, with demand and supply peaking in the winter months, an interesting application is to test the theory of storage and its predictions using data for natural gas.

4.2 Theoretical Foundations.

Early attempts to test the theory of storage used inventory data and market prices. However, this method was inherently flawed as it was difficult to define what inventory actually was and how to measure it. As an alternative method of measurement and testing, Fama and French (1988) developed a method relating the relative variation of spot and futures prices.

Following Fama and French (1988), we can define the price of storage as the difference between future and spot prices. If we let F(t,T) be the futures price at time t for delivery of the natural gas at time T, and the spot price of natural gas at time t be S(t), we can define the price differential to be:

$$F(t,T)-S(t) \tag{1}$$

Equation (1) is also known as the basis. If the futures price is above the spot price (a positive basis), we say that the return on storage is positive. This should act as a signal to store the commodity, in this case, the natural gas, as one can expect a higher price in the future. Conversely, a negative basis should be a signal to draw natural gas out of storage as prices will probably fall; yet, this is not always observed.

In the case of a negative storage price, Working (1949) explains that the commodity may actually be stored because the non monetary benefits of storage are substantial (a convenience yield is present). For example, Cho and McDougal (1990) point out that a sizable inventory may allow a firm to change its production schedule without incurring extraordinary marginal or adjustment costs.

Continuing on, Fama and French (1988) believe that the current futures-spot price differential should equal the interest forgone during storage, S(t)R(t,T), plus the marginal warehousing cost, W(t,T), minus the marginal convenience yield, C(t,T). Thus, the price differential equation will take the following form:

$$F(t,T) - S(t) = S(t)R(t,T) + W(t,T) - C(t,T) \tag{2}$$

where R(t,T) is the interest rate that any participant in the natural gas market can borrow at. If we divide both sides of the equation by S(t) and rearrange, we arrive at the following equation:

$$\frac{F(t,T) - S(t)}{S(t)} - R(t,T) = \frac{W(t,T) - C(t,T)}{S(t)} \tag{3}$$

The left hand side of equation (3) is known as the interest adjusted basis (IAB) and is equal to the relative warehousing cost minus the relative convenience yield. See figure 4.1 (next page) for a graphical representation of the IAB as it corresponds to inventory levels.

Brennan (1958) and Tesler (1958) note that since storage costs increase as storage space declines, we know that the convenience yield is inversely related to the total stock available and will tend to fall at a decreasing rate ($\partial C/\partial I<0$ and $\partial^2 C/\partial I^2>0$, where I is the inventory level). See figure 4.2 (page 42) for a graphical representation of the convenience yield as it corresponds to inventory levels. The authors assume that the marginal warehousing cost for natural gas is relatively constant over the relevant range of inventory, that marginal convenience yield dominates variation in the marginal warehousing cost and then combine this with the fact that the marginal convenience yield falls at a decreasing rate as inventory rises. This enables us to test some hypotheses about the convenience yield and the impact of supply and demand shocks.

Figure 4.1

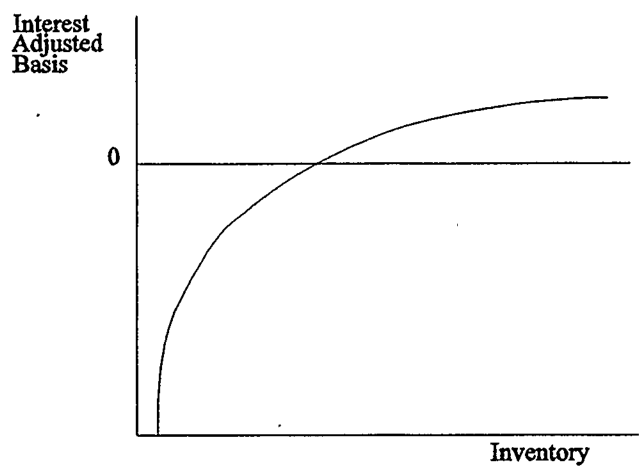

Figure 4.2

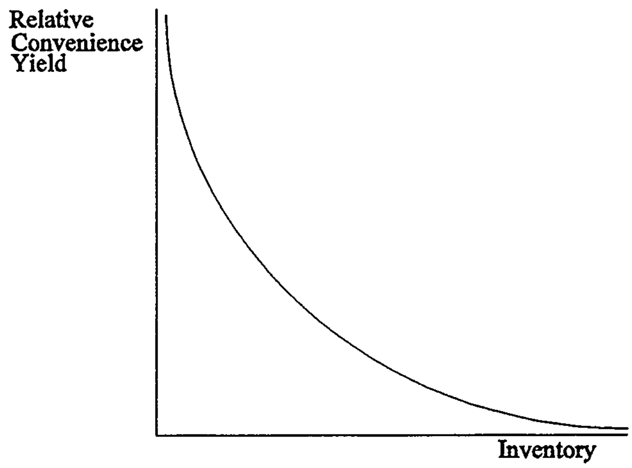

With the knowledge that price elasticities of supply and demand are smaller in the short run than in the long run, we expect that supply and demand shocks will generate greater variability in spot prices than futures prices. The relative differences in variability will be the greatest when inventories are low (convenience yield is high and the IAB is negative) and unable to accommodate the shock. Conversely, the relative variability in spot and futures prices will be the smallest when inventories are high (the convenience yield in low and the IAB is positive) as the high inventory will help to accommodate this shock; this would indicate that changes in spot prices are permanent, that is, they closely match those of futures price.

Three factors will tend to dictate how much a demand shock will affect spot and futures prices: the size of the shock, the levels of inventory at the time of the shock and the shape of the convenience yield curve.

As an example, suppose a large demand shock hits when inventories are quite high. Since a high inventory level corresponds with us being on the "flat" portion of both the convenience yield and IAB curves (see figures 4.1 and 4.2) the effect of the demand shock will have little effect on either the convenience yield or the IAB. Inventories are drawn down in response to the shock but we remain on or close to the flat portion of the curves. Since the current inventory will satisfy the amount required by the demand shock, this indicates that changes in spot prices are permanent.

However, if we take this scenario of a demand shock but apply it to a situation where inventory levels are low, spot and futures prices should behave quite differently. From figure 4.1, we see that a low inventory level corresponds to a high convenience yield. Further, as inventory levels drop, the convenience yield curve rises quickly. Thus, when faced with a demand shock, those firms with small inventories will be reluctant to deplete them further. As a short term result, spot prices will change dramatically to compensate for the excess demand for natural gas. However, the change in futures prices should be less volatile than the change in spot prices since market participants will be able to anticipate supply and demand responses.

With these scenarios in mind, Fama and French (1988), use the sign of the IAB as a proxy for the inventory level. When inventories are high, the marginal convenience yield is low and the IAB is positive. Conversely, when inventories are low, the marginal convenience yield is high and the IAB is negative. By using the IAB as a proxy for inventory levels, we get around the problem of the measurement of actual inventory.

Fama and French (1988) go on to test three variations of the theory of storage, all of which I will reproduce with respect to natural gas prices.

4.3 Tests on the Theory of Storage.

The first test developed by the authors tests the theory of storage's prediction that supply and demand shocks cause more volatility in spot and futures prices when inventory levels are low. This implies that the IAB is more variable when it is negative. To test this, they compare the standard deviations of daily changes in the negative IAB sample to the standard deviation of daily changes in the positive IAB sample. They use a F-test to see if the null hypothesis of equal variance of changes in the positive and negative IAB samples holds. If we reject the null hypothesis of equal variances, the theory of storage has made a correct prediction.

The next test tests the prediction that supply and demand shocks cause approximately equal changes in spot and futures prices when inventory levels are high (positive IAB), but shocks cause spot prices to change more than futures prices when inventory levels are low (negative IAB). To accomplish this task, they compare the ratios of the standard deviations of percent futures prices changes to the standard deviation of percent spot price changes and compare these across the positive and negative IAB samples. The theory of storage predicts that the calculated ratio should be close to one when the IAB is positive and less than one when the IAB is negative

Finally, the authors test the prediction that supply and demand shocks produce larger changes in near term futures as compared to longer term futures. This is done to test the work originally developed by Samuelson (1965) who believed that futures prices were less volatile than spot prices since market participants will have time to anticipate supply and demand responses after a shock. Once again, Fama and French use the ratio of the standard deviation of percent futures prices change to the standard deviation of the percent spot price change. The theory of storage predicts that this ratio should fall as the maturity date increases for both the positive and negative IAB samples.

4.4. Data

For the purposes of this chapter, daily one, two, four and seven month natural gas futures traded on the NYMEX were used. Also, daily one, three and six month maturity U.S. T-bill rates, obtained from the Bank of Canada were used. I matched these two samples, eliminating any trading date with data missing. The sample period is from April 1990 to May 1995.

For this chapter, it was necessary to use proxy futures prices. Specifically, the one months futures price was used as a proxy for the spot price, the two months futures price as a proxy for the one month futures price, and so on. This is done as the spot price is not available for many trading days. Further, Cho and McDougal (1990) note that:

> "the spot price and the first nearby futures price are so close that the first nearby month contract is referred to as the 'spot contract'. Thus, the difference between two nearby futures prices, rather than the difference between the spot and the first nearby futures price is an appropriate measure of time basis."

To test the theory of storage, both daily and weekly series are analyzed. To construct a weekly series, I averaged each of the futures price and the T-bill rates over a calender week. Weekly series are needed for several reasons. Since, on a number of occasions, some trading days had to be eliminated, one would expect abrupt price movements around the data point that was eliminated as compared to a week where no data point was eliminated. By averaging the series over a week, it tends to smooth these sometimes volatile price movements. Secondly, Cho and Mcdougal (1990) note that there tends to be more volatile price movement as futures contracts come close to (or shortly after) reaching their maturity date. This is because traders are trying to offset their positions. Once again, averaging the series over the calender week will smooth out these price movements. Further, studies such as those by Chang and Kim (1988) show that price volatility tends to be highest on Mondays as compared to the other days of the week. Averaging eliminates this characteristic. Although averaging is the superior method to use for this kind of testing, I have also included daily data in my analysis in order to be as complete as possible.

Looking at table 4.1, we can see the number of positive, negative and total observations of the natural gas IAB for both daily and weekly data.

Table 4.1

Number of Positive, Negative and Total Observations
of the Interest Adjusted Basis
Daily and Weekly Data

	Daily Data			Weekly Data		
Contract	Positive	Negative	All	Positive	Negative	All
1-Month	762	511	1273	165	103	268
3-Month	801	472	1273	165	103	268
6-Month	773	500	1273	160	108	268

In order to conduct any testing on the theory of storage, it is necessary to define our variables and show how they were derived. I calculated the one (IAB1), three (IAB3) and six (IAB6) month IAB utilizing the following formulas:

$$IAB1 = 12\ [Ln(F2/F1)] - R1/100$$

$$IAB3 = 04\ [Ln(F4/F1)] - R3/100$$

$$IAB6 = 02\ [Ln(F7/F1)] - R6/100$$

where F1, F2, F3, and F7 correspond to the number's monthly maturity futures price (for example, F7 equals the seven month futures price). Also, R1, R3, and R6 correspond to the numbers monthly T-bill rate (for example, R3 equals the three month T-bill rate). Figure 4.3 shows the weekly one, three, and six month futures prices. Figures 4.4, 4.5, and 4.6 show the one, three, and six month weekly IAB's respectively.

4.5 Test Results.

4.5.1 Variability of the IAB and Inventory Levels.

The first test of the theory of storage is based on the prediction that supply and demand shocks cause more variability in spot and futures prices when inventory levels are low. Since inventory levels are proxied using the IAB, this implies that the IAB is more variable when it is negative. To test this hypothesis, we simply compare the standard deviation of changes in negative IAB observations to the standard deviation of changes in positive IAB observations. Table 4.2 shows the average values of the IAB for both daily and weekly data, and table 4.3 shows the standard deviation of changes in the IAB for both daily and weekly data.

Table 4.2

Average Values of the Interest Adjusted Basis
Daily and Weekly Data

	Daily Data			Weekly Data		
Contract	Positive	Negative	All	Positive	Negative	All
1-Month	0.63	-0.56	0.15	0.60	-0.58	0.14
3-Month	0.47	-0.56	0.09	0.47	-0.54	0.08
6-Month	0.32	-0.37	0.05	0.32	-0.37	0.04

Table 4.3

Standard Deviations of Changes in the Interest Adjusted Basis
Daily and Weekly Data

	Daily Data			Weekly Data		
Contract	Positive	Negative	All	Positive	Negative	All
1-Month	12.92	2.45	10.11	4.48	1.83	3.68
3-Month	6.05	2.05	4.97	2.74	7.55*	5.13
6-Month	1.29	15.56*	9.92	9.43	2.06	7.38

Note: Table 4.2 and 4.3 statistics are for observations when the interest adjusted basis is positive (Positive) and observations when the interest adjusted basis is negative (Negative). An asterisk indicates that an F-test rejects the hypothesis of equal variances at the 1 and 5 percent levels.

To be consistent with this test of the theory of storage, the standard deviations of changes in the IAB should be larger when the IAB is negative. However, the standard deviations are larger for the negative IAB sample in only two of the six cases; the six month daily contract, and the three month weekly contract. Remembering that the weekly data set is better suited for testing purposes, the weekly data set supports the prediction that the IAB is more variable when it is negative in only one of three cases.

A more robust test of seeing if natural gas passes the test of the theory of storage is to apply an F-test. Following Hulleman (1993), we construct the F-test to take the following form:

$$H_o: VAR1 = VAR2$$

$$H_1: VAR1 \neq VAR2$$

$$F_{CALC} = VAR1 / VAR2$$

$$\text{Reject } H_o \text{ if } F_{CALC} > F_{TABLES}$$

In essence, the F-test is testing if the variance in the positive IAB sample is equal to the variance in the negative IAB sample. If we reject the null hypothesis at a reasonable level of significance, we are rejecting the hypothesis that the variance of the IAB does not depend on its sign. If we reject the null hypothesis and the negative IAB value is larger than the positive IAB value, we can say that the market passes this indirect test of the theory of storage. Looking at the weekly sampled series in table 4.3, we can see that only the three month market passes this theory of storage, rejecting the null hypothesis of equal variances at the one percent level of significance. The daily sampled series performs similarly, only rejecting the null hypothesis on the six month market. However, as mentioned earlier, the daily sampled series gives less reliable results.

Thus, we can say that the natural gas market does not pass this test of the theory of storage.

4.5.2 Variability of Spot and Futures Prices, as Dependant upon Inventory Levels.

The second test of the theory of storage tests the prediction that supply and demand shocks cause approximately equal changes in spot and futures prices when inventory levels are high (positive IAB), but cause spot prices to change more than futures prices when inventory levels are low (negative IAB).

In order to test this, we compare the ratios of the standard deviations of percent futures price changes to the standard deviation of percent spot price changes and compare these between the positive and negative IAB samples. The theory of storage predicts that the calculated ratio should be close to one when the IAB is positive and somewhat less than one when the IAB is negative.

Table 4.4

Ratios of the Standard Deviation of Percent Futures Price Changes
to the Standard Deviation of Percent Spot Price Changes
Daily and Weekly Data

	Daily Data			Weekly Data		
Contract	Positive	Negative	All	Positive	Negative	All
1-Month	0.84*	0.80*	0.83	0.87	0.86	0.87
3-Month	0.72*	0.68*	0.70	0.76	0.68*	0.73
6-Month	0.78*	0.41*	0.61	0.88	0.33*	0.60

Note: Statistics are for when the interest adjusted basis is positive (Positive) and observations when the interest adjusted basis is negative (Negative). An asterisk indicates that an F-test rejects the null hypothesis that the variances of spot and futures prices are equal at the ten percent level.

Simply eyeballing the data in table 4.4 would give us an indication that the theory of storage prediction is correct in this case. In both the daily and weekly data IAB samples, the ratio for the positive IAB sample is always greater than that of the negative IAB sample. However, it is necessary to use an F-test to test if the null hypothesis that the standard deviation of percent price changes for negative IAB samples is the same as the positive IAB samples. Hulleman (1993) explains: If, for positive IAB samples, we reject the null hypothesis, we can conclude that the market does not pass this test (since the variances are not equal, the ratios of the variances could not equal one). Conversely, for negative IAB samples, if we reject the null hypothesis, we can conclude that the market does pass this indirect test (since we require a ratio of less than one for negative IAB samples, rejection indicates that the variances are not equal and therefore could not possibly equal one).

Once again looking at table 4.4. we can see that, in the weekly sampled series, the three and six months contracts pass the hypothesis testing. The daily results show that the positive IAB samples fail the hypothesis testing for all three contracts. Once again though, we must realize that the weekly sampled data holds more weight when it come to analysis.

4.5.3 Time to Maturity Effects on Price Volatility.

The third and final test of the theory of storage is to see if demand shocks produce larger changes in near term futures as opposed to longer term futures. To test this, we once again look at the ratio of the standard deviation of the percent futures price change to the standard deviation of the percent spot price changes. The theory of storage predicts that this ratio should fall as the maturity date increases for both the positive and negative IAB samples. Looking at table 4.4, we can see that this holds for both the negative daily and weekly IAB samples, but does not hold for either of the positive IAB samples as the six month contract ratio is higher than the three month ratio in both daily and weekly samples. Overall (see the column labeled "All"), we see that the natural gas market passes this test of the theory of storage. However, we need both the positive and negative IAB columns to act in an appropriate manner, and thus, natural gas fails this test of the theory of storage.

4.6 Comparisons With Other Energy Commodities.

Since natural gas has been traded for only five years on the NYMEX, it is interesting to see how the tests of the theory of storage conducted here compare to those conducted on other energy products which have been traded over a longer period of time. Serletis and Hulleman (1994) test the theory of storage on crude oil, heating oil and unleaded gasoline.

Their first test on the theory of storage (the test to see if shocks produce more variation in spot than futures prices when inventory levels are low, implying that the IAB is more variable when negative) produces disimilar results to the ones I report for natural gas. For crude oil, heating oil, and unleaded gasoline, an F-test rejected the null hypothesis of equal variances in two of three, three of three, and two of three cases respectively. From the results reported in this thesis, we reject the null hypothesis of equal variances in one of three cases. We must remember that to pass this test of the theory of storage, in addition to rejecting the null hypothesis of equal variances, the standard deviation of changes in the negative IAB sample must be greater than in the positive IAB sample. This is not the case in either of the two crude oil contracts that were rejected by the F-test. However, for heating oil, unleaded gasoline, and natural gas, any contract which rejected the null hypothesis of equal variances also had a larger value for standard deviations of changes in the IAB in the negative IAB sample as compared to the positive IAB sample. Thus we can say that, in general, the results of this test of the theory of storage as applied to natural gas fit the theory of storage less when than testing done on other energy products.

Their second test of the theory of storage (to test if shocks produce roughly equal changes in spot and futures prices when inventory is high, but more variation in spot prices then in futures prices when inventory is low) produces similar results to the ones I produce for natural gas. To pass this test of the theory of storage, we need to reject the null hypothesis of equal variances between the positive and negative IAB samples when looking at the ratio of the standard deviations of percent futures price changes to spot price changes. For crude oil, heating oil, and unleaded gasoline, Serletis and Hulleman (1994) reject the null hypothesis in two of three, three of three, and two of three cases respectively. The results for natural gas I derived here rejected the null hypothesis in two of three cases, so we can conclude that my results are generally consistent with those of other energy products.

Serletis and Hullemans final test of the theory of storage (to see if demand shocks produce larger changes in near term futures as compared to longer term futures) gives quite different results than those I derived for natural gas. To pass this test of the theory of storage, the ratio of the standard deviation of percent futures price changes to spot price changes should fall with increasing maturities for both positive and negative IAB samples. Each of crude oil, heating oil and unleaded gasoline pass this test. However, for the natural gas positive IAB sample, this does not hold. Thus, the natural gas market does not pass this theory of storage.

4.7. Conclusion.

The theory of storage was developed to explain why various commodities were being stored even when the return on storage was negative. Those doing work on the subject explained this phenomenon as a result of the commodity providing "convenience yield" to the party storing it. While the theory was originally developed to be applied to agricultural commodities subject to a harvest, there is no reason that it could not be applied to energy products which also exhibit seasonal variations in supply.

With this in mind, I applied the various tests of the theory of storage and the predictions it entails to natural gas, following Fama and French (1988). The test to see if shocks provided more variation in spot than in futures prices when inventory levels are low produced results which rejected this prediction. When testing to see if shocks produce equal changes in spot and futures prices when inventory levels are high but more variation in spot than futures prices when inventory levels are low, I found evidence supporting this prediction. Finally, the test to see if shocks produce larger changes in shorter maturity futures as compared to longer maturity futures produced results which did not support this prediction. Thus, of the three tests of the theory of storage, the natural gas market passed only one of them.

After completing the testing, an interesting application is to see how the results compare to testing done on other energy products. To achieve this end, I compared my results to that of Serletis and Hulleman (1994) who engaged in the same testing of the theory of storage discussed within this chapter but applied to crude oil, heating oil and unleaded gasoline. When doing so, I found my results to be less than consistent with theirs on two of three tests of the theory of storage. Overall, the testing results reported in this chapter are less than close to that reported by Serletis and Hulleman. After natural gas has traded for a longer period, giving us more observations, it would be interesting to see if the results for the natural gas market more closely match those of other energy markets. Also, due to the changing nature of regulation in the industry another anlysis in a few years time may be warranted.

Figure 4.3. Natural Gas Weekly Spot, 1-Month, 3-Month, and 6-Month Futures Prices: 04/03/90 - 05/17/95

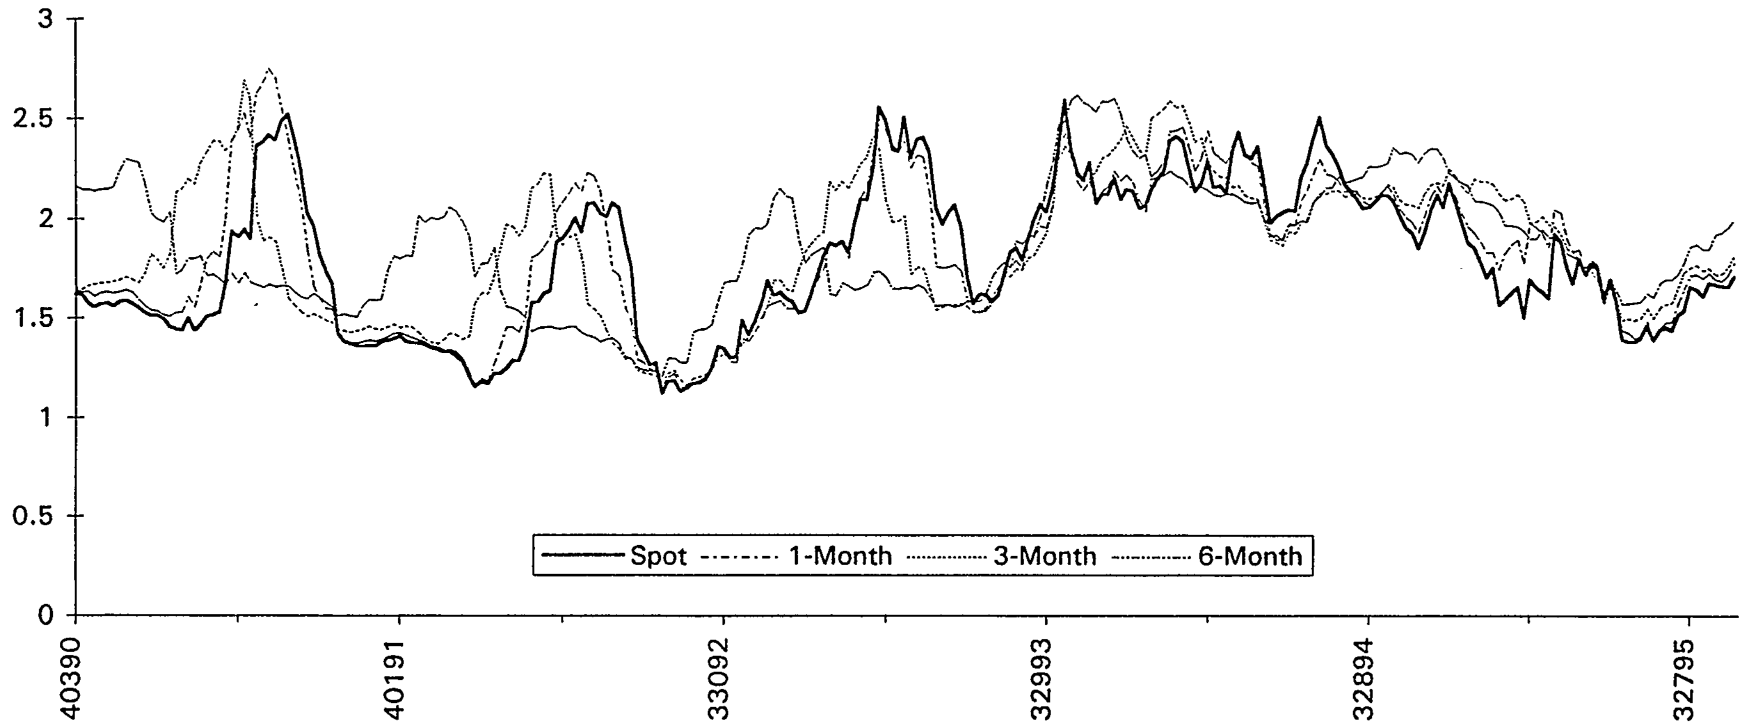

Figure 4.4. 1-Month Natural Gas Interest Adjusted Basis

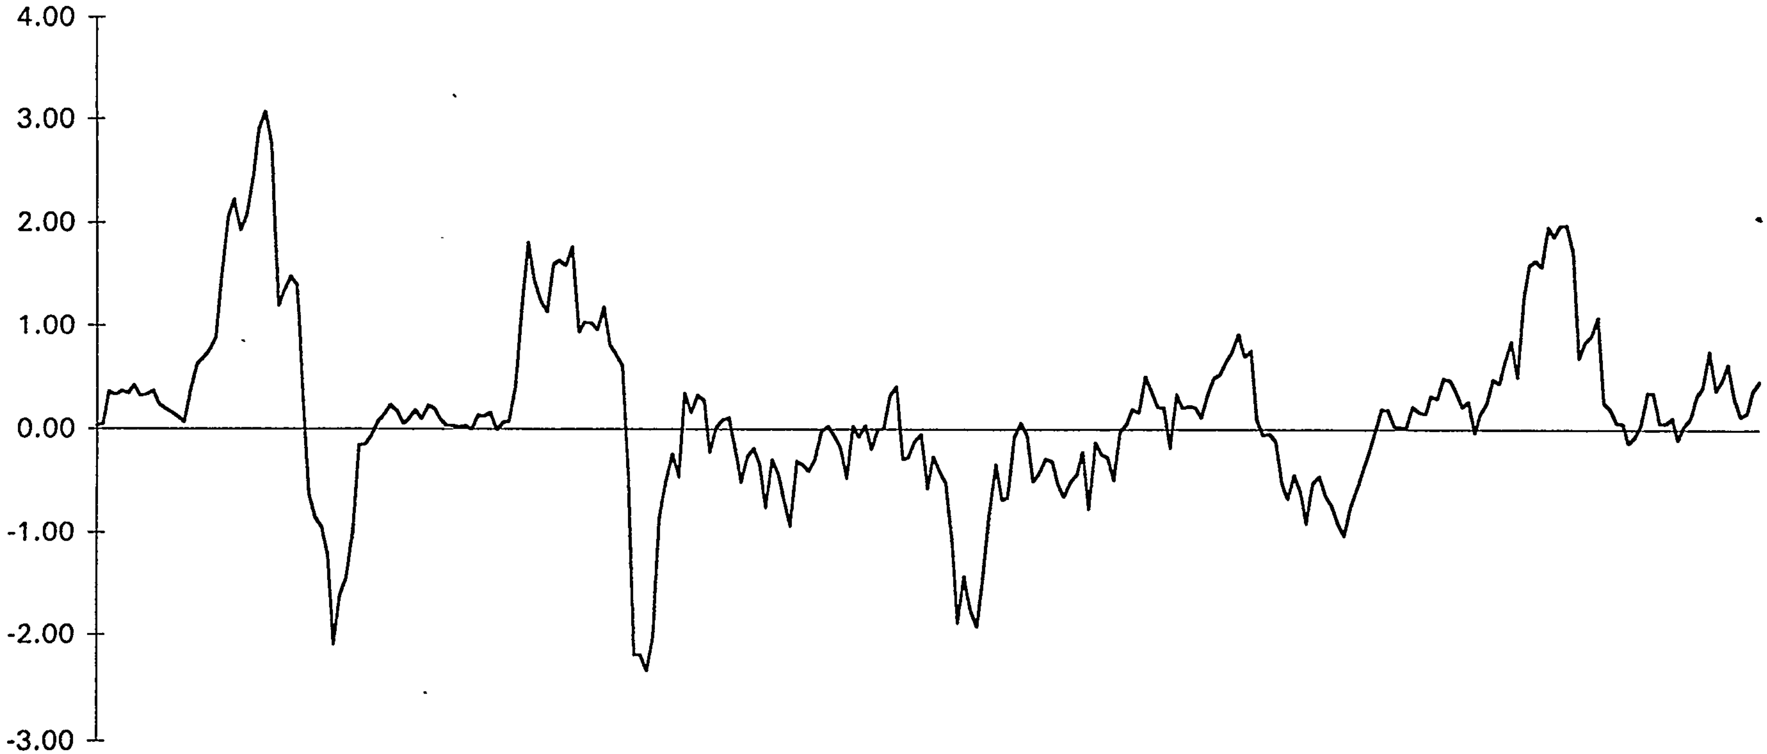

Figure 4.5. 3-Month Natural Gas Interest Adjusted Basis

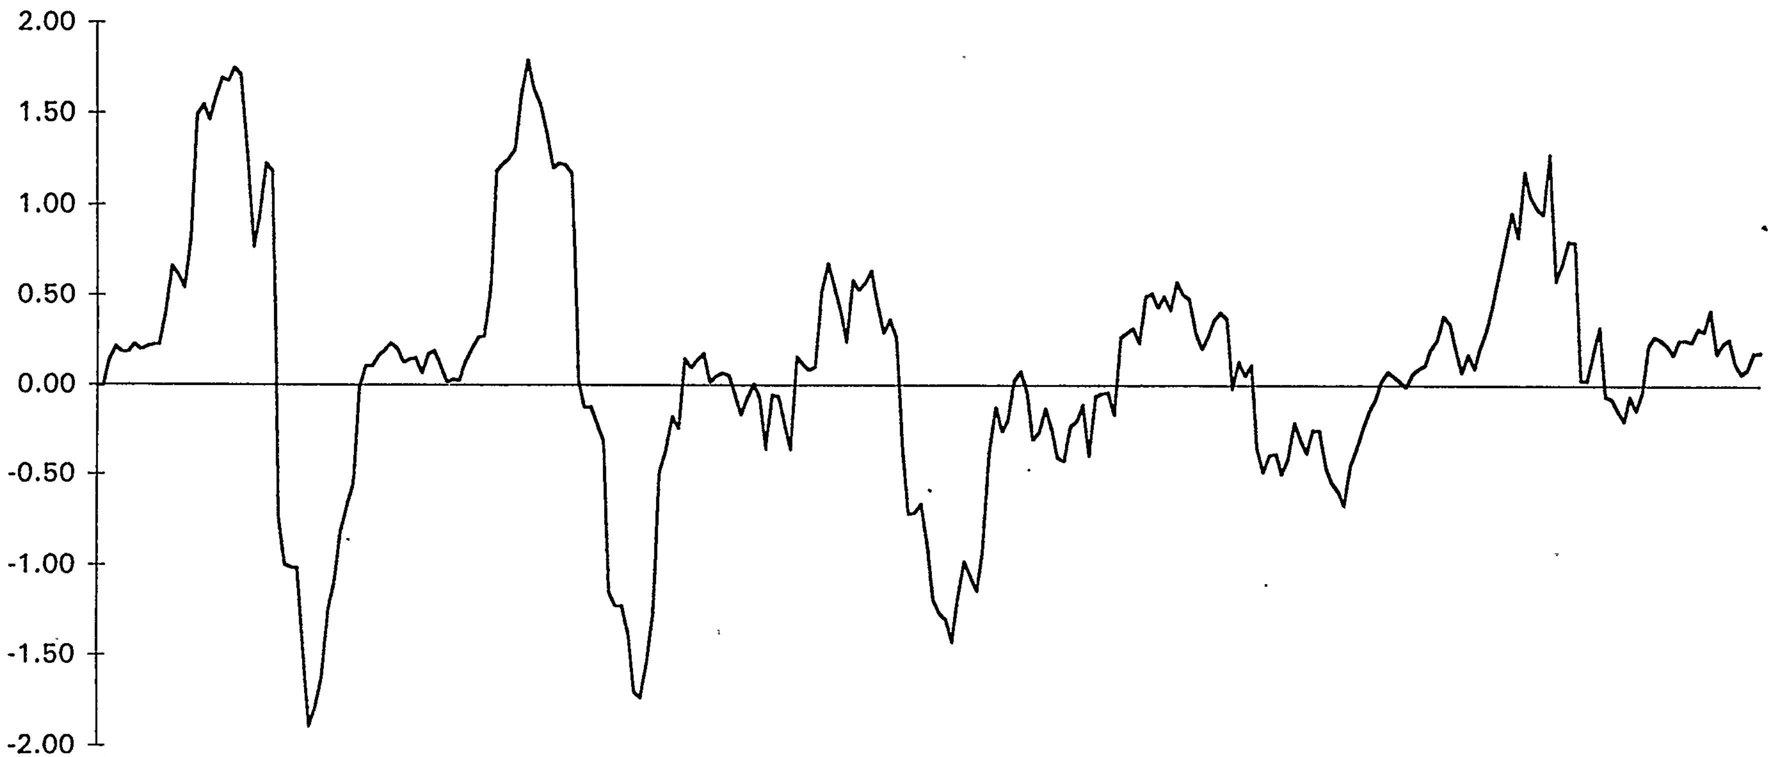

Figure 4.6. 6-Month Natural Gas Interest Adjusted Basis

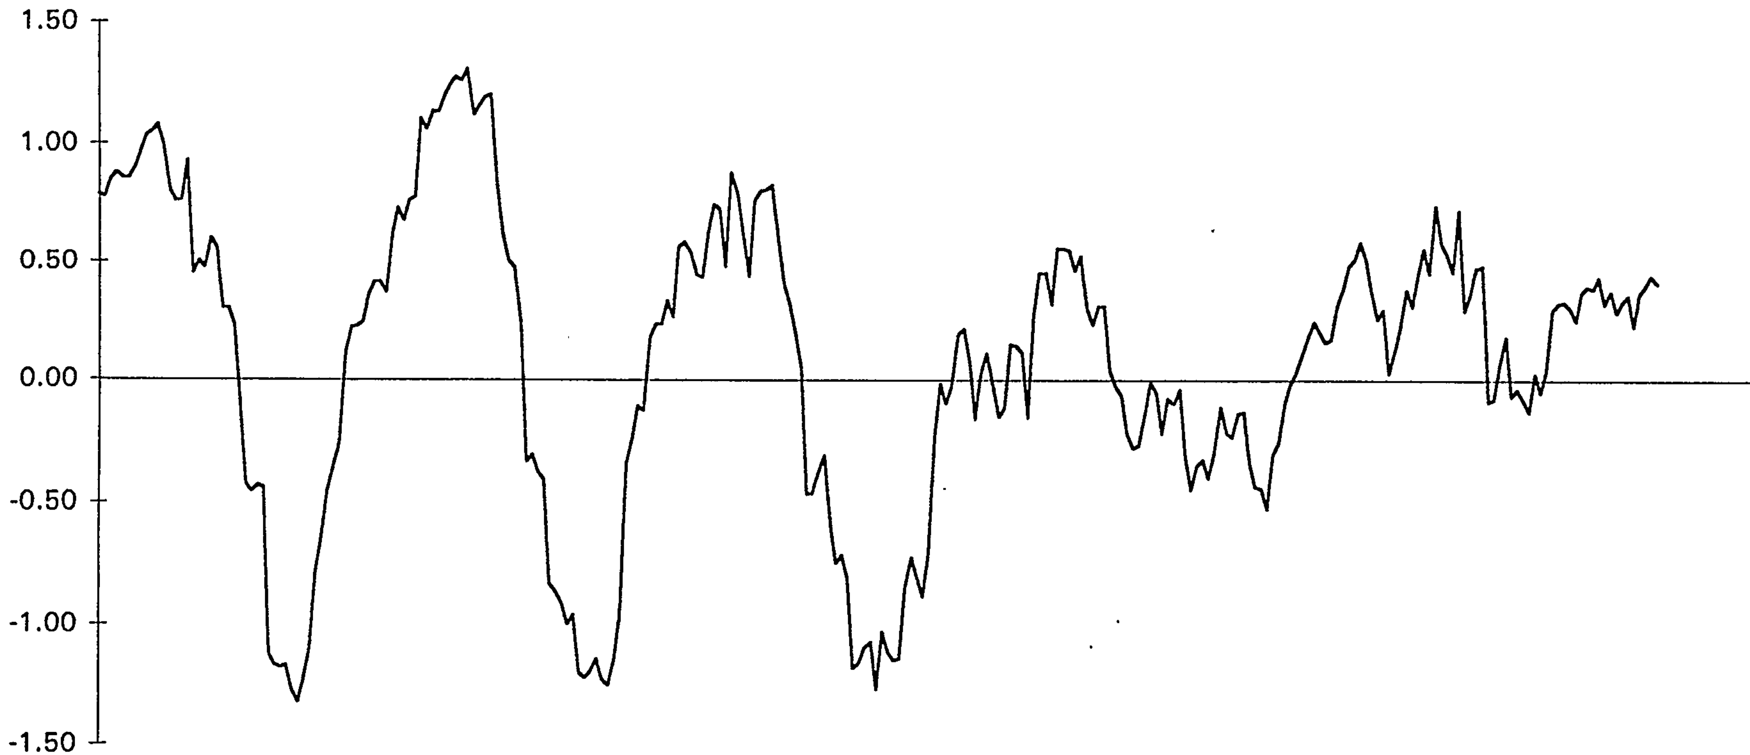

Chapter 5

Tests for Market Efficiency

5.1 Introduction.

As discussed earlier in this thesis, it is generally accepted that futures markets provide two basic functions. The first is to transfer risk from those who wish to shed it (namely hedgers) to those who wish to accept it (speculators). Almost without question, this basic function of futures markets is satisfied. On a rare occasion, one may find difficulty closing a position in the market, but this is quite unlikely to happen. Secondly, it is said that futures markets can be used as a forecasting tool. Specifically, the current futures price should be an unbiased predictor of the future spot price if the market is functioning efficiently. However, one may witness that this is not always the case, and thus, the legitimacy of this function of futures markets is debatable. The goal of this chapter is to test to see if natural gas markets do behave in an efficient manner and if, indeed, the current futures price can reliably predict the future spot price.

To engage in such a task, I will use the methodology of Dickey and Fuller (1979), which will enable me to determine if the variables we are looking at are stationary or not. If they are indeed stationary, I will follow Fama's (1984) approach to determine if current futures prices can predict the future spot price. If the variables are non stationary, I will use the cointegration methodology of Engle and Granger (1987) to determine if some sort of long run relationship exists between the variables in question.

5.2 Non Stationarity and Unit Roots.

From the literature on the efficient market hypothesis (see, for example, Serletis (1992)), we learn that an efficient market is a market in which price changes are uncorrelated. In other words, we question whether price changes fully reflect economic information. If price changes are uncorrelated, this implies that a unit root is present in the level of that particular price series (or the natural log of the price series).

The presence of a unit root in a time series simply implies that the series in question follows a stochastic trend. Conversely, if there is not a unit root present in the time series, the series is said to be "stationary", meaning linear properties exist and are time invariant (see Granger 1986).

Typically, a series with a unit root is said to be difference stationary, meaning that the series needs to be differenced once to become stationary. A series such as this is denoted as I(1). In general, a series needing to be differenced 'd' times to become stationary is denoted I(d). Thus, a stationary series is denoted as an I(0) series.

The differences between an I(0) series and an I(1) series are stark. An I(0) series is characterized as a series which possesses a mean and will tend to fluctuate around its mean. Any deviation from the mean will be relatively short as the series is "drawn back" to its mean value. Autocorrelations in such a series will diminish rapidly as one increases the lag length, thus we can say that an I(0) series has a finite memory. Conversely, an I(1) series will tend to follow a somewhat smooth pattern, covering a large range of values, rarely returning to an earlier value. Autocorrelations in a series such as this are usually close to the value of one, even for a large number of lags. Thus, we can say that an I(1) series is characterized as having an infinite memory.

If two series in question are I(1) (ie., a unit root is present), then it is possible that the two series may be cointegrated. This simply implies that there is some sort of long run relationship between the two. If the two series are I(0), the two series cannot be considered cointegrated, but standard testing methodologies, such as ordinary least squares, can be incorporated to see if there is any relationship between the two series.

To test for the presence of a unit root , we utilize the Dickey-Fuller (DF) and the Augmented Dickey-Fuller (ADF) testing procedure. The ADF equation is shown below, in equation (1).

$$\Delta Z_t = \alpha + \beta t + \rho Z_{t-1} + \sum_{i=1}^{r} \beta_i \, \Delta Z_{t-i} + \epsilon_t \tag{1}$$

where Z_t is the time series being considered and r is selected to take in a value which is large enough to ensure that ϵ_t is white noise. To choose r, we follow Said and Dickey (1984) who show that the ADF test is asymptotically valid if r is increased with sample size (N) at a controlled rate ($N^{1/3}$). For my sample size of sixty observations, this implies that r should take on a value of four. If we select r to have a value of zero, the ADF equation collapses to that of the DF equation.

We run ordinary least squares to test the null hypothesis: Ho: $Z_t \sim I(1)$. This null hypothesis is rejected if ρ is negative and significantly different from zero. In this case, the test statistic is not distributed as t; however, Dickey and Fuller (1979) have provided tables of significance. Incorporated in equation (1) is a time trend (t), which can be included or excluded in the ordinary least squares calculation, depending on which specification is most valid for the data at hand. For the purposes of this thesis, I will both include and exclude the time trend to see how robust the results are.

By testing for a unit root in a series, we are simply trying to determine if the series is stationary or not. If the two series in question are non stationary, we can use the theory of cointegration to see if a relationship exists between the two. If the two series are stationary, we can use a standard infrencing procedure such as ordinary least squares to see what relationship, if any, holds between the two respective series.

My hypothesis is that the approach taken hereafter will yield series' that are stationary. With this in mind, I will now move on to discuss the methodology introduced by Fama (1984), to determine the relationship between spot and futures prices. To have any value, his methodology requires stationarity of the series in question.

5.3. Theoretical Foundations.

For the presentation purposes, let the futures price at time t for delivery at T be F(t,T) and let the spot price at time t be S(t). If we then assume that the market will act in such a way that the current futures price, F(t,T), will be the certainty equivalent of the future spot price, S(T), we can divide the certainty equivalent into two separate parts. The first part signifies a premium, the second being an expected future spot price. In the form of an equation, this will look like:

$$F(t,T)=P(t)+E\{S(T)\} \tag{2}$$

In this case, E{S(T)} is a rational forecast which will depend upon all the information available at time t. P(t) is the bias of the futures price, F(t,T), as the forecast of the future spot price S(T).

By rearranging equation (2), we can formulate an expression which will allow us to test some hypotheses about the basis which is commonly defined as the difference between the current futures price and the spot price. If we subtract the current spot price, S(t), from both sides of equation (2), we arrive at:

$$F(t,T)-S(t)=P(t)+E\{S(T)\}-S(t) \tag{3}$$

Looking at equation (3), one can see that it represents the basis as being split between a premium component, P(t), and an expected change in the spot price component, E{S(T)-S(t)}.

Following Fama (1984), we can investigate various hypotheses related to the variability of risk premiums and expected spot price changes. Specifically, Fama (1984) suggests that the following two regressions be considered:

$$F(t,T)-S(T)=\alpha_1+\beta_1[F(t,T)-S(t)]+u(t,T) \tag{4}$$

$$S(T)-S(t)=\alpha_2+\beta_2[F(t,T)-S(t)]+\epsilon(t,T) \tag{5}$$

By estimating equation (5), we will be able to tell whether the current futures-spot price differential has the power to predict the future change in the spot rate. If we find evidence that β_2 is reliably non zero, we can take this as meaning that the futures price observed at time t contains information about the spot price observed at time T. Similarly, estimation of equation (4) and the corresponding value of β_1 will give us information about whether the premium component of F(t,T)-S(t) exhibits variation that shows up reliably in F(t,T)-S(T).

Upon further inspection of equations (4) and (5), one can see that the equations are dependent. This is due to the fact that the regressor in both equations is the same. Further, one can also see that the sum of the dependent variables also equals the regressor. With this fact in mind, this implies that $\alpha_1 = -\alpha_2$, $\beta_1 = 1 - \beta_2$ and $u(t,T) = -\epsilon(t,T)$. Realizing this, one can see that there is no need to run both regressions, since both contain the same information. However, I will run both regressions to see if the relationships above do hold.

Both of the regressions allocate the variation in the basis to variation in premiums, expected spot price changes, and some mix of the two. However, Serletis (1990) explains that, "the allocation may be statistically unreliable when the premium and the expected change in the spot price components of the basis are correlated." To better explain the situation, Fama (1984) shows that under the assumption that the expected future spot rate is efficient, the appropriate regression coefficients are correctly stated as:

$$\beta_1 = \frac{COV[F(t,T)-S(T), F(t,T)-S(t)]}{VAR[F(t,T)-S(t)]} \quad (6)$$

$$\beta_2 = \frac{COV[S(T)-S(t), F(t,T)-S(t)]}{VAR[F(t,T)-S(t)]} \quad (7)$$

where COV(.,.) and VAR(.) stand for the covariance and variance respectively. Substituting equation (3) into equations (6) and (7), we arrive at the following expressions:

$$(8) \qquad \beta_1 = \frac{VAR[P(t)]+COV[P(t),E\{S(T)-S(t)\}]}{VAR[P(t)]+VAR[E\{S(T)-S(t)\}]+2COV[P(t),E\{S(T)-S(t)\}]}$$

$$(9) \qquad \beta_2 = \frac{VAR[E\{S(T)-S(t)\}]+COV[P(t),E\{S(T)-S(t)\}]}{VAR[P(t)]+VAR[E\{S(T)-S(t)\}]+2COV[P(t),E\{S(T)-S(t)\}]}$$

The reason we engage in this transformation is to show that if the premium, P(t), is constant over time, β_1 and β_2 will be equal to zero and one respectively. From this, we can see that the two coefficients will roughly describe the degree of variability in the components that make up the basis. If we are faced with a situation where the premium and the expected change in the spot price are uncorrelated, then our β_1 coefficient would be equal to the proportion of the variance on the basis due to variation in the risk premium. Along similar lines, β_2 would be equal to the proportion of the variance of the basis due to the variance of the expected change in the spot price. However, we must realize that it is quite unlikely that the components of the basis are uncorrelated, and the covariance terms must be incorporated.

5.4 Data and Results.

The data used here were daily natural gas futures prices, beginning on April 3, 1990 and ending on May 17, 1995. Once again, since cash prices are difficult to obtain, the spot months futures price is used as a proxy for the current cash price and the second months futures price is used as the current futures price.

Since we have assumed that the futures prices converges to the future spot price on the settlement date of the contract, I matched the current spot and futures price with the spot price on the settlement date of the contract. Doing so gave me a sample containing 60 observations.

Table 5.1
Summary Statistics for the Basis, the Premium and Spot Price Changes

	Observations	Mean	Standard Deviation
F(t,T)-S(t)	60	0.0086	0.0804
F(t,T)-S(T)	60	0.0154	0.1657
S(T)-S(t)	60	-0.0068	0.1677

Referring to Table 5.1, we can see the mean and the standard deviations of the basis (F(t,T)-S(t)), the premium (F(t,T)-S(T)), and the change in the spot price (S(T)-S(t)). From the results, we see that the basis variation, as measured by the standard deviation, is low relative to the variations of the premiums and spot price changes. This is consistent with Serletis (1991). He interprets this finding as meaning it is "unlikely that the regressions (equations (4) and (5)) will reliably assign basis variation to premiums and expected spot price changes."

With this in mind, we can now move on to the unit root testing that was discussed earlier. Reported in Table 5.2, are the "t" values for ρ obtained by running the Simple Dickey Fuller (DF) and the Augmented Dickey Fuller (ADF) (see equation (1)) tests with both a constant and a time trend incorporated. Reported in Table 5.3 are test results incorporating the constant term, but not a time trend.

Table 5.2
Test Results for Unit Roots Incorporating a Constant and a Time Trend

	DF	ADF
F(t,T)-S(t)	-9.11	-4.45
F(t,T)-S(T)	-9.56	-5.14
S(T)-S(t)	-10.22	-6.01

Table 5.3
Test Results for Unit Roots Incorporating a Constant

	DF	ADF
F(t,T)-S(t)	-9.19	-4.40
F(t,T)-S(T)	-9.60	-5.19
S(T)-S(t)	-10.28	-6.09

To reject the null hypothesis of a unit root (H_0: $Z_t \sim I(1)$), the "t" value of ρ must be negative and significantly different from zero. The critical values of the DF and ADF tests at the 1%, 5% and 10% levels are -2.62, -1.95, and -1.61 respectively. One can see that the values of ρ reported in both Table 5.2 and Table 5.3 are all negative and all significant different from zero. Thus, we can say that all of the variables are stationary, or I(0), as expected. Since the series' are all stationary, taking the approach described by Fama (1984) is appropriate. We can now utilize the regression analysis described earlier to see what sort of predictive power futures prices hold.

In Table 5.4, I report the regression results of the premium F(t,T)-S(T) and the change in spot price S(T)-S(t) on the basis F(t,T)-S(t). In other words, these are equations (4) and (5) revisited:

(4) $$F(t,T)-S(T)=\alpha_1+\beta_1[F(t,T)-S(t)]+u(t,T)$$

(5) $$S(T)-S(t)=\alpha_2+\beta_2[F(t,T)-S(t)]+\epsilon(t,T)$$

Table 5.4
Regression Results

α_1	β_1	α_2	β_2	$S(\alpha)$	$S(\beta)$	R_1^2	R_2^2	DW
0.0116	0.4479	-0.0116	0.5521	0.0212	0.2641	0.047	0.070	1.5468

As was discussed before, the sum of the β coefficients should be and is one and the sum of the intercepts is zero since the regressor in both equations is the same. Note also that in both equations, the R^2 statistics are very small. R_1^2 denotes the coefficient of determination for the premium regression (equation (4)) and R_2^2 denotes the coefficient of determination for the changes in spot price regression (equation (5)). In both cases, the R^2 statistic is low, which was expected due to the fact that the basis has low variation as compared to both the premium and the change in spot price (as was reported in Table 5.1). Further, due to the complementary nature of the two regressions, only one set of coefficients for the standard errors is reported. However, this is sufficient since these coefficients are the same for both regressions.

From this, we can also see that the hypothesis that the β_2 coefficient equals one is rejected (or, equivalently, that the β_1 coefficient equals zero is rejected). This implies that the premium varies over time and is not constant.

Since both my β_1 and β_2 coefficients are positive and reliably non zero, we can also say that the futures price has reliable power to forecast spot prices and the futures price contains a time varying premium that shows up reliably in F(t,T)-S(T).

Thus, the evidence from these tests supports the hypothesis that natural gas markets are efficient and that the current futures price has reliable power to predict the future spot price.

5.5 Conclusion.

The goal of this chapter was to test for efficiency in natural gas markets. When speaking about testing for an efficient market, I mean that we are testing to see if the market's current futures price can reliably predict the future spot price. To engage in such a task requires several steps.

After creating a sample by matching the current spot and futures price with the spot price on the settlement date of the contract, I tested the simple statistical properties of the market. Doing so showed that basis variation was low relative to variation in premiums and spot price changes which indicated that it was unlikely that basis variation was related to premium and spot price changes.

Next, using the methodology of Dickey and Fuller (1979), I tested for the presence of unit roots in the basis, the premium, and spot price changes. Doing so, tells us if our series' are stationary and if we can utilize "normal" testing procedures and will not have to alter our series. Evidence from this testing showed that all three series were indeed stationary.

Since the series were deemed stationary, standard regressions were run to see if the current futures price had any power to predict the future spot price. After running these regressions, the results showed evidence of a time varying risk premium, and also, that the current futures price did indeed have the power to predict the future spot price on a reliable basis. Therefore, we can conclude that the natural gas market did behave in an efficient manner over this particular time period.

Chapter 6

Conclusion

This thesis originated with the goal of examining several different properties of natural gas prices and the market in which natural gas operates. The reason that this is interesting is that the natural gas industry is experiencing widespread growth, and it is of some interest to obtain better knowledge of how this market behaves.

I began my analysis with an introduction, for the uninitiated, of the origin and operation of futures markets. This discussion encompassed the basics of futures markets, the criteria needed for a successful futures market, the specifics of the natural gas futures contract and how trades occur. Evolving from this discussion was the conclusion that natural gas met all the criteria needed for a successful futures market and showed its necessity for hedging and speculative purposes.

Following this, chapter 3 examined the cyclical properties of natural gas prices. This was done to get an idea of how gas prices tended to move with other variables and what the implications might be for hedgers, speculators, and the economy in general. Making use of Prescott's (1986) methodology, I employed the Hodrick-Prescott (HP) filter to detrend several different time series in order to observe their cyclical component. Specifically, I detrended a monthly natural gas price series and compared its cyclical component with that of a monthly United States production index, price index, and unemployment rate. Here, it was found that natural gas prices and output were related in a weak procyclical fashion, natural gas and unemployment were also related in a weak procyclical fashion, and finally, it was established that natural gas prices and other prices had no discernible relationship with each other.

To test the robustness of these results, I removed the trends of the series in question using first differencing as opposed to using the HP filter. Using this method provided results that natural gas prices were acyclical with all of output, unemployment and prices. Since two different methods of trend elimination gave quite different results, the reliability of the results are quite questionable.

Next, running a similar exercise, I compared crude oil prices with natural gas, unleaded gasoline, and heating oil. This was done in order to see how energy prices tended to move with respect to crude oil prices, the most widely used energy product. Once again, I used both the HP filter and first differencing to test the robustness of the results.

The results showed that crude oil and heating oil, and crude oil and unleaded gasoline moved in a strong procyclical fashion. Crude oil and natural gas moved in a weak procyclical fashion. The results taken together indicate that energy prices in general tend to move in the same direction. These results were consistent using both methods of trend elimination, and thus, we can be quite comfortable with the reliability of these results.

Following this cyclical analysis of natural gas prices, chapter 4 turned to examine the theory of storage and the various predictions it encompasses. The works of Working (1949), Brennan (1958), and Tesler (1958) introduce us to the concept of a convenience yield which helps to explain why commodities were being stored when the price of (or return on) storage was negative. Their work led to three predictions, all of which I tested with respect to natural gas. Utilizing weekly data to smooth erratic price movements and constructing a variable called the Interest Adjusted Basis (IAB) which was used to approximate inventory levels, enabled me to test the various predictions.

First, the authors predicted that supply and demand shocks cause more variability in spot and futures prices when inventory levels are low (negative IAB) as opposed to when inventory levels are high (positive IAB). This implies that the IAB is more variable when it is negative. This was found to be true in one of three cases, in the three month IAB.

The second test of the theory of storage tests the prediction that supply and demand shocks cause approximately equal changes in spot and futures prices when the inventory levels are high, but cause spot prices to change more than futures prices when inventory levels are low. To accomplish this, I compared the standard deviation of percent futures price changes to that of spot price changes. The prediction is that the ratio should be close to one when the IAB is positive and somewhat less than one when the IAB is negative. This was found to hold in two of three cases, in the three and six month IAB samples.

The third test of the theory of storage tests the prediction that demand shocks produce larger changes in nearer term futures than in longer term futures. I once again look at the ratios of the standard deviations discussed above. These ratios should fall as the maturity date increases for both the positive and negative IAB samples. Here, it was found that the negative IAB sample passes this test, while the positive IAB sample does not, giving mixed results.

Overall, it seems that natural gas acts in a manner that is inconsistent with the predictions that the theory of storage supplies, and the results are generally inconsistent with work done on other energy commodities such as that done by Serletis and Hulleman (1994).

Chapter five looks at the natural gas market with goal of finding whether or not it has operated efficiently over its history. To be considered an efficient market, the market's current futures price should have the power to predict the future spot price. Using the methodology of Dickey and Fuller (1979), I tested for unit roots in the basis, the premium, and spot price changes. Finding that the series were stationary, regressions were run to see if indeed the current futures price had the power to predict the future spot price. These regressions gave evidence of a time varying risk premium and that futures prices did in fact have the power to predict future spot prices. Thus, it can be concluded that the natural gas market has been efficient in its operation.

In conclusion, I have investigated various properties of natural gas markets. Specifically, I looked at the form and function of futures markets, the cyclical properties of natural gas prices, the theory of storage, and market efficiency. Hopefully, this has shed new light on a growing market and will be of use to those involved in natural gas in any fashion.

Bibliography

Banack, D. (1990) *The Message in Petroleum Futures Prices,* M A (Econ) thesis, University of Calgary. © David Banack.

Baxter, M., and R.G. King (1995). "Approximate Band-Pass Filters for Economic Time Series." NBER Working Paper No. 5052.

Brennan, M.J. (1958). "The Supply of Storage." *American Economic Review* 48:50-72.

Chang, E. and C. King (1988). "Day of the Week Effects and Commodity Price Changes." *The Journal of Futures Markets* 4,2: 513-521.

Cho, D.W., and G.S. McDougall (1990). "The Supply of Storage in Energy Futures Markets." *The Journal of Futures Markets* 10: 611-21.

Cogley, T., and J.M. Nason (1995). "Effects of the Hodrick-Prescott Filter on Trend and Difference Stationary Time Series" *Journal of Economic Dynamics and Control* 19:253-278.

Daigler, R.T. (1993). *Managing Risk with Financial Futures.* Chicago: Probus Publishing Company.

Dickey, D.A., and W.A. Fuller (1979). "Distribution of the Estimators for Autoregressive Time Series with a Unit Root." *Journal of the American Statistical Association* 74:427-431.

Engle, R.F., and C.W. Granger (1987). "Cointegration and Error Correction: Representation, Estimation, and Testing." *Econometrica* 55:251-276.

Fama, E. F. (1984). "Forward and Spot Exchange Rates." *Journal of Monetary Economics* 14:319-388.

Fama, E.F., and K. R. French (1988). "Business Cycles and the Behaviour of Metals Prices." *The Journal of Finance* 43:1075-93.

Fiorito, R. and T. Kollintzas (1994). "Stylized Facts of Business Cycles in the G7 from a Real Business Cycles Perspective." *European Economic Review* 38:235-269.

Granger, C.W. (1986). "Developments in the Study of Cointegrated Economic Variables." *Oxford Bulletin of Economics and Statistics* 48:213-28.

Hulleman, V.W. (1993) *The Theory of Storage and Petroleum Futures,* M A (Econ) thesis, University of Calgary. © Vaughn W. Hulleman.

King, R.G., and S. T. Rebelo (1993). "Low Frequency Filtering and Real Business Cycles." *Journal of Economic Dynamics and Control* 17:207-231.

Kydland, F.E., and E.C. Prescott (1990). "Business Cycles: Real Facts and a Monetary Myth." Federal Reserve Bank of Minneapolis *Quarterly Review* 26:3-18.

Prescott, E.C. (1986). "Theory Ahead of Business Cycle Measurement." Federal Reserve Bank of Minneapolis *Quarterly Review* 10:9-22.

Said, S.E., and D. A. Dickey (1984). "Testing for Unit Roots in Autoregressive-Moving Average Models of Unknown Order." *Biometrica* 71:599-607.

Samuelson, P.A. (1965). "Proof that Properly Anticipated Prices Fluctuate Randomly." *Industrial Management Review* 6:41-49.

Serletis, A. (1991). "Rational Expectations, Risk and Efficiency in Energy Futures Markets." *Energy Economics* 4:111-115.

Serletis, A. (1992). "Unit Root Behaviour in Energy Futures." *The Energy Journal* 13:119-128.

Serletis, A. (1994). "A Cointegration Analysis of Energy Futures Prices." *Energy Economics* 16:93-97.

Serletis, A., and D. Banack (1990). "Market Efficiency and Cointegration: An Application to Petroleum Markets." *The Review of Futures Markets* 9:372-380.

Serletis, A., and V. Hulleman (1994). "Business Cycles and the Behaviour of Energy Prices." *The Energy Journal* 15:125-134.

Tesler, L.G. (1958). "Futures Trading and the Storage of Cotton and Wheat." *Journal of Public Economy* 66:233-55.

Treat, J.E. (1990). *Energy Futures: Trading Opportunities for the 1990's.* Oklahoma: Penwell Publishing Company.

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