THE SIMPLE ECONOMICS OF COMMODITY PRICE SPECULATION

Christopher R. Knittel and Robert S. Pindyck*

This draft: March 12, 2013

Abstract

The price of crude oil in the U.S. had never exceeded $40 per barrel until mid-2004. By 2006 it reached $70 per barrel, and in July 2008 it reached a peak of $145. By the end of 2008 it had plummeted to about $30 before increasing again, reaching about $110 in 2011. Are "speculators" to blame for at least part of the volatility and sharp run-ups in price? We clarify the potential and actual effects of speculators, and investors in general, on commodity prices. We focus on crude oil, but our approach can be applied to other commodities. We first address the question of what is meant by "oil price speculation," and how it relates to investments in oil reserves, oil inventories, or oil price derivatives (such as futures contracts). Next we outline the ways in which one could speculate on oil prices. Finally, we turn to the data, and calculate counterfactual prices that would have occurred from 1999 to 2012 in the absence of speculation. Our framework is based on a simple and transparent model of supply and demand in the cash and storage markets for a commodity. It lets us determine whether speculation as the driver of price changes is consistent with the data on production, consumption, inventory changes, and changes in convenience yields given reasonable elasticity assumptions. We show speculation had little, if any, effect on prices and volatility.

JEL Classification Numbers: Q40; L71; G13

Keywords: Commodity prices, crude oil, speculation, futures markets, storage markets, inventories.

*Knittel: William Barton Rogers Professor of Energy Economics, Sloan School of Management, MIT and NBER, knittel@mit.edu. Pindyck: Bank of Tokyo-Mitsubishi Professor of Economics and Finance, Sloan School of Management, MIT and NBER, rpindyck@mit.edu.

Commodities have become an investment class: declines in their prices may simply reflect the whims of speculators.

The Economist, June 23, 2012.

Tens of billions of dollars went into the nation's energy commodity markets in the past few years, earmarked to buy oil futures contracts. Institutional and hedge funds are investing increasingly in oil, which has prompted President Obama and others to call for curbs on oil speculation.

The New York Times, September 4, 2012.

Federal legislation should bar pure oil speculators entirely from commodity exchanges in the United States.

Joseph Kennedy II, *New York Times*, April, 10, 2012.

1 Introduction.

The price of crude oil in the U.S. had never exceeded $40 per barrel until mid-2004. By 2006 it reached $70 per barrel, and then in July 2008 it reached a peak of over $140. As shown in Figure 1, by the beginning of 2009 it had plummeted to about $38 and then started increasing again, reaching about $110 in 2011. What caused these sharp changes in oil prices since 2004? Were they due to fundamental shifts in supply and demand, or are "speculators" to blame for at least part of the volatility and sharp run-ups in price? This question is important: the wide-spread claim that speculators have caused price increases has been the basis for attempts to limit—or even shut down—trading in oil futures and other energy-based derivatives.

Other commodities have also experienced large price swings, as shown in Figure 2. Note that on several occasions during the past decade the prices of industrial metals such as copper, aluminum, and zinc have more than doubled over periods of just a few months, often followed by sharp declines. Also, as the figure suggests, price changes across commodities tend to be correlated. For example, the correlation coefficients for crude oil and aluminum, copper, gold, and tin range from .74 to .89 in levels, and .52 to .71 in monthly first differences. Should we infer from the volatility of commodity prices (and there correlations) that commodities have indeed "become an investment class?"

The claim that speculators are to blame and futures trading should be limited is well exemplified by a recent Op-Ed piece in the *New York Times* by Joseph P. Kennedy II. He wrote that “the drastic rise in the price of oil and gasoline” is at least partly attributable to “the effect of pure speculators—investors who buy and sell oil futures but never take physical possession of actual barrels of oil.” He argues that “Federal legislation should bar pure oil speculators entirely from commodity exchanges in the United States. And the United States should use its clout to get European and Asian markets to follow its lead, chasing oil speculators from the world’s commodity markets.”

Was Kennedy on to something? Unfortunately there is a great deal of confusion—even among economists—over commodity price speculation and how it works. In fact, even identifying “speculators,” as opposed to long-term investors or firms hedging risk, is not so simple. Claiming, as Kennedy did, that anyone who buys or sells futures but does not take possession of the physical commodity is a “pure speculator” is nonsensical. Hardly any person, firm, or other entity that buys or sells commodity futures contracts ever takes possession of the commodity, and we know that a substantial fraction of futures contracts are held by producers and industrial consumers (e.g., refineries) to hedge risk.

This paper attempts to clarify the potential and actual effects of speculators, and investors in general, on the prices of storable commodities. We will focus on the price of crude oil because it has received the most attention as the subject of speculation. More than other commodities, sharp increases in oil prices are often blamed, at least in part, on speculators. (Interestingly, however, speculators are rarely blamed for sharp *decreases* in oil prices.) But our theoretical and empirical approach can be applied equally well to other commodities.

We begin by addressing the question of what is meant by “oil-price speculation,” and how it relates to investment in oil reserves, oil inventories, or oil-price derivatives (such as futures contracts). Given that oil-price speculation is just an investment designed to pay off if the price of oil goes up (or alternatively, if it goes down), we outline the ways in which one could actually engage in speculation. What kind of instruments are available for speculation, and how costly and effective are they? Most importantly, how can these various forms of price speculation affect the current price of oil? We try to clarify the mechanisms by which speculators (and investors) affect oil prices, production, and inventories, and thereby provide a “simple” explanation of the economics of price speculation. Finally, we turn to the data, and answer the questions raised in the first paragraph: What role did speculation have in the sharp changes in oil prices that have occurred since 2004?

Other researchers have also investigated the causes of oil price changes and the possible role of speculation. Fattouh et al. (forthcoming) summarize the literature. We briefly discuss the most relevant papers here. Kilian and Murphy (forthcoming) note (as we do) a connection between speculative activity and inventory changes, and estimate a vector-autoregressive (VAR) model that includes inventory data to identify the "asset price component" of the real price of oil. They find no evidence that speculation increased prices. But Juvenal and Petrella (2011) estimate an alternative VAR model, also using inventory data, and conclude that "speculation played a significant role in the oil price increase between 2004 and 2008 and its subsequent collapse." Hamilton (2009a,b) provide an overview of possible causes of oil price changes and concludes that speculation played some role in the price increase in the summer of 2008. Smith (2009) does not find any evidence that speculation increased prices between 2004 and 2008 noting that inventories were drawn down during this time and there was no evidence that non-OPEC producers reduced output.[1] Below, we also discuss whether production or drilling activity subsided during this time.

Our framework is based on a simple and transparent model of supply and demand in the cash and storage markets for a commodity. Using that model, we can determine whether speculation as the driver of price changes is consistent with the data on production, consumption, inventory changes, and spot and futures prices, given reasonable assumptions about elasticities of supply and demand. We believe that the simplicity and transparency of our approach makes our results quite convincing.

We will show that although we cannot rule out that speculation had *any* effect on oil prices, we can indeed rule out speculation as an explanation for the sharp changes in prices since 2004. Unless one believes that the price elasticities of both oil supply and demand are close to zero, the behavior of inventories and futures-spot spreads are simply inconsistent with the view that speculation has been a significant driver of spot prices. Across the entire sample, speculation *decreased* prices on average or left them essentially unchanged, and

[1] In a more dynamic model, inventories may be drawn down in the presence of speculation on net, if shocks to the market would have led to increases in inventories in the absence of speculation. Pirrong (2008) and Kilian and Murphy (forthcoming) make this point.

reduced peak prices by roughly 5 percent. When we focus on four specific periods of price run-ups, we find that speculation decreased prices by, on average, 1.5 percent.

In the next section we clarify the meaning of speculation (versus, e.g., an "investment"), and discuss the ways in which speculation can occur. We note that the simplest and lowest-cost way to speculate on the price of a commodity is to buy or sell futures contracts. In Section 3 we lay out a simple analytical framework that shows how production, consumption, inventories, and spot and futures prices are connected. Sections 4 and 5 show how this framework can be used to distinguish the effects of speculation (via the purchase or sale of futures) from the effects of shifts in the fundamental drivers of supply and demand. In Section 6 we present our empirical results, and show that there is no evidence that speculation contributed to the observed sharp increases in oil prices. If anything, speculation slightly reduced oil price volatility.

2 Some Basics.

We begin by addressing some basic issues. First, what exactly is meant by "oil price speculation," and how does it differ from a hedging operation, or an investment made as part of a diversified portfolio? Second, how can one actually go about speculating on oil prices? The answers to these questions provide a foundation for an economic analysis of oil price speculation and its effects on production, inventories, and prices themselves.

2.1 What is Meant by "Oil Price Speculation?"

We define oil price speculation as *the purchase (or sale) of an oil-related asset with the expectation that the price of the asset will rise (or fall) to create the opportunity for a capital gain.* A variety of oil-related assets are available as instruments for speculation; oil futures, shares of oil companies, and reserves of oil itself are examples. Thus a speculator might take a long position in oil futures because she believes that the price is more likely to rise than fall, and hopes to "beat the market." (But note that for every long futures position there is an off-setting short position, held by someone betting that the price is more likely to fall than rise.)

An important issue is how to distinguish speculation from an oil-related *investment*, which we define as the purchase or sale of an asset such that the expected net present value (NPV) of the transaction is positive. One example of such an investment is the purchase or sale of oil futures (or other derivatives), not to "beat the market," but instead to hedge against price fluctuations that, if large enough, could lead to bankruptcy. A second example is the purchase of oil-related financial assets, such as futures or oil company shares, to diversify a portfolio.[2] Still another example is the accumulation of oil inventories by producers or industrial consumers of oil as a way to facilitate deliveries and reduce the risk of stock-outs.

As a practical matter, it is often difficult or impossible to differentiate between a speculative activity and an investment. For example, mutual funds, hedge funds, and other institutions often hold futures positions as well as oil company shares, and might do so to make a "naked" (unhedged) bet on future prices, or instead as a way to diversify or to hedge against other oil-related risks. Sometimes it is possible to clearly identify a hedging activity, but more often it is not. So in most cases, what we call an "investment" and what we call "speculation" are likely to be the same thing, or at best ambiguous. Thus when we examine the impact of, e.g., purchases of futures contracts, we will not be concerned about whether the purchase represents an investment or pure speculation.

Although we will not try to distinguish among motivations for purchases of oil-related financial assets, we can be clear about what oil price speculation is *not*: a shift in fundamentals. This could include a shift in consumption demand for oil (e.g., a short-term shift resulting from unusually cold weather, or a long-term shift resulting from increased use of oil in China) or a shift in the supply of oil (e.g., because of a strike or hurricane that shuts down some output). A shift in fundamentals can certainly lead to a change in the price of oil, and we want to distinguish that from a price change caused by speculators or investors betting on a change in price that is not already accounted for by expected shifts in demand and supply (e.g., an expected shift in demand because of seasonal variation in the weather).

We must also be clear about what price or prices we are referring to. When speculation

[2]Note that in this example the expected return on the asset would account for systematic risk and thus would equal the opportunity cost of capital, making the expected NPV of the investment just zero. However, by helping to diversify the portfolio, purchasing the asset would reduce the portfolio's risk.

is blamed for pushing oil prices up or down, it is usually the *spot price* that is being referred to; i.e., the price for immediate delivery. By contrast, the *futures price* is the market price of a futures contract for oil to be delivered at some future point in time. When speculators (or investors) buy and sell futures contracts, the futures price may change, and we will be concerned with whether and how that change can affect the spot price.

2.2 How to Speculate on Oil Prices.

There are several methods by which an individual or firm could speculate on the price of oil. As we will see, some methods are more feasible and less costly than others. We will also explain what each method would do to oil prices, production, and inventories.

The claim that speculation is to blame for changes in the price of oil is usually made when the price is or has been rising, not when the price is falling. But of course one could speculate in either direction, i.e., one could bet that prices will rise or fall. In fact, market equilibrium requires an equal number of "long" bettors (those betting the price will rise) and "short" bettors. For simplicity (and in keeping with the popular press), we will focus on the "long" side, i.e., on speculating that oil prices will rise.

Buy Stocks of Oil Companies. Holding shares of oil company stocks is probably the most common way to speculate (or invest) in oil, even though it is not what most critics have in mind when they call for a ban on speculation. A speculator or investor can focus on companies that are largely in exploration and production, companies that hold large reserves (developed or undeveloped), or integrated companies. This is, of course, what many mutual funds do, and also some hedge funds and individual investors.

Suppose speculators become "bullish" on oil and start to buy oil company stocks, driving up the stock prices for these companies. What would this do to the price of oil itself? In the short run (less than a year or two), it would have no effect on oil production or consumption and thus do nothing to the price of oil. In the longer run, to the extent that the stock prices of oil companies are higher than they would be otherwise, it would lower oil companies' cost of capital. This in turn would encourage investment in exploration and development, and eventually lead to more oil production and thus *lower prices*. But this impact on prices could take several years. In any case, this is certainly not an explanation for the sharp changes in oil prices since 2004.

Hold Physical Oil in situ. An owner of in-ground oil reserves can "speculate" on higher prices by keeping the oil in the ground rather than producing and selling it. Clearly this is something oil companies can do, but not hedge funds, mutual funds, or individual investors. How easily an oil company can "speculate" this way depends on whether the in-ground reserves are *undeveloped* or *developed*. Let's look at each in turn.

Undeveloped reserves are reserves that have been discovered and are owned by the oil company, but the oil cannot be produced until very large sunk cost investments in development are made. Normally, development (construction of production wells, maybe offshore platforms, etc.) will take at least a year or two. Thus an oil company that wants to "bet" on a sharp price increase over the next few years could delay development, and simply keep the reserves undeveloped. There could be (and usually are) other reasons for delaying development, e.g., because of the reserve's option value.[3] Although it is unlikely oil companies would do this, in principle it is a feasible way to "bet" on higher prices.

What would happen to oil markets if companies actually did this? Suppose that around 2004 or 2005, oil companies "bet" on prices rising by withholding development of undeveloped reserves. This would indeed imply lower production levels and higher prices—but only after at least one or two years had passed, given the time it takes to develop a reserve. How would we identify this kind of activity? Normally rising oil prices increase the return from development, and lead to rising rig rental rates and rig utilization. If rig rates and utilization were instead falling, this would provide some evidence that companies were holding back on development. We examine this possibility later in the paper.

Developed reserves have the production wells, pipelines, platforms, and other infrastructure needed to actually produce oil. However, once a reserve is developed and production begins, the rate of production cannot be easily varied. Production usually follows a decline path largely determined by the internal pressure and other physical characteristics of the

[3]An undeveloped reserve gives the owner an option to develop the reserve, the exercise price for which is the cost of development. If there is considerable uncertainty over future oil prices, the option value will be high, so that there is an incentive to keep the option open by delaying development. Siegel et al. (1987) were among the first to calculate the option value of an undeveloped reserve. See Dixit and Pindyck (1994) for a discussion of "real options," including a treatment of the option to develop an oil reserve.

reserve, the size of the wells, etc. Reducing or temporarily stopping production from a fully developed reserve can reduce the total quantity of oil that can potentially be produced, and thus is usually not an economical option. Nonetheless, later in the paper we will examine whether production has fallen below trend during periods of suspected speculation.

Hold Physical Oil Above Ground. Producers and consumers of oil normally hold inventories, which can serve a number of functions. Producers hold them to facilitate production and delivery scheduling and avoid stockouts. If marginal production costs are increasing with the rate of output and if demand is fluctuating, producers can reduce their costs over time by selling out of inventory during high-demand periods, and replenishing inventories during low-demand periods. Industrial consumers also hold inventories, and for the same reasons—to reduce adjustment costs and facilitate production (i.e., when the oil is used as a production input), and to avoid stockouts. In principle, however, inventories could also be held to speculate: if you think the price will rise sharply, you could buy oil and store it in tanks, oil tankers, etc.

Clearly this form of speculation is not something that hedge funds, mutual funds, or individual investors can do. Although it is unlikely that oil companies would speculate in this way, in principle it is feasible, at least if sufficient storage capacity is available. Were oil companies (or industrial consumers of oil) indeed accumulating "excess" inventories during periods of suspected speculation? Here, "excess" means that barring a speculative motive, the marginal unit of inventory is worth less than would normally be the case. We will see how we can test for this possibility using futures price data.

Hold Oil Futures. This is the easiest, lowest cost, and most common way to speculate on oil prices. One would hold a long futures position to speculate on prices going up, or a short futures position to speculate on prices going down. (Remember that every long position must be matched by a short position.) Holding a futures position involves very low transaction costs, even for an individual investor. This is an important means of investment for hedge funds, some ETFs, mutual funds, and also individuals. It is also the most common explanation for how oil price speculation takes place, and is usually the focus of those who criticize the activities of speculators (and investors).

If more people want to go long than short at the current futures price, the futures price will rise. What would that do to the spot price, which is the price we care about? In principle it could push the spot price up, but only under certain conditions. Since the use of futures contracts is the most important means of speculation, we need to look at it in detail.

Hold Other Oil Derivatives. A futures contract is itself a derivative, but other derivatives could be used to speculate on oil prices. Examples include call or put options on oil futures prices. Buying or selling such derivatives is commonly done by hedge funds and also by individuals, and is easy and relatively low-cost. There are also more complex derivatives that are sometimes held by hedge funds. The impact of these derivatives on oil prices is closely related to the impact of futures contracts, so we will ignore the considerable variety of derivatives and focus below on futures.

3 Analytical Framework.

For any storable commodity that is subject to stochastic fluctuations in production and/or consumption, producers, consumers, and possibly third parties will hold inventories. Producers hold inventories to facilitate production and delivery scheduling and avoid stock-outs. If marginal production costs are increasing with the rate of output and if demand is fluctuating, producers can reduce their costs over time by selling out of inventory during high-demand periods, and replenishing inventories during low-demand periods. Industrial consumers also hold inventories, to reduce adjustment costs and facilitate production (i.e., when the commodity is used as a production input), and to avoid stock-outs.

Thus there are two interrelated markets for a commodity: the *cash market* for immediate, or "spot," purchase and sale, and the *storage market* for inventories of the commodity. Although the price of storage is not directly observed, it can be determined from the spread between futures and spot prices. As with any good or service sold in a competitive market, the price of storage is equal to the marginal value of storage, i.e., the flow of benefits to inventory holders from a marginal unit of inventory, and is termed the *marginal convenience yield*. In what follows, we will present a framework that describes the cash market, the market for storage, and the futures-spot spread.[4] We will then use this framework to show

[4] This framework is presented in more detail in Pindyck (2001).

diagrammatically how speculative activity in the futures market—as well as fundamental shifts in supply or demand—can affect spot prices, inventories, and convenience yield.

3.1 The Cash Market.

In the cash market, purchases and sales for immediate delivery occur at a price that we will refer to as the "spot price."[5] Because inventory holdings can change, the spot price does not equate production (which might include imports) and consumption (which might include exports). Instead, the spot price determines "net demand," i.e., the difference between production and consumption. To see this, note that demand in the cash market is a function of the spot price, and perhaps other variables such as the weather and aggregate income, and random shocks reflecting unpredictable changes in tastes and technologies. Because of these shocks and the fact that some of the variables affecting demand (such as the weather) are themselves partly unpredictable, demand will fluctuate unpredictably. We can therefore write the demand function for the cash market as $Q = Q(P; z_1, \epsilon_1)$, where P is the spot price, z_1 is a vector of demand-shifting variables, and ϵ_1 is a random shock.

Supply in the cash market is also a function of the spot price, a set of (partly unpredictable) variables affecting the cost of production (e.g, wage rates and capital costs), and random shocks reflecting unpredictable changes in operating efficiency, strikes, etc. Thus supply also shifts unpredictably, and can be written as $X = X(P; z_2, \epsilon_2)$, where z_2 is a vector of supply-shifting variables, and ϵ_2 is a random shock.

Letting N_t denote the inventory level, the change in inventories is given by the accounting identity:

$$\Delta N_t = X(P_t; z_{2t}, \epsilon_{2t}) - Q(P_t; z_{1t}, \epsilon_{1t}) \ . \qquad (1)$$

We will refer to ΔN_t as *net demand*, i.e., the demand for production in excess of consumption. Thus eqn. (1) says that the cash market is in equilibrium when net demand equals net supply.

[5] The spot price is a price for immediate delivery at a specific location of a specific grade of oil, where the location and grade are usually specified in a corresponding futures contract. In contrast, the "cash price" refers to an average transaction price, usually averaged over a week or a month, and that might include discounts or premiums resulting from relationships between buyers and sellers. We will ignore this difference for now, and use "spot price" and "cash price" interchangeably.

We can rewrite eqn. (1) in terms of the following *inverse net demand function*:

$$P_t = f(\Delta N_t; z_{1t}, z_{2t}, \epsilon_t) \ . \tag{2}$$

Market clearing in the cash market is therefore a relationship between the spot price and the change in inventories.

Because $\partial X/\partial P > 0$ and $\partial Q/\partial P < 0$, $f(\Delta N_t; z_{1t}, z_{2t}, \epsilon_t)$ is upward sloping in ΔN. This is illustrated by the left panel of Figure 3, where $f_1(\Delta N)$ is the inverse net demand function for some initial set of values for z_1 and z_2, and $f_2(\Delta N)$ is the inverse net demand function following an increase in the demand for the commodity (i.e., an increase in z_1), or a decrease in supply (i.e., a decrease in z_2). Note that this upward shift in the inverse net demand function represents a *structural*—as opposed to speculative—change in the market. For crude oil, it might occur because of an increase in Chinese oil demand, or a strike or similar disruption that reduces supply. In Figure 3, we assume that this upward shift is permanent.

3.2 The Market for Storage.

At any instant of time, the supply of storage is simply the total quantity of inventories held by producers, consumers, or third parties, i.e., N_t. In equilibrium, this quantity must equal the quantity demanded, which is a function of price. The price of storage is the implicit "payment" by inventory holders for the privilege of holding a unit of inventory. As with any good or service sold in a competitive market, if the price lies on the demand curve, it is equal to the marginal value of the good or service, i.e., the utility from consuming a marginal unit. In this case, the marginal value is the value of the flow of services accruing from holding the marginal unit of inventory, and is referred to as *marginal convenience yield*. We denote the price of storage (marginal convenience yield) by ψ_t, so the demand for storage can be written as $N(\psi)$.

The storage market is illustrated by the right-hand panel of Figure 3, where N_1 is the supply of storage and ψ_1 is the corresponding price (convenience yield). Note that the marginal value of storage is small when the total stock of inventories is large, because in that case an extra unit of inventory will be of little value, but it rises sharply when the total stock becomes small. Thus $N'(\psi) < 0$ and $N''(\psi) > 0$.

In addition to the price ψ, the demand for storage will depend on other variables. For example, it will depend on expected *future* rates of consumption or production; if a seasonal increase in demand is expected, the demand for storage will increase because producers will want greater inventories to avoid sharp increases in production cost and to make timely deliveries. The demand for storage will also depend on the spot price of the commodity; one should be willing to pay more to store a higher-priced good than a lower-priced one. And the demand for storage may depend on the volatility of price; greater volatility increases the demand for storage by making scheduling and stock-out avoidance more costly.[6] Letting z_3 denote these demand-shifting variables and including a random shock, we can write the inverse demand function as:

$$\psi = g(N; z_3, \epsilon_3) \quad . \tag{3}$$

A structural change could also take the form of a shift in the demand for storage. For example, if oil supply and demand become more volatile, e.g., because of increased volatility of GDP or weather conditions, the demand for storage curve on the right-hand side of Figure 3 will shift upwards, so that if that supply of storage remains fixed at N_1, the price (convenience yield) ψ will increase. The demand for storage curve could also shift for reasons related to speculation, as we explain later.

3.3 Spot Price, Futures Price, and Convenience Yield.

Speculative activity most commonly occurs via the futures market. Thus it is important to understand how the futures market operates and how the futures price can affect the spot price. A *futures contract* is an agreement to deliver a specified quantity of a commodity at a specified future date, at a price (the futures price) to be paid at the time of delivery. Futures contracts are usually traded on organized exchanges, and as a result tend to be more liquid than *forward contracts*, which are also agreements to deliver a specified quantity of a commodity at a specified future date, at a specified price (the forward price). A futures contract differs from a forward contract only in that the futures contract is "marked to

[6] Pindyck (2004) estimates the impact of changes in volatility on inventories and price for crude oil, heating oil, and gasoline.

market," which means that there is a settlement and corresponding transfer of funds at the end of each trading day.

It is not necessary to actually take delivery on a futures (or forward) contract; in fact, the vast majority of futures contracts are "closed out" or "rolled over" before the delivery date, so the commodity does not change hands. The reason is that these contracts are usually held for hedging, investment, or speculation purposes, so there would be no reason to take delivery of the commodity.

Once we know the spot price at time t and the futures price for delivery at time $t+T$, we can determine the convenience yield. Let $\psi_{t,T}$ denote the (capitalized) flow of marginal convenience yield from holding a unit of inventory over the period t to $t+T$. Then to avoid arbitrage, $\psi_{t,T}$ must satisfy:

$$\psi_{t,T} = (1+r_T)P_t - F_{t,T} + k_T \ , \tag{4}$$

where P_t is the spot price at time t, $F_{t,T}$ is the futures price for delivery at $t+T$, r_T is the risk-free T-period interest rate, and k_T is the T-period per-unit cost of physical storage.[7]

We will be interested in how changes in the futures price affect the spot price, so it is useful to rewrite eqn. (4) with the spot price on the left-hand side:

$$P_t = \frac{1}{1+r_T}[F_{t,T} + \psi_{t,T} - k_T] \ . \tag{5}$$

Thus an increase in $F_{t,T}$ will lead to an increase in P_t—unless there is a equivalent decrease in $\psi_{t,T}$ and/or increase in k_T. The drop in $\psi_{t,T}$ could occur if N_t increases. But what if N_t increases to the point that there is almost no more storage capacity? In that case k_T would increase sharply, which once again would limit the impact of an increase in the futures price on the spot price.

As we will see, an increase in the futures price can lead to an increase in the spot price of a commodity, but any impact will be limited by activity in the market for storage. In

[7] To see why eqn. (4) must hold, note that the (stochastic) return from holding a unit of the commodity from t to $t+T$ is $\psi_{t,T} + (P_{t+T} - P_t) - k_T$, i.e., the convenience yield (like a dividend) plus the capital gain minus the physical storage cost. If one also shorts a futures contract at time t, which would yield a return $F_{t,T}F_{T,T} = F_{t,T}P_{t+T}$, one would receive a total return over the period of $\psi_{t,T} + F_{t,T} - P_t - k_T$. The futures contract requres no outlay and this total return is non-stochastic, so it must equal the risk-free rate times the cash outlay for the commodity, i.e., $r_T P_t$, from which eqn. (4) follows.

addition, we can look to the storage market (i.e., the behavior of inventories and convenience yield) to determine whether changes in the spot price are due more to structural shifts in demand and supply, or instead to speculative activity in the futures market.

3.4 Example: Permanent versus Seasonal Shifts in Demand.

The interaction of the cash and storage markets can be seen in Figures 3 and 4, which illustrate the impact of an upward shift in demand. In Figure 3, the shift in demand is expected to be—and is—permanent, e.g., a permanent increase in oil Chinese oil demand. The net demand curve shifts up and the spot price increases from P_1 to P_2. The demand for storage curve remains fixed, and assuming the shift in the net demand curve occurs slowly, there would be no reason for producers or consumers of oil to change their inventory holdings, so the total inventory level remains fixed at N_1.

Figure 4 illustrates an anticipated shift in demand that is expected to be—and is—temporary. For example, this could be a seasonal increase in the demand for oil. Because the increase is anticipated, inventories are accumulated ahead of time (so that N increases from N_1 to N_2), and the spot price increases (from P_1 to P_2) before there is any shift in the net demand curve. When the net demand curve does shift up, inventories are drawn back down, as producers and consumers anticipate that net demand will shift down again. Thus the spot price stays at or near P_2, rather than rising to P_3, the level that would have been reached had there been no changes in inventories. Finally, the net demand curve shifts back down and the spot price returns to P_1.

Note that the spot price changes illustrated in Figures 3 and 4 and the inventory changes in Figure 4 are the result of *structural shifts* in the cash market for oil, as opposed to speculation. In the next section we examine the impact of speculative activity.

4 The Impact of Speculation.

The easiest way to speculate on the price of a commodity, and the one that receives the most attention from those claiming that price changes are caused by speculators, is to buy or sell futures contracts. Thus we focus mostly on futures as the instrument of speculation.

However, we will also consider what happens if producers and/or consumers of the commodity accumulate inventories for speculative purposes.

4.1 Speculation via the Futures Market.

Suppose speculators (and investors) decide to buy futures, thereby driving up the futures price $F_{t,T}$. What will this do to the spot price, inventories, and convenience yield? We will begin by assuming that demand and supply are moderately price elastic, so that the net demand curve slopes up, but not sharply. This is illustrated in Figure 5. Although speculators have pushed up the futures price, there is no shift in net demand $f(\Delta N)$ because there has been no change in the fundamentals driving demand and supply. Nor will there be any shift in the demand for storage curve.

From eqn. (5) we know that equilibrium in the spot and futures markets requires a reduction in $\psi(N)$ and/or an increase in the spot price. Given that $F_{t,T}$ is now high relative to P_t, the payoff from holding inventories is large, so inventories will increase. Thus $\Delta N > 0$, and as shown in Figure 5, the spot price increases from P_1 to P_2. Eventually inventories reach N_2 and convenience yield falls from ψ_1 to ψ_2. At that point, with no further inventory buildup, ΔN falls to zero and the spot price must fall to its original level, P_1. This can be consistent with a higher futures price because $\psi_2 < \psi_1$, so the futures price can remain high even though the spot price falls to where it started.

As the futures contracts reach expiration, the futures price must approach the spot price (and at expiration must equal the spot price). If speculators remain optimistic about prices, they might "roll over" their futures contracts, i.e., sell the near-term futures and buy longer-term futures. In that case inventories will remain at N_2 and the convenience yield will remain at ψ_2, keeping the spot price at P_1. But it is likely that speculative buying of futures will eventually diminish, so that the futures price falls, reducing the expected payoff from holding inventories. Inventories are then sold off, pushing the spot price down (to P_3 in Figure 5). Eventually inventories fall back to N_1 and convenience yield increases to ψ_1, at which point ΔN approaches zero, and the spot price returns to its original level, P_1.

What would we observe in this speculative scenario, and how would it differ from what we would observe in the case of the fundamental shifts illustrated in Figures 3 and 4? In Figure 3 there is a substantial increase in the spot price, but no change in inventories or convenience yield. In Figure 4 there is a temporary increase in the spot price and the inventory level, but those changes follow seasonal patterns. Thus if we deseasonalized the spot price and inventory data, we would observe no change in either the price or inventories. The situation in Figure 5 is quite different. First, the increase in the spot price requires a large increase in inventories, and the spot price would fall back to its original level once inventories peaked and the inventory buildup stopped. Second, as speculative buying of futures slowed or reversed, the spot price would fall below its original level, and inventories would fall back to N_1. Third, there would be no seasonal pattern in either the price or inventory changes.

Note that if demand and supply are very price-inelastic, the impact of speculative buying of futures on the spot price can be much greater, as illustrated in Figure 6. In this case the net demand curve $f(\Delta N)$ is very steep, so a small (but positive) ΔN will be sufficient to cause the spot price to rise considerably. Then inventories will increase only slowly (because ΔN is small), and the higher spot price can be sustained longer. But even if speculation is sufficient to keep inventories at the higher level (N_2 in Figure 6), once inventories stop growing the price will have to decline to its original level (P_1).

4.2 Correct and Incorrect Predictions of Demand and Supply Shocks.

Speculation is often based on beliefs about price changes, rather than a blind gamble. Sometimes those beliefs turn out to be correct and sometimes not. Suppose speculators buy futures contracts because they believe that there will be a change in fundamentals, namely a supply or demand shock that will cause an increase in price. How would such speculation affect the spot price and inventories?

Figure 7 illustrates what happens when speculators *correctly* anticipate a change in fundamentals that causes an upward shift in the the net demand curve. Speculators buy futures *before* the shift occurs, pushing up the futures price. This leads to an increase in inventories (from N_1 to N_2) and an increase in the spot price (from P_1 to P_3). Later the net demand curve shifts up, speculators sell their futures, the futures price falls relative to the (now higher) spot price, so inventories decline back to their original level, and the price declines.

Eventually the inventory sell-off stops and the spot price settles at its new equilibrium level (P_2 in the diagram).

Figure 8 illustrates what happens when speculators *incorrectly* expect a shift in net demand. Once again they buy futures and there is a build-up of inventories as the spot price increases (from P_1 to P_2). But the inventory build-up eventually stops as the spot price declines back to its original level. Inventories are then sold off, pushing the spot price down (to $P_3 < P_1$). Finally the inventory sell-off ends as both inventories and the spot price return to their original levels.

4.3 Speculative Inventory Holdings.

In principle, oil companies (and some oil consumers) could speculate by accumulating above-ground inventories. In effect, this would cause an upward shift in the demand for storage curve. The reason is that now there would be a speculative benefit from holding inventories in addition to the usual benefit of stock-out avoidance, etc.

Suppose oil companies (or individuals with very large bathtubs) decide to accumulate inventories as a speculative bet on rising prices. Figure 9 illustrates this scenario. The demand for storage curves shifts upward. Assuming no change in holdings of futures contracts, as inventories increase (from N_1 to N_2 in the figure), eqns. (2) and (4) imply that the spot price will increase (from P_1 to P_2) because $\Delta N_t > 0$, and therefore the convenience yield must increase (to ψ_2). In the figure, inventories peak at N_3, and as ΔN_t drops to zero, the spot price returns to P_1 and the convenience yield returns to ψ_1.

If the high inventory level N_3 is maintained, there will be no further changes in P_t or ψ_t. If, on the other hand, this speculative episode ends with companies selling off part of their inventories, the spot price and convenience yield will fall (to P_3 and ψ_3 in Figure 9) as inventories decline. Eventually the inventory sell-off ends as N returns to N_1, and both price and convenience yield return to their original levels. Depending on when they bought and sold, some of the speculators may have made money and other lost money. On average speculators will have lost, however, because they will have had to incur the additional costs of physical storage.

4.4 The Limitations of Speculative Effects for Oil.

During the last decade we have seen very large movements in the spot price of oil. For example, from 2007 to 2008, the spot price of WTI crude more than doubled, from about \$60 per barrel to about \$130. Could this have been the result of speculation? In other words, could this price change have occurred with no shift in the net demand curve, as in Figures 5 or 6? One way to answer this question—using data based only on the cash market—is by calculating the change in inventories that would have had to occur as a result.

To do this, we write an equation for net demand, and calibrate it to data for 2007, and apply alternative estimates of supply and demand elasticities. We will assume that supply and demand are isoelastic, and that supply includes imports. Then supply is given by $X = k_S P^{\eta_S}$ and demand is $Q = k_D P^{\eta_D}$. We can express the change in inventories as:

$$\Delta N_t = k_S P_t^{\eta_S} - k_D P_t^{\eta_D} \ . \qquad (6)$$

We will use one month as our time unit, and calibrate to a total US average monthly consumption of 540 million barrels, and price of \$60. For the elasticities of supply and demand, we will use $\eta_S = 0.2$ and $\eta_D = -0.2$. These elasticity numbers are in line with various econometric estimates that have appeared in the literature in recent years. For the period in question, the constants k_S and k_D are then $k_S = 238.1$ and $k_D = 1224.7$. With these constants and elasticities, and with a price of \$60, the quantities demanded and supplied roughly match U.S. data for 2007.

Now, what would it take to reach a price of \$130 *with no shifts in the demand or supply curves*? At a price of \$130, the quantity supplied would rise to 630.3 million barrels per month, and the quantity demanded would fall to 426.1 million barrels per month. This means that inventories would have to increase at a rate of 168 million barrels *per month*! To put this in perspective, the total stock of commercial inventories in the U.S., i.e., excluding the Strategic Petroleum Reserve (SPR), was about 286 million barrels in 2007, and the SPR held around 700 million barrels. A rate of inventory buildup of 168 mb/m would fill the entire SPR in just over four months, and would double total commercial inventories in less than two months. A rate of inventory buildup this large is almost inconceivable, and certainly bears no resemblance to the data. (Over the entire calendar year 2007, commercial inventories *fell* by 28 million barrels, and rose by 55 mb over calendar year 2008.)

Even with less elastic supply and demand, attributing the price increase to something other than a shift in fundamentals is implausible. For example, if we use $\eta_S = 0.1$ and $\eta_D = -0.1$, the constants k_S and k_D become $k_S = 358.6$ and $k_D = 813.2$. Then an increase in price to \$130 with no shift in the supply or demand curves would imply an inventory buildup of 84 million barrels per month, which would result in a doubling of commercial inventories in 3.4 months.

5 Evaluating the Impact of Speculation.

In this section, we lay out a simple and transparent method for decomposing changes in prices into a component coming from changes in market fundamentals and a component resulting from speculative activity. We do this in two ways. The first uses the relationship between supply and demand elasticities, changes in inventories, and prices. The second relies on the relationship between convenience yields, changes in oil stocks, and prices.

5.1 Speculative Changes in Spot Prices.

To begin, we focus on the cash market, and maintain two simplifying assumptions: (i) the supply of oil includes imports, and domestic production and imports are indistinguishable; and (ii) the supply and demand functions are isoelastic, so that eqn. (6) holds.[8] Furthermore, we assume that market fundamentals are incorporated in the supply and demand parameters k_S and k_D, so that a shift in supply or demand would imply a change in one or both of these parameters, rather than in the elasticities η_S and η_D.

Dividing both sides of eqn. (6) by Q_t:

$$\frac{\Delta N_t}{Q_t} = \frac{X_t}{Q_t} - 1 = \frac{k_S}{k_D} P_t^{\eta_S - \eta_D} - 1 \ . \tag{7}$$

Now rearrange and take logs of both sides:

$$(\eta_S - \eta_D) \log P_t = \log k_D - \log k_S + \log\left(\frac{\Delta N_t}{Q_t} + 1\right) \ . \tag{8}$$

[8] We are also assuming that demand includes exports, which in any case are negligible.

If the demand and supply curves are stable over the period that the price is changing, i.e., there is no change in fundamentals, then the supply and demand parameters k_S and k_D are constant, so that taking first differences yields:

$$(\eta_S - \eta_D)\Delta \log P_t = \Delta \log \left(\frac{\Delta N_t}{Q_t} + 1 \right) \ . \tag{9}$$

Since $\Delta N_t = X_t - Q_t$, eqn. (9) can be written equivalently as:

$$(\eta_S - \eta_D)\Delta \log P_t = \Delta \log (X_t / Q_t) \ . \tag{10}$$

This equation explains a price change ΔP_t that results from speculation or investment, as opposed to a change in fundamentals. It says that the sum of the absolute values of the elasticities must equal the percent change in the production-to-sales ratio divided by the percent change in price. Again, we are assuming that the supply and demand parameters, k_S and k_D, incorporate fundamentals. Thus a shift in the demand curve resulting from an increase in Chinese oil consumption, for example, would imply an increase in k_D , but no change in the elasticity η_D. We use eqns. (9) and (10) to test for speculation in the following three ways.

Price Behavior. Beginning with a set of plausible values for the sum of the supply and demand elasticities, $\eta_S - \eta_D$, we can decompose a price change over any period of time into fundamental and speculative components: $\Delta \log(P_T) = \Delta \log(P_S) + \Delta \log(P_F)$. Consider any three-month period, for example. Summing the monthly inventory changes over the three months and dividing by the initial consumption Q_0, eqn. (9) gives the price change that can be attributed to speculation/investment. Subtracting that from the total price change gives the portion that is due to a shift in fundamentals. A comparison of the two components provides a picture of the relative importance of speculation as a driver of price.

Inventory Behavior. We again begin with a set of plausible values for the sum of the supply and demand elasticities, $\eta_S - \eta_D$. Now suppose the price change over some period (say three months) is entirely due to speculation. Rearranging eqn. (9), this would imply:

$$\frac{\Delta N_t}{Q_t} + 1 = \left(\frac{\Delta N_0}{Q_0} + 1 \right) \left(\frac{P_t}{P_0} \right)^{\eta_S - \eta_D} \ . \tag{11}$$

If speculation was a substantial cause of the price changes, this inventory change should be close to the actual inventory change.

Elasticities. Finally, given the data for price and inventory changes, we can use eqn. (9) to determine the sum of the elasticities that would be required to reconcile actual price and inventory changes with pure speculation:

$$\eta_S - \eta_D = \frac{\log(\Delta N_t/Q_t + 1) - \log(\Delta N_0/Q_0 + 1)}{\log P_t - \log P_0} \ . \tag{12}$$

5.2 Speculative Inventory Holdings and Convenience Yield.

The tests described above are all based on equilibrium in the cash market. They rely on the link between price changes and inventory changes that must hold if there are no shifts in supply or demand that are due to fundamentals, i.e., no changes in the parameters k_S and k_D. However, as explained earlier, speculation via inventory accumulation can manifest itself in the market for storage.

To see this, write the (inverse) demand for storage curve as:

$$\psi(N_t) = P_t g(N_t) = k_N P_t N_t^{-1/\eta_N} \ . \tag{13}$$

where $\eta_N > 0$ is the price elasticity of demand for storage. This is a standard specification for the demand for storage, and has been estimated in the literature for a wide variety of commodities. As discussed later, we estimated this equation using our data for crude oil and found that $\eta_N \approx 1$, consistent with other econometric studies. Note that the marginal value of storage is proportional to the price, P_t, of the commodity being stored.

The parameter k_N captures other factors that might affect the demand for storage. Those factors might reflect fundamentals; for example, an increase in market volatility or an increased threat of war in the Persian Gulf would cause an increase in k_N. But a change in k_N might also (or instead) reflect speculation. Earlier we considered the possibility that oil producers decide to accumulate inventories as a means of speculating on price increases. As illustrated in Figure 9, this would cause a shift in the demand for storage curve because there would now be a speculative benefit from holding inventories in addition to the usual benefit. In other words, speculative inventory accumulation would be reflected by an increase in k_N.

Taking logs and first differences of eqn. (13) gives:

$$\Delta \log \psi_t = \Delta \log P_t - (1/\eta_N)\Delta \log N_t + \Delta \log k_N \ . \tag{14}$$

Absent any substantial change in volatility or the threat of war (which we will assume to be the case), the last term in eqn. (14) would reflect a shift in the demand for storage attributable to speculation. As explained earlier, marginal convenience yield can be measured directly from the spread between spot and futures prices. Thus, as with price behavior in the cash market, we can use eqn. (14) to compare the behavior of the actual convenience yield with what it would be in the absence of speculation.

To do this comparison, we use eqn. (14), with $\Delta \log k_N = 0$, to compute a counterfactual series for ψ_t, i.e., values of ψ_t that would we would observe if there were no speculation-induced changes in P_t, N_t, and in the demand for storage curve. We compare this to the actual series for ψ_t to assess the possibility of speculation-driven inventory accumulation.

6 Were Oil Prices Driven by Speculation?

We now turn to the data to test whether changes in oil prices after 2000 can be attributed, even in part, to speculation. As explained earlier, although speculation is most easily done using futures contracts, in principle oil companies could speculate on rising prices by accumulating above-ground inventories, by stopping or slowing down the development of undeveloped reserves (which would result in a drop in the rental and utilization rates of drilling rigs), or by slowing down the production from developed reserves. We examine these last two possibilities first, and then turn to the use of futures contracts as the vehicle for speculation, and to the use of inventory accumulation.

Our data come from the Energy Information Administration (EIA). We collected monthly data on US production, commercial stocks, imports, and exports. We construct consumption as the change US production plus net imports minus changes in commercial stocks. The EIA also report spot and futures prices for oil. We use the WTI price although the results change little if we instead use Brent crude prices. Our sample runs from January 1999 to June 2012.

6.1 Speculation by Oil Companies.

Might oil companies have contributed to the sharp price increases we have seen by delaying the development of undeveloped reserves? If this were the case, we would expect to see a drop in the utilization of drilling rigs in advance of the observed price increases. Figure 10 shows average and utilization rates from 2000 onwards in the Gulf of Mexico, plotted along with the WTI crude spot price.[9] Clearly these data are inconsistent with the view that development delays drove price increases. Note that rig utilization rates were roughly level during 2004–2007, increased in early 2008 as the crude spot price increased, and then dropped shortly after the steep plunge in the spot price.

Might oil companies have contributed to price increases by reducing production from developed reserves (even though doing so would damage the reserves and is thus costly)? We address this possibility by looking at the behavior of production. Figure 11 plots U.S. crude production along with the WTI spot price, again from 2000 onwards. We also include the predicted-production levels based on production prior to 2007. The smooth downward sloping curve is a quadratic trend line fit to the production series over 1999 through the end of 2006, and then extrapolated forward through 2012. Observe that this downward trend ended by 2007, well before the 2008 price spike. Production leveled out during 2006 to 2008.[10]

6.2 Speculation via the Futures Market.

As discussed above, we can calculate counterfactual prices that would have occurred in the absence of speculation by decomposing observed prices changes into a component attributable to speculative activity, resulting from either changes in inventories or convenience yields, and a component due to changes in market fundamentals. Using eqn. (9), along with numbers for supply and demand elasticities, we can also calculate the inventory changes that would be required if the observe price changes are the result of speculation. And given observed price and inventory changes, eqn. (9) also yields the sum of supply and demand elasticities required for speculation to have led to the change in price. Finally, we can use eqn. (14) to calculate counterfactual prices based on changes in convenience yields. In this

[9] We purchased these data from RigZone. The data report utilization of jackups, semi-subs, and drill-ships in the Gulf of Mexico.

[10] The downward spike 2005 was the result of Hurricane Katrina (both a supply and demand shock). We are unable to find any weather-related cause of the drop in production in February 2009; however, this drop corresponds to a large drop in consumption (11 percent).

section we discuss the details of these exercises.

We examine price and inventory changes for non-overlapping three-month and twelve-month intervals. Each price and inventory change is calculated on a moving month-to-month basis. For example, for three-month intervals, we compare the average price for April, May, and June 2005 to the average price for January, February and March 2005. We then compare the average price for May, June, and July 2005 to February, March, and April 2005, and so on. Thus we have a different set of price and inventory changes for each month in our sample. We use intervals of varying length because we are interested in whether speculation may have short-term effects that dissipate over longer periods.[11]

For any given time interval, we calculate the consumption-weighted spot price, average consumption, average stock levels, and the change in inventories over the interval. When our analysis focuses on X-month intervals, the differences in eqns. (9) and (14) are defined as X-month differences. We calculate these X-month differences for every month in our sample. We are interested in speculative activity beyond the normal response to seasonal patterns in the demand for oil. Therefore, we de-seasonalize inventories by first regressing changes in inventories on a full set of month dummies and take the residuals as our measure of inventory changes.[12] Thus we measure speculative activity in terms of how changes in inventories over any X-month interval differ from their average changes during the same interval across the entire sample.

Generating counterfactual prices using eqn. (9) requires estimates of supply and demand elasticities. These elasticities will obviously vary depending on the amount in time over which supply and demand can adjust to price changes. Studies by Dahl (1993), Cooper (2003), and Hughes et al. (2008) suggest that the short-run demand elasticity is roughly -0.1. Although, Kilian and Murphy (forthcoming) estimate a short-run elasticity of roughly -0.25. Dahl (1993) and Cooper (2003) find that the long-run demand elasticity is in the range of -0.2 to -0.3. The literature on supply elasticities is more sparse. Dahl and Duggan (1996) summarize the literature on supply elasticities and find that many estimates, for both

[11] We have also done the analysis using monthly intervals and obtained similar results, but the counterfactual prices are much more volatile.

[12] We observe no seasonality in the convenience yield.

short- and long-run elasticities, are noisily estimated and often the wrong sign. Hogan (1989) estimates a short-run elasticity of 0.09 and a long-run elasticity of 0.58. It is easy to see how short-run supply elasticities would be close to zero. We note, however, that what matters for generating counterfactual prices is the *sum* of elasticities.

We begin by showing results based on a sum of supply and demand elasticities of 0.2, consistent with a supply and demand elasticity of 0.1 and -0.1 respectively, for the three-month intervals. For the twelve-month intervals we use a sum of elasticities of 0.4, consistent with elasticities of 0.2 and -0.2. We investigate the implications of alternative elasticity assumptions in Appendix A. We use these same elasticity assumptions when we construct the inventory changes that would be required for speculation to have led to the observed price changes.

To calculate counterfactual values for the convenience yield using eqn. (14), we need a times series for the convenience yield, along with an estimate of the price elasticity of the demand for storage, η_N. Using eqn. (4), the components of the convenience yield are the risk-free rate of interest, the spot price, a futures price, and the cost of storage (i.e., the cost of storing a barrel of oil over the length of the futures contract). We use the 3-month T-bill rate for the risk-free rate of interest and the price of the three-month futures contract as our futures price. This gives a three-month gross convenience yield, which is plotted in Figure 12. There is little data on the cost of storage; a rough estimate is \$1.50 per barrel per month, but the cost can rise when inventory levels are large and storage facilities fill up, and for some periods the value of the cost of storage can significantly affect our counterfactual prices. We begin the analysis using a net convenience yield based on a monthly storage cost of \$1.50 per barrel. However, there are occasions when futures prices were much larger than spot prices, so a constant storage cost of \$1.50 would imply a negative net convenience yield, violating the arbitrage condition. This occurs in 5 of the 162 months in our sample. For example, in December 2008, the gross convenience yield is $-\$6.08$. When we aggregate to three-month periods, we are left with three consecutive periods where a constant storage cost of \$1.50 yields a negative net convenience yield. We therefore truncate the three-month net convenience yield below at \$1.50.

A reasonable value for the price elasticity of demand for storage, η_N, is 1.0.[13] However, we estimate this elasticity and also investigate our assumption that changes in the convenience yield are proportional to changes in prices. We estimate eqn. (14) for both the three-month and 12-month intervals over our sample (1999 to 2012) assuming an AR(2) process for the error term. Table 1 reports the results. Both the three- and twelve-month data are consistent with our assumption that changes in convenience yields are directly proportional to changes in spot prices; we cannot reject a coefficient of 1 at any conventional level (p-values of 0.18 and 0.64, respectively). In addition, we cannot reject a coefficient of -1 for the change in the log of stocks (p-values of 0.78 and 0.61, respectively).[14]

6.3 Results: Prices, Inventories, and Elasticities.

Price Changes. We begin by calculating and plotting estimates of spot prices that would have been observed had no speculative (or investment) activity taken place, i.e., counterfactual prices that would have changed only in response to changes in fundamentals. Figures 13 and 14 plot actual prices and counterfactual prices using changes in inventories for the three- and twelve-month intervals, respectively.

For the three-month intervals, using changes in inventories to construct price movements arising from only changes in fundamentals leads to counterfactual prices that are very close to the actual prices; the correlation is 0.96. The average spot price over this period was \$55.37 per barrel, and the average counterfactual spot price is \$55.34. The peak counterfactual price is 7 percent higher than the actual price, \$144.90 compared to \$130.85, and the volatility of the counterfactual prices is essentially the same, with the standard deviation in the counterfactual prices being \$29.04, while the standard deviation of actual prices \$28.79. These results show that (i) speculation can account for very little of the observed price changes; (ii) speculation did not cause an increase in price volatility; and (iii) price spikes would have been slightly higher absent speculation.

For the twelve-month intervals, the resulting counterfactual prices are even closer to the

[13] Pindyck (1994) estimated the storage price elasticity to be about 1.1 for copper and 1.2 for heating oil.

[14] It is clear, however, that the confidence intervals around the change in the log of stocks are quite wide. This is likely due to the small amount of variation in stocks.

actual prices. The correlation between the two price series is 0.9997. The average spot and counterfactual prices over this period were both \$53.06 per barrel, with a standard deviation of \$27.05. The peak counterfactual price is slightly lower, \$107.50 compared to \$107.85. The reason for this high correlation is the average change in inventories across 12-month intervals is only 1.01 million barrels (compared to an average commercial stock of oil of 319 million barrels), implying counterfactual prices over a 12-month period will mirror actual prices.

Inventory Changes. We now turn to inventory changes as the basis for assessing the possible role of speculation. To do this, we calculate the inventory changes that would have been needed if the observed changes in actual prices were due to speculation rather than changes in fundamentals, i.e., we use eqn. (11).

Figure 15 plots the actual and simulated (counterfactual) inventory changes for the three-month intervals. The most striking result is that these two series are *negatively* correlated; the correlation between the two is -0.54. Also, the actual inventory changes are much larger in magnitude than would have had to occur if price changes were completely due to speculation, as opposed to shifts in fundamentals. The average change in actual inventories over the sample is 0.98 million barrels, compared to an implied mean change of 6.16. As with price behavior, the observed changes in inventories are inconsistent with speculation.

Figure 16 plots the actual and counterfactual inventory changes for the twelve-month intervals. Once again, the two series are negatively correlated (-0.23). The implied changes in inventories swamp the actual changes. The mean implied inventory change is nearly 600 million barrels compared to an actual mean of 1.62 million barrels.

Sum of Elasticities. Finally, using eqn. (12), we calculate the sum of elasticities (i.e., $\eta_S - \eta_D$) required to rationalize the observed changes in inventories and prices as due purely to speculation. The three-month interval results are shown in Figure 17, truncated at $+/-$ 0.4 for visual ease. Observe first that the sum of the elasticities fluctuates wildly, with no consistent pattern. In fact, nearly half of the time, the sum is negative—i.e., it has the wrong sign. Also, note that the sum of the elasticities is on average very close to zero (0.04 for our sample). Quarterly elasticities of supply and demand this small are simply implausible.

The twelve-month interval results are shown in Figure 18, also truncated at $+/-$ 0.4. Given the longer time intervals, we would expect implied elasticities to be larger compared to the three-month interval results. In fact, the mean implied sum of elasticities is an order of magnitude smaller (0.003).

The results shown in Figures 15 and 16 for inventory changes, and Figures 17 and 18 for the required sum of elasticities, are completely inconsistent with the notion that speculation has been a major driver of oil prices.

6.4 Results: Changes in Convenience Yield.

We now turn to the possibility that speculators drove up oil prices by accumulating above-ground inventories, with the hope of selling them at a higher price. Recall that this would imply a change in k_N in eqn. (13). Thus by holding k_N fixed, we can generate a counterfactual series for convenience yield (for which there is no speculation) and compare it to the actual series (for which there might have been speculation).

As discussed above, a one-time increase in the speculative demand for above-ground inventories will shift the demand for storage curve, $\psi(N)$ upwards, so that both inventories and the convenience yield increase. Thus if speculative inventory accumulation was at work, we would observe counterfactual convenience yields—those that would have occurred in the absence of speculation—that are below the actual convenience yields. Figures 19 and 20 show the actual and counterfactual convenience yields (the latter implied by eqn. (13)) for three-month and 12-month intervals, respectively.

These results are inconsistent with speculative inventory accumulation. In fact, the average fundamentals-only convenience yield is slight larger than the actual for both the three-month and twelve-month periods. For the three-month intervals, the counterfactual convenience yield is on average roughly 5 percent higher than the actual, while it is 19 percent higher for the twelve-month intervals. Furthermore, the volatility in the counterfactual convenience yields, measured by its standard deviation, exceeds the actual convenience yield by 16 and 46 percent for the three- and twelve-month periods, respectively. These results suggest that if anything, “speculation” tended to *decrease* the demand for storage and reduce the volatility of convenience yields.

6.5 Focusing on Specific Periods.

Next, we focus on specific time periods during which prices increased sharply and there was intensive public concern over speculation. Figure 21 plots WTI spot prices and Google search intensity for the term "oil speculation."[15] Because search may occur with some lag, we begin the "epochs" at the beginning of the price run-up and end at the maximum price.

We analyze four epochs, for which the beginning and end points are shown in Figure 21 by a green-solid line and red-dotted line respectively. Note that the last two epochs are subsets of the second one. The epochs are:

1. January 2007 to July 2008
2. February 2009 to April 2011
3. February 2009 to April 2010
4. September 2010 to April 2011

We chose these epochs because they encompass periods of sustained prices increases as well as heavy Google search activity. We split the second interval into two sub-epochs because of the leveling off of prices in the middle of the interval.

We examine the behavior of price, inventories, and convenience yield for these epochs using the same methods as before: (a) We generate a counterfactual final price for the epoch, i.e., the price that would prevail absent speculation; (b) we calculate the required inventory changes for speculation to have caused the observed price increase, given assumptions on supply and demand elasticities; (c) we calculate the sum of supply and demand elasticities required for the observed changes in inventories to have caused the price increase; and (d) we calculate the no-speculation change in the convenience yield and compare it to the actual. We use our long-run supply and demand elasticity assumptions for the first three periods (0.2 and -0.2), since they exceed a year in length. To be conservative we use our short-run elasticities (0.1 and -0.1) for the final period, which is seven months long. The results are shown in Table 2.

[15] Google Insights data allow one to track the intensity of search for a particular term. Within the time period specified, the Insights data report the relative intensity of search for that term. So, the week with maximum search intensity is scaled at 100, and all other weeks are a percentage of the maximum week. Figure 21 plots the weekly average within a particular month.

We begin with prices. Observe that for all four epochs, the counterfactual prices that remove speculative activity are extremely close to the actual ending prices. In three of the four epochs the fundamentals-only price is higher than the actual price. This is consistent with the previous sets of results which show that speculation had almost no impact on prices, and if anything, speculation dampened price spikes.

The next panel of Table 2 shows the build up in inventories required for speculation to have caused the observed price increase. For all four epochs, huge inventory increases would have been required had price increases been drive by speculation, whereas the actual inventory changes were very small. In the first epoch, the required increase in inventories is nearly as large as the level of commercial inventories present at the end of the epoch, whereas actual inventories fell slightly. (The required increase is also more than one-third of the Strategic Petroleum Reserve.) The required inventory build-ups in the other three epochs are also unrealistically large.

The implied elasticities consistent with speculation-induced price increases are likewise unreasonable. For the first three epochs, the sum of elasticities would have to be negative. In the fourth, the sum is close to zero.

Finally, we calculate the change in the convenience yield that would result from only changes in fundamentals. Had the price increases been partly a result of speculative inventory accumulation, we would have observed an increase in the convenience yield larger than what would be justified by fundamentals. Instead, for the first three epochs, the actual increase in the convenience yield was *smaller* than what would be justified by fundamentals. (For the fourth epoch, the actual change is only slightly larger than that justified by fundamentals.) Again, these results are inconsistent with the notion that speculation drove up spot prices through the storage market.

7 Conclusions.

We have shown how a simple model of equilibrium in the cash and storage markets for a commodity can be used to assess the role of speculation as a driver of price changes. With reasonable assumptions about elasticities of supply and demand, the model can be used to determine whether speculation is consistent with the data on production, consumption, inventory changes and spot and futures prices. Given its simplicity and transparency, we believe that our approach yields results that are quite convincing. We have focused on the price of crude oil because sharp increases in oil prices have often been blamed on speculators, but our approach can be applied equally well to other commodities.

We found that although we cannot rule out that speculation had *any* effect on oil prices, we can indeed rule out speculation as an explanation for the sharp changes in prices since 2004. Unless one believes that the price elasticities of both oil supply and demand are close to zero, the behavior of inventories and futures-spot spreads are simply inconsistent with the view that speculation has been a significant driver of spot prices. If anything, speculation had a slight stabilizing effect on prices.

The simplicity of our approach to speculation is a benefit, but also implies limitations, and it yields results that are in some sense qualitative in nature. For example, we have not tried to estimate the specific fraction of each price increase in Table 2 that is attributable to speculation. Instead, we conclude that any effects of speculation were negligible or at most very small. Why don't we use the equations of our model to come up with precise estimates of speculative effects? Because the model is too simple. It assumes that demand and supply in the cash market are isoelastic functions of price, and that the elasticities do not change over time. It also assumes that imports can be combined with domestic supply and respond to price changes in the same way. And it assumes that apart from shifts in the multiplicative parameter k_N, the demand for storage is completely stable. We believe these assumptions are all reasonable—as long as we acknowledge the inability of our approach to estimate speculative effects with precision.

Finally, as we explained at the outset, it is difficult or impossible to distinguish "speculation" from an "investment." The latter might involve buying or selling futures, not to "beat the market," but instead to hedge against large price fluctuations. Or it might involve the purchase or sale of commodity-related financial assets, such as futures or company shares, to diversify a portfolio. Mutual funds, hedge funds, and other institutions often hold futures positions, but it is usually impossible to know whether they are doing so to make a "naked" (unhedged) bet on future prices, or instead to diversify or hedge against other commodity-related risks. Thus when we examined the impact of increased purchases of futures contracts, we were not concerned with whether this represented an investment or pure speculation, and our use of the word "speculation" should always be interpreted as including investment activities—but not a shift in fundamentals.

References

Cooper, J. C. (2003): "Price Elasticity of Demand for Crude Oil: Estimates for 23 Countries," *OPEC Review*, 27, 1–8.

Dahl, C. (1993): "A Survey of Oil Demand Elasticities for Developing Countries," *OPEC Review*, 17, 399–419.

Dahl, C. and T. E. Duggan (1996): "U.S. energy product supply elasticities: A survey and application to the U.S. oil market," *Resource and Energy Economics*, 18, 243 – 263.

Dixit, A. K. and R. S. Pindyck (1994): *Investment Under Uncertainty*, Princeton University Press.

Fattouh, B., L. Kilian, and L. Mahadeva (forthcoming): "The Role of Speculation in Oil Markets: What Have We Learned So Far?" *Energy Journal*.

Hamilton, J. D. (2009a): "Causes and Consequences of the Oil Shock of 2007-08," *Brookings Papers on Economic Activity*, 40, 215–283.

——— (2009b): "Understanding Crude Oil Prices," *The Energy Journal*, 30, 179–206.

Hogan, W. W. (1989): "World Oil Price Projection: A Sensitivity Analysis," Tech. rep., Harvard University, Cambridge, MA.

Hughes, J. E., C. R. Knittel, and D. Sperling (2008): "Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand," *Energy Journal*, 29.

Juvenal, L. and I. Petrella (2011): "Speculation in the Oil Market," Tech. rep., Research Division, Federal Reserve Bank of St. Louis, St. Louis, MO.

Kilian, L. and D. P. Murphy (forthcoming): "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil," *Journal of Applied Econometrics*.

Pindyck, R. S. (1994): "Inventories and the Short-Run Dynamics of Commodity Prices," *RAND Journal of Economics*, 25, 141–159.

——— (2001): "The Dynamics of Commodity Spot and Futures Markets: A Primer," *The Energy Journal*, 22, 1–29.

——— (2004): "Volatility and Commodity Price Dynamic," *The Journal of Futures Markets*, 24, 1029–1047.

Pirrong, S. C. (2008): "Stochastic Fundamental Volatility, Speculation, and Commodity Storage," Tech. rep., University of Houston, Houston, TX.

Siegel, D., J. Smith, and J. Paddock (1987): "Valuing Offshore Oil with Option Pricing Models," *Midland Corporate Finance Journal*, 5, 22–30.

Smith, J. L. (2009): "World Oil: Market or Mayhem?" *Journal of Economic Perspectives*, 23, 145–64.

8 Figures

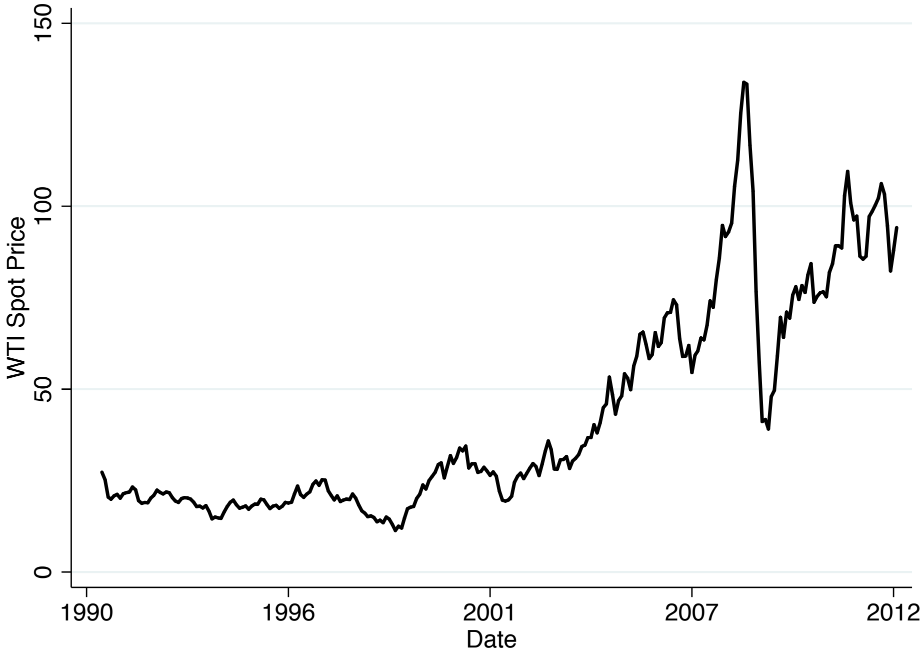

Figure 1: Monthly Spot Price of WTI Crude Oil, 1990–2012

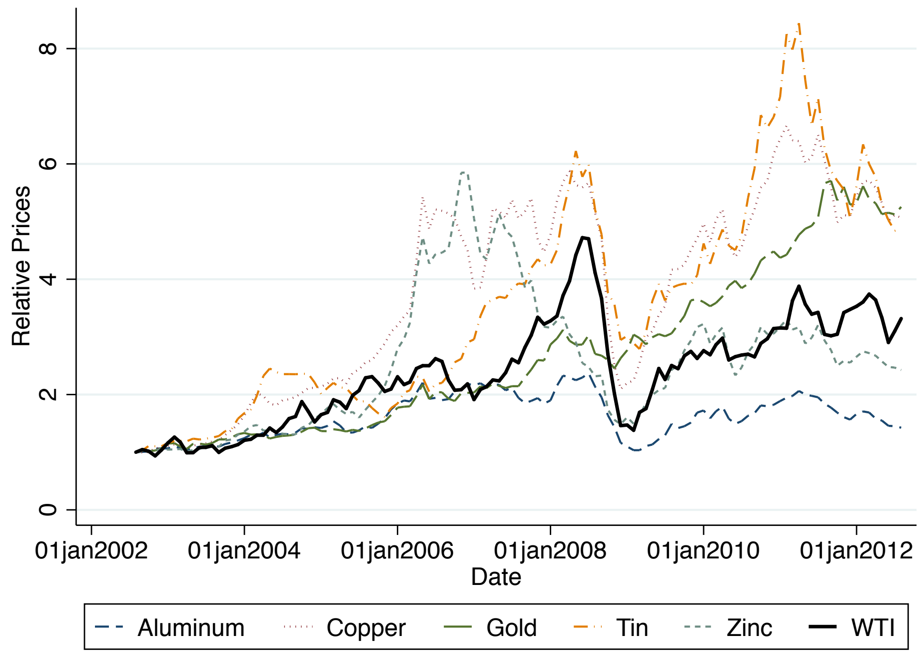

Figure 2: Normalized Commodity Prices Since 2002

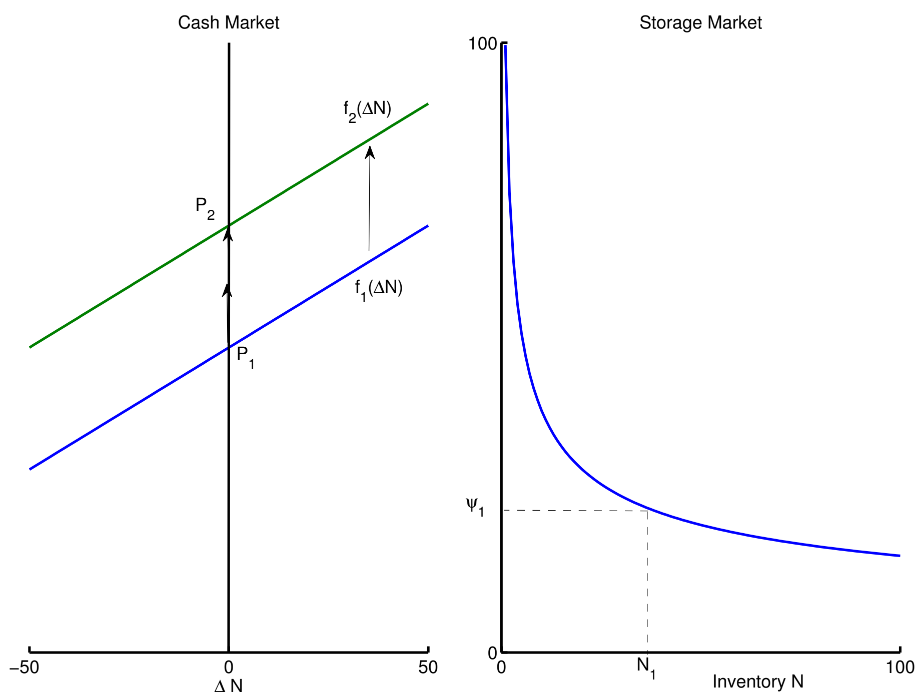

Figure 3: Permanent Increase in Demand for Oil

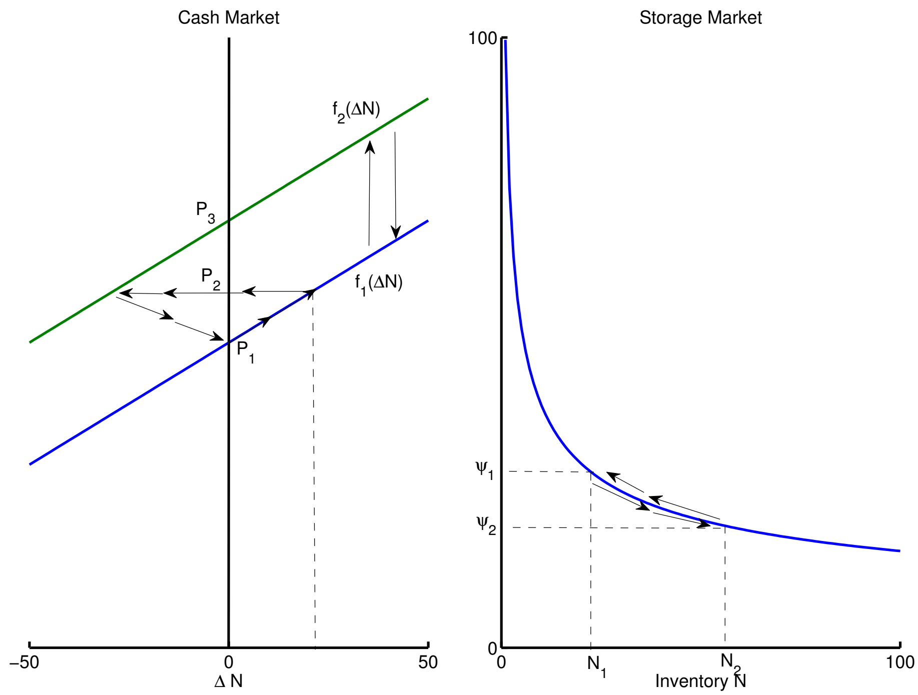

Figure 4: Seasonal (and Anticipated) Changes in Demand

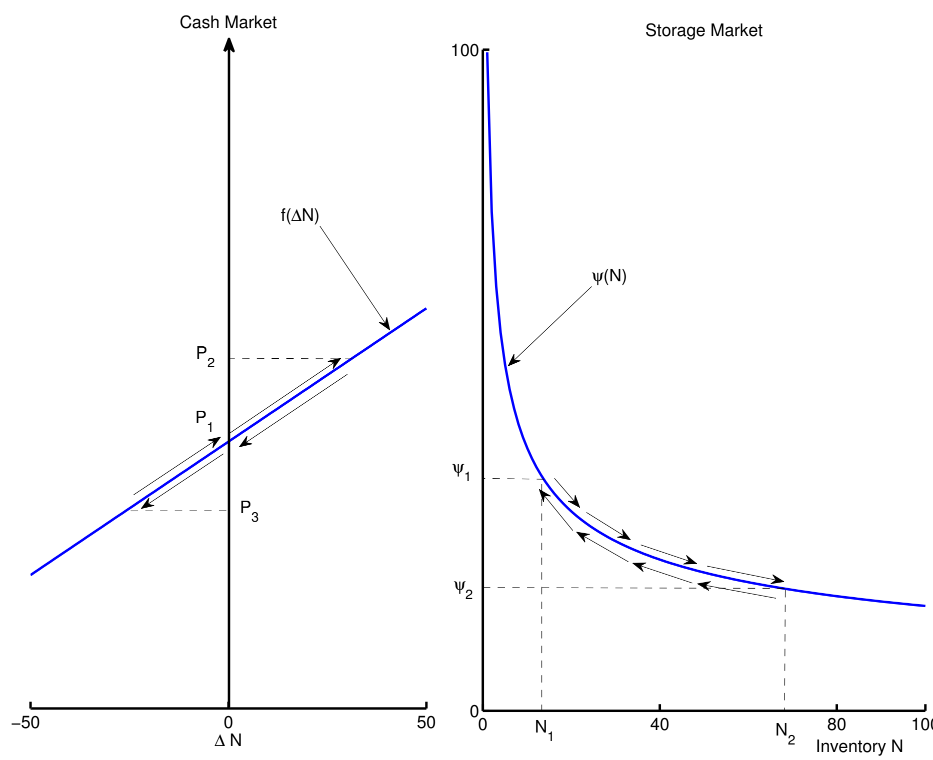

Figure 5: Impact of Speculation on Cash and Storage Markets

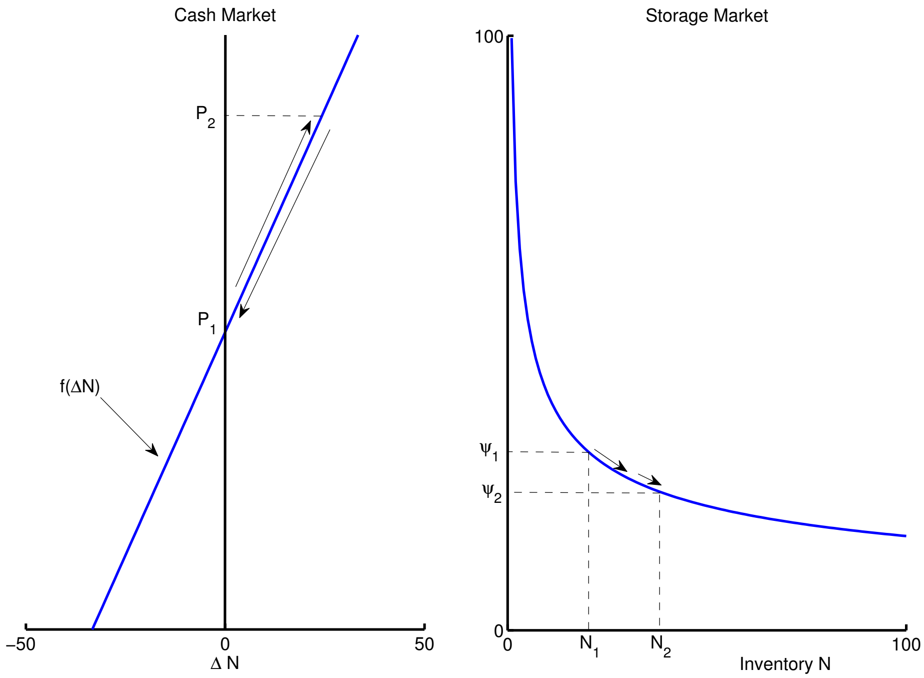

Figure 6: Impact of Speculation on Cash and Storage Markets—Inelastic Demand and Supply

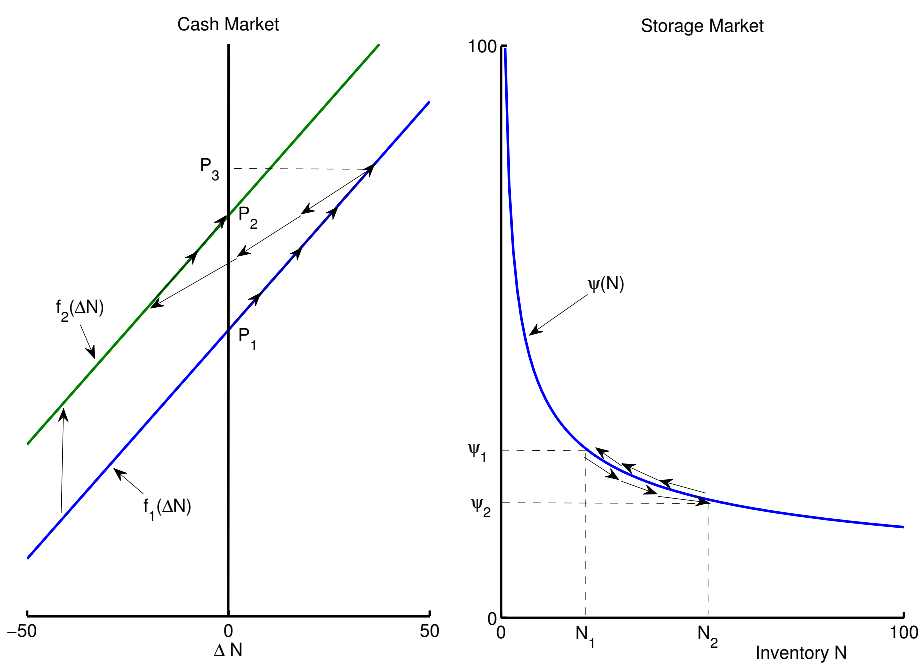

Figure 7: Speculators *Correctly* Predict a Demand or Supply Shock

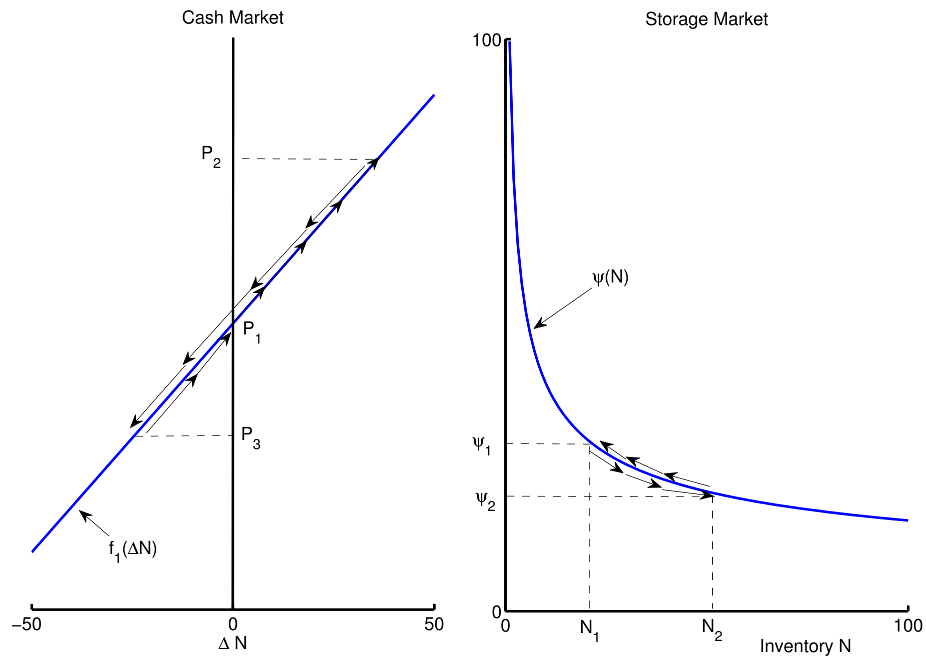

Figure 8: Speculators *Incorrectly* Predict a Demand or Supply Shock

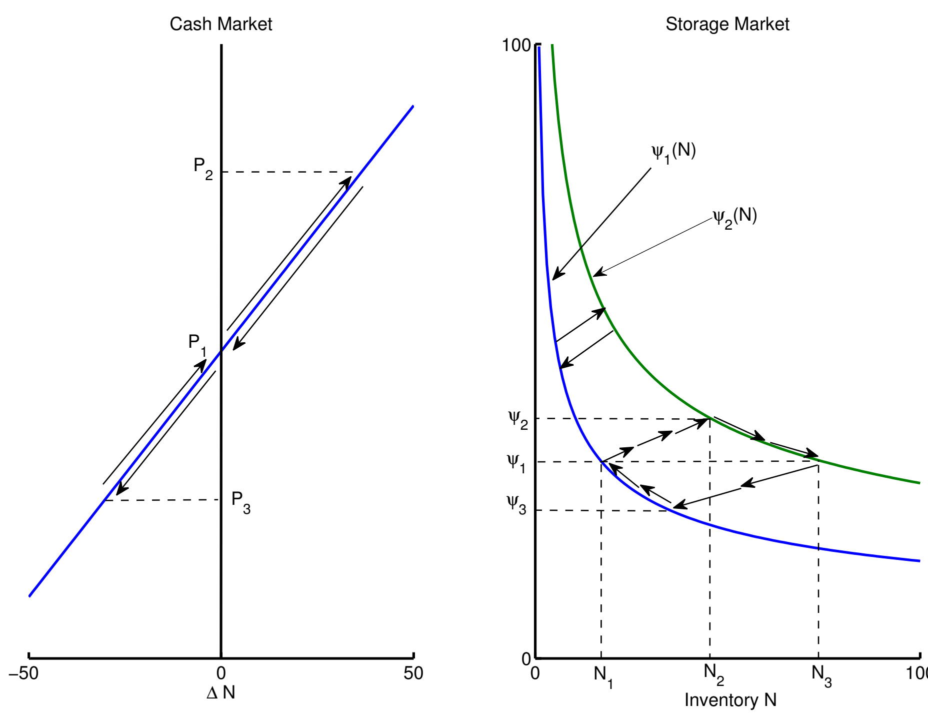

Figure 9: Speculation via Inventory Accumulation

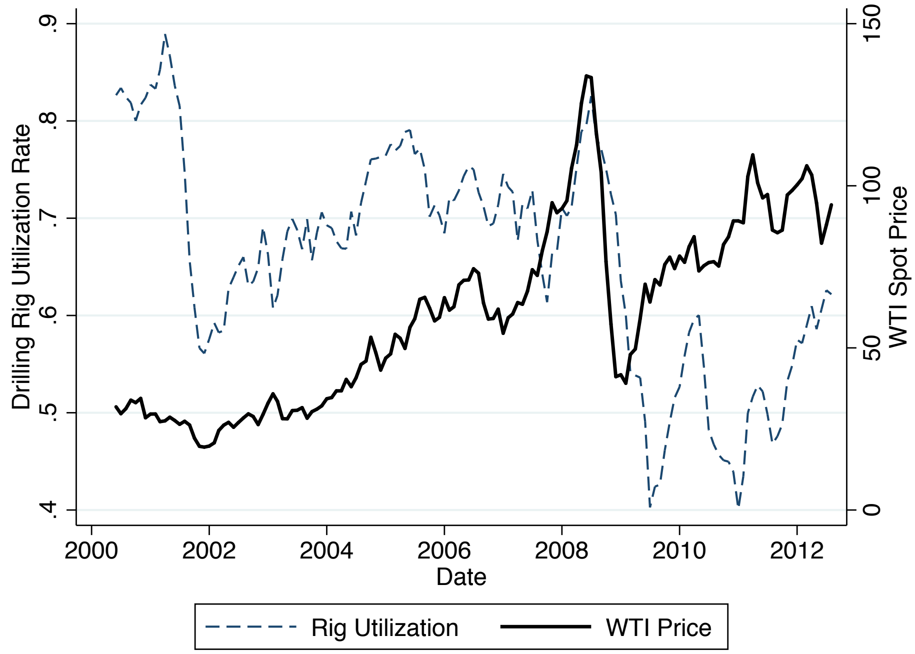

Figure 10: Gulf of Mexico Rig Utilization Rates and WTI Spot Prices

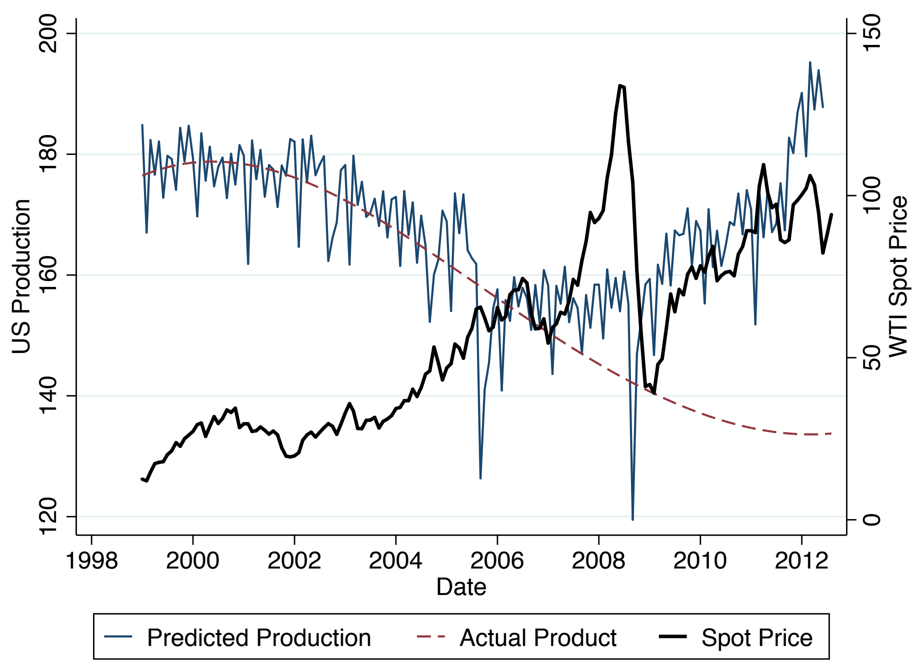

Figure 11: Production, Predicted Production, and WTI Spot Prices

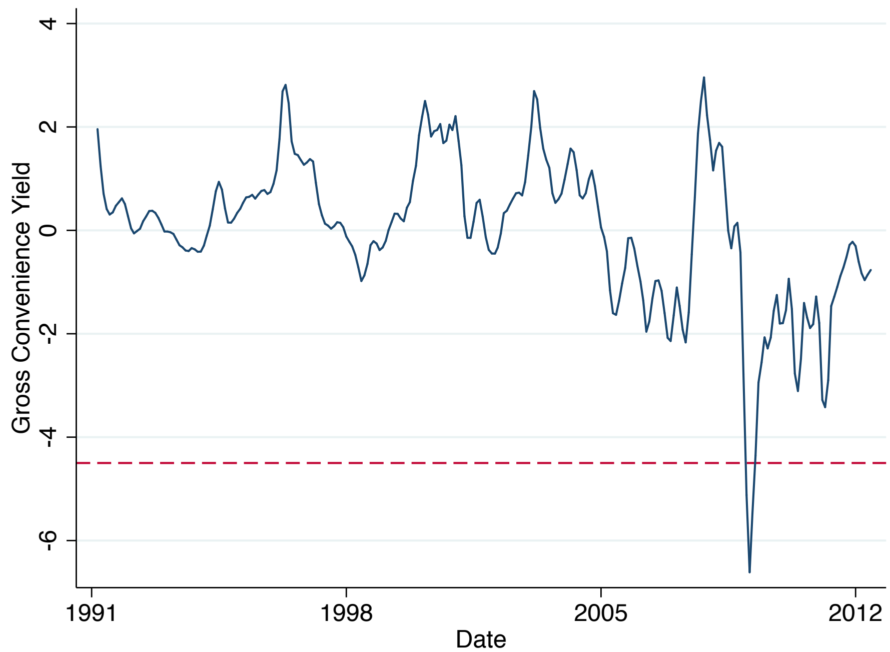

Figure 12: Gross Convenience Yield over Time

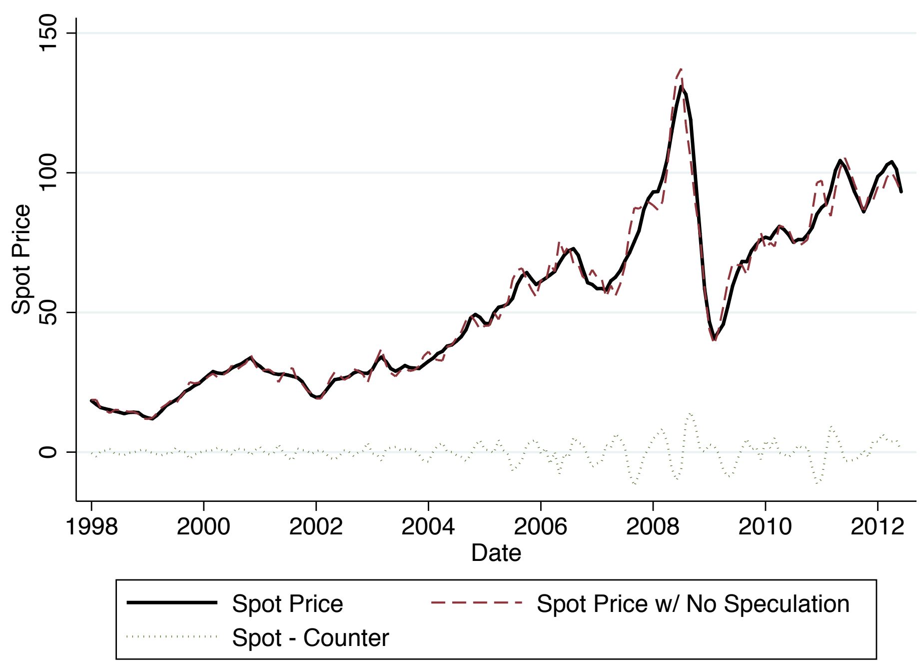

Figure 13: Actual Prices and Implied Prices with No Speculative Activity: Using Inventory Changes and Three-Month Periods

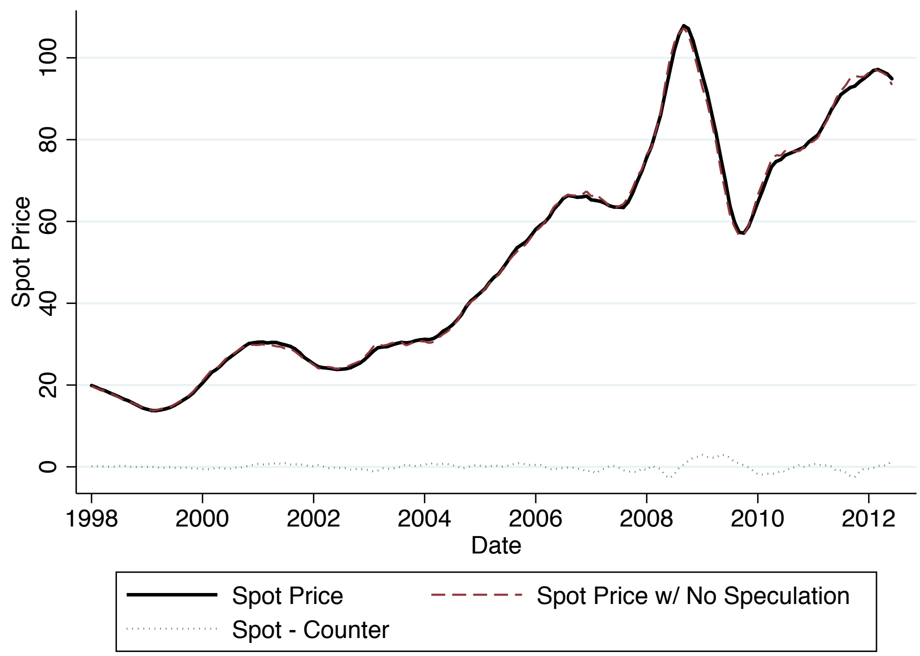

Figure 14: Actual Prices and Implied Prices with No Speculative Activity: Using Inventory Changes and Twelve-Month Periods

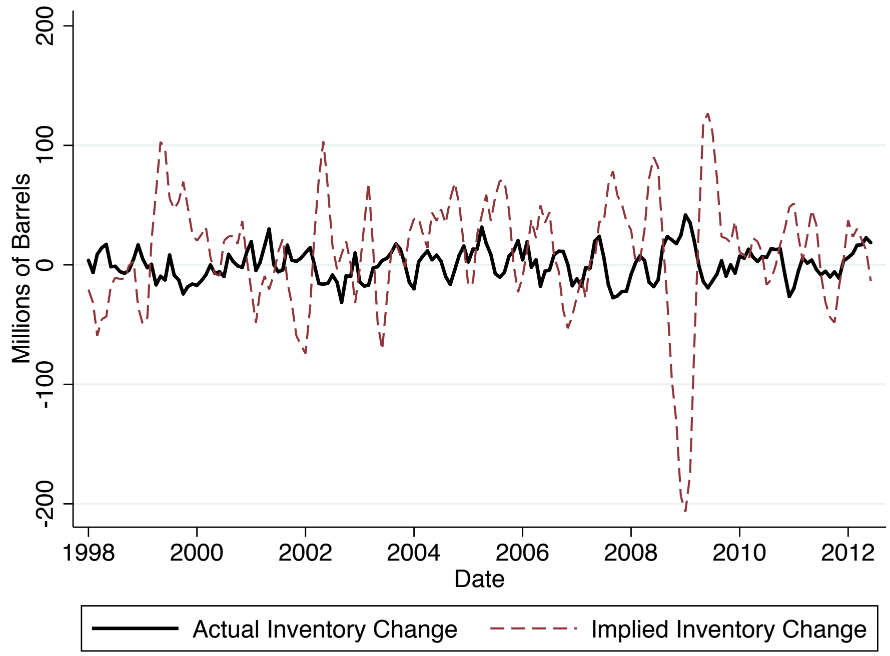

Figure 15: Actual Inventory Changes and Implied Inventory Changes if Prices Changes are Due to Speculation: Three-Month Periods

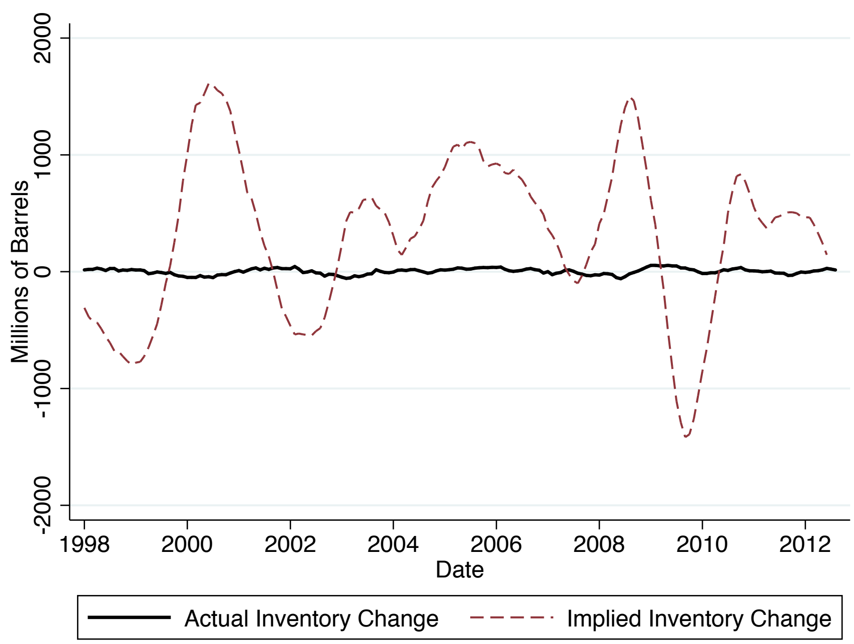

Figure 16: Actual Inventory Changes and Implied Inventory Changes if Prices Changes are Due to Speculation: Twelve-Month Periods

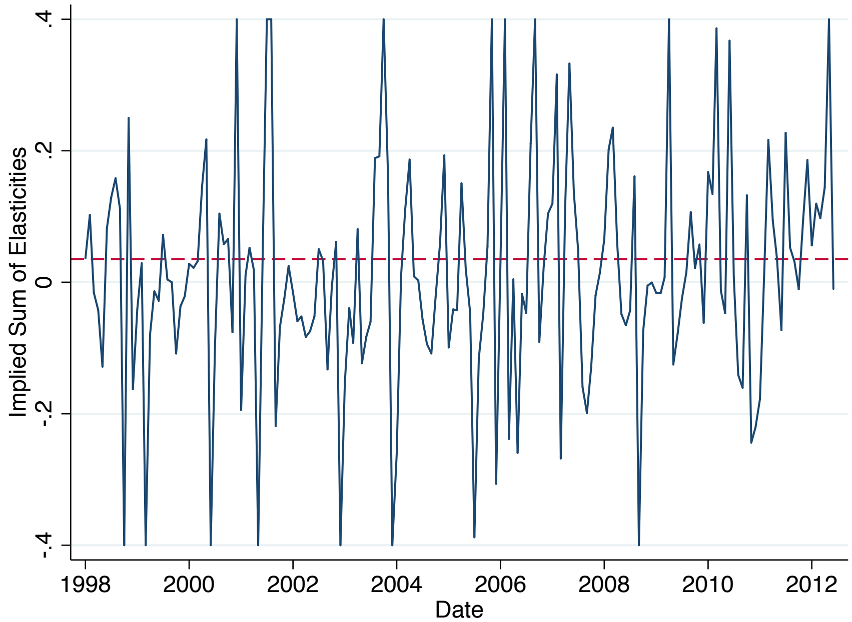

Figure 17: Implied Sum of Elasticities if Prices Changes are Due to Speculation: Three-Month Periods (red-dashed line represents the mean sum of elasticities)

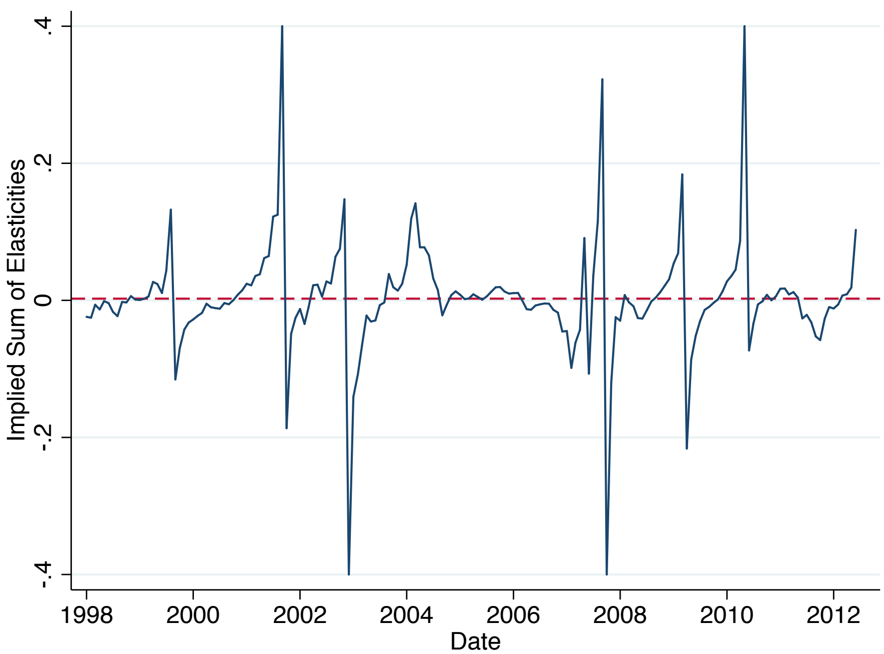

Figure 18: Implied Sum of Elasticities if Prices Changes are Due to Speculation: Twelve-Month Periods (red-dashed line represents the mean sum of elasticities)

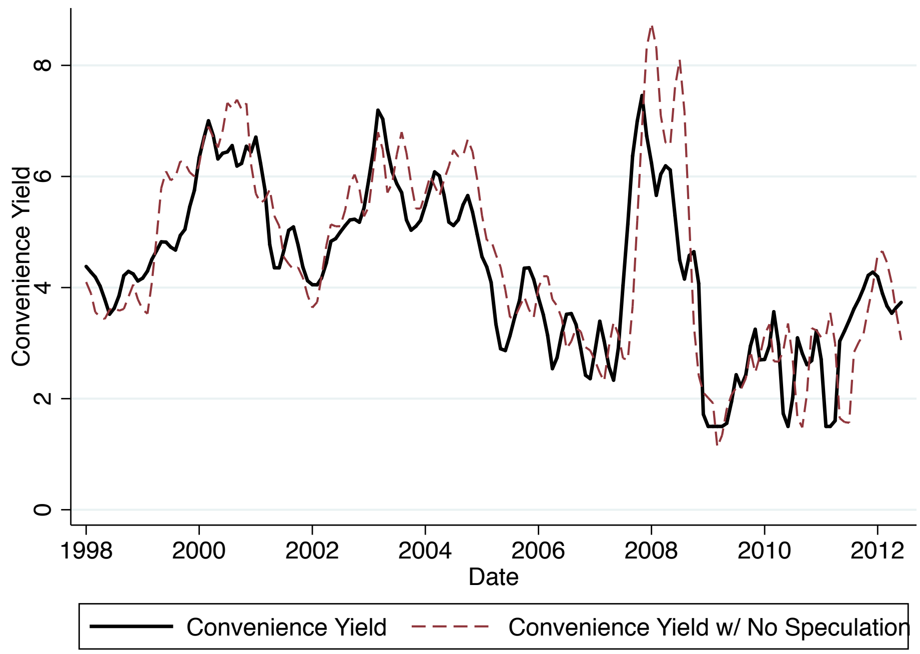

Figure 19: Actual Convenience Yields and Implied Convenience Yields with No Speculative Activity: Three-Month Periods

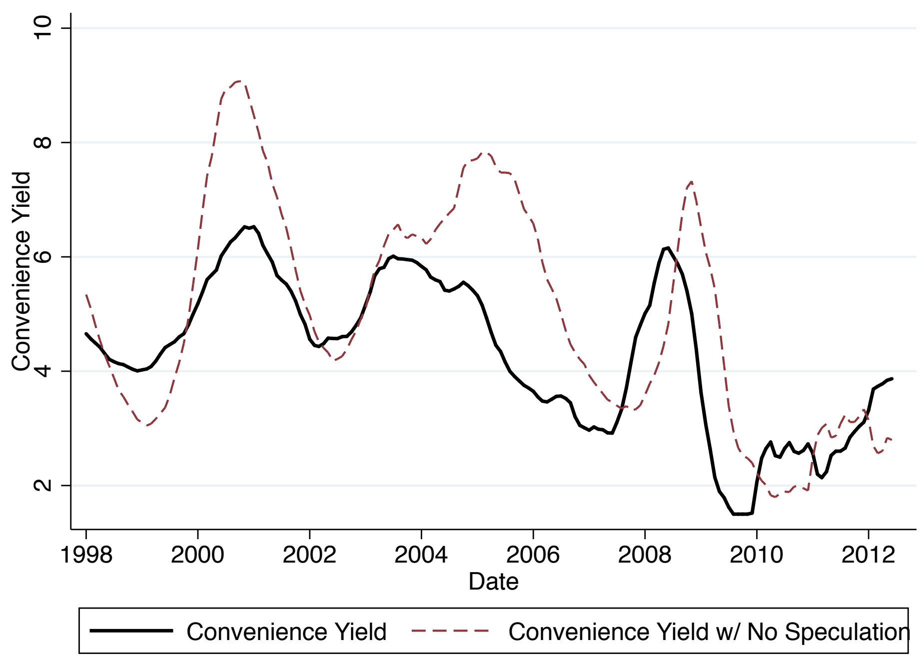

Figure 20: Actual Convenience Yields and Implied Convenience Yields with No Speculative Activity: Twelve-Month Periods

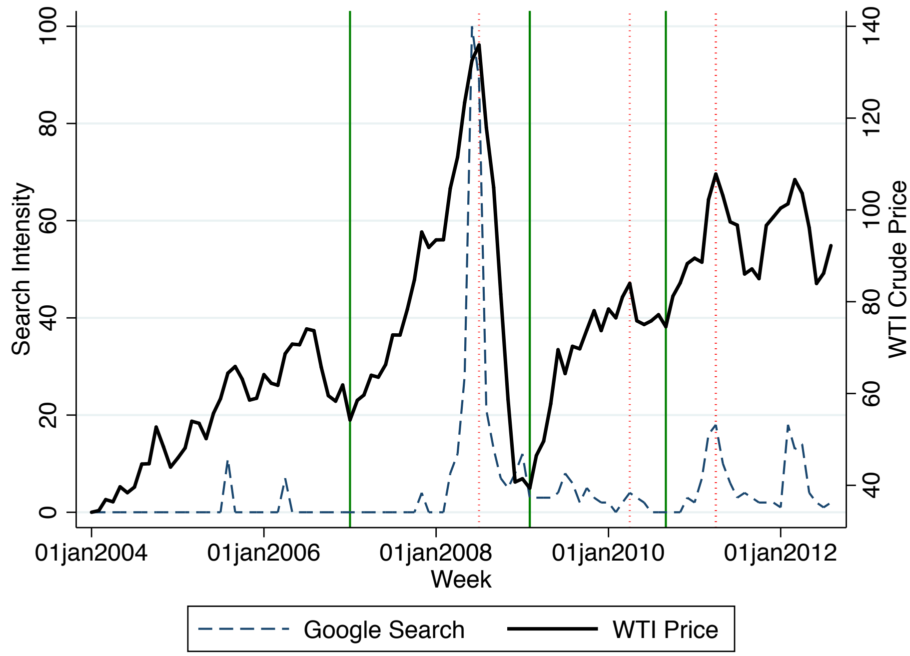

Figure 21: Monthly WTI Spot Prices and Google Search Intensity for "Oil Speculation"

9 Tables

Table 1: Estimation of Demand for Storage Curve

	(1) Three-Month	(2) Twelve-Month
$\Delta \ln(Spot)$	0.8221***	0.8337**
	(0.1328)	(0.3519)
$\Delta \ln(Stock)$	-1.1735*	-1.7479
	(0.6202)	(1.4742)
AR1	0.9749***	1.5671***
	(0.0508)	(0.0514)
AR2	-0.5866***	-0.6221***
	(0.0428)	(0.0526)
Constant	-0.0299	-0.0874
	(0.0224)	(0.0949)
σ	0.1614***	0.0596***
	(0.0060)	(0.0025)
Observations	162	162

Table 2: Epoch Analysis

Epoch	1 1/07-7/08	2 2/09-4/11	3 2/09-4/10	4 9/10-4/11
Beginning Price	$ 54.51	$ 39.09	$ 39.09	$ 75.24
Ending Price	$ 133.37	$ 109.53	$ 84.29	$ 109.53
Fundamentals-Only Price	$ 140.99	$ 109.90	$ 86.60	$ 106.56
Ending Inventories (Millions of Barrels)	295.23	366.54	363.27	366.54
Actual Inventory Build up	-0.54	8.98	5.70	3.71
Implied Inventory Build up*	261.47	125.59	96.45	43.76
Implied Sum of Elasticities*	-0.025	-0.001	-0.007	0.015
Ending Convenience Yield	$ 3.89	$ 3.04	$ 1.77	$ 3.04
Actual Change in Convenience	$ 0.37	$ 1.54	$ 0.27	$ 1.35
Fundamentals-Only Change in Convenience Yield	$ 5.94	$ 2.60	$ 1.68	$ 0.74

[1] * Consistent with speculation causing the price change.

A Alternative Demand and Supply Elasticity Assumptions

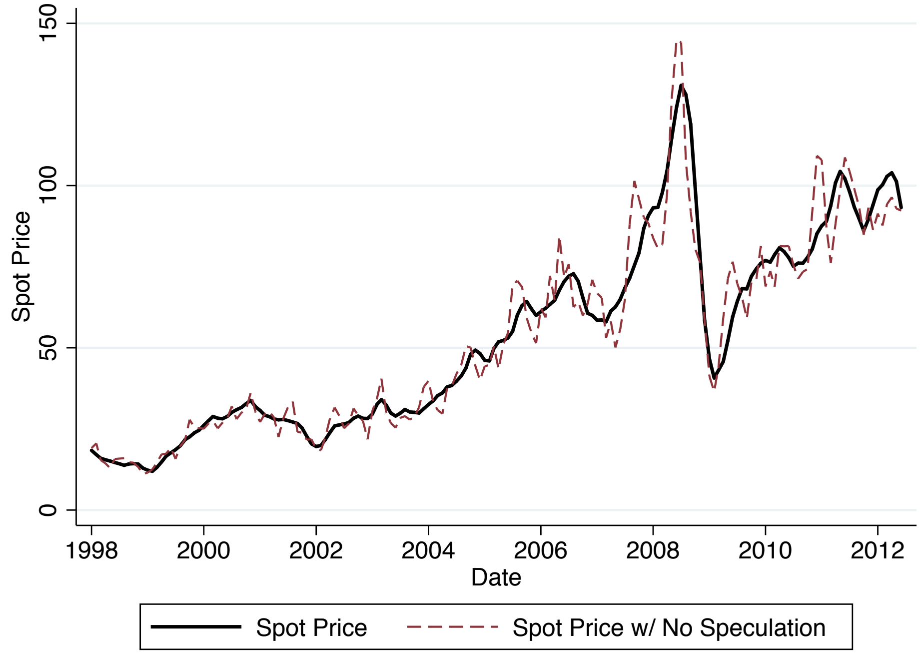

Figure 22: Actual Prices and Implied Prices with No Speculative Activity: Using Inventory Changes and Three-Month Periods using 0.05 for Supply and -0.05 for Demand

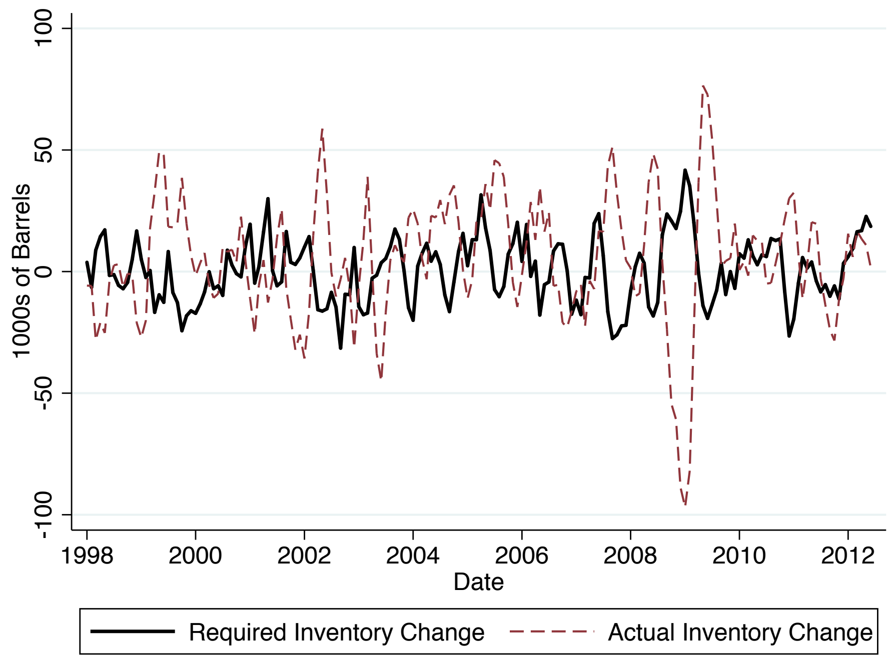

Figure 23: Actual Inventory Changes and Implied Inventory Changes if Prices Changes are Due to Speculation: Three-Month Periods using 0.05 for Supply and -0.05 for Demand

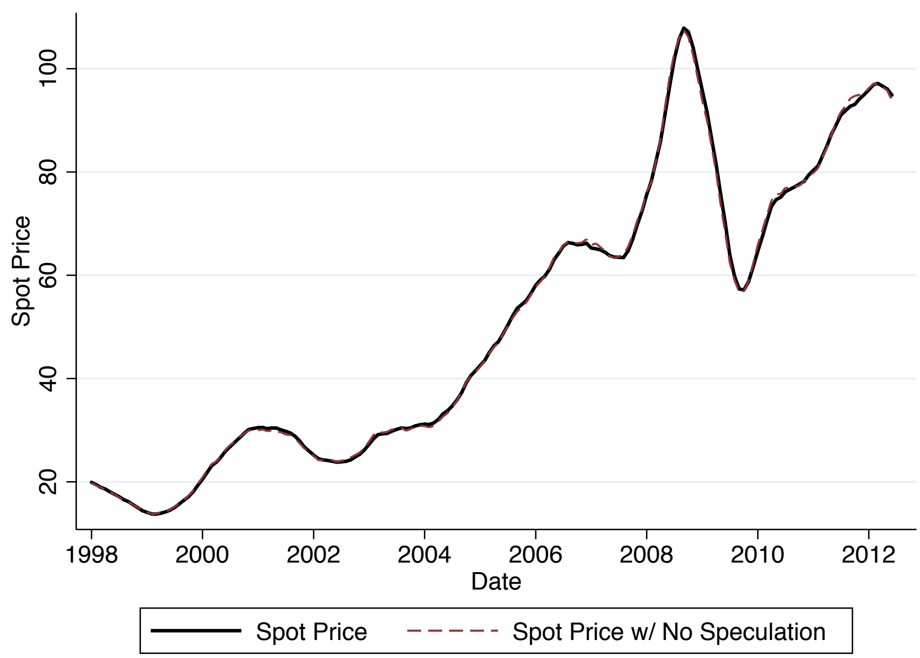

Figure 24: Actual Prices and Implied Prices with No Speculative Activity: Using Inventory Changes and Twelve-Month Periods using 0.30 for Supply and -0.30 for Demand

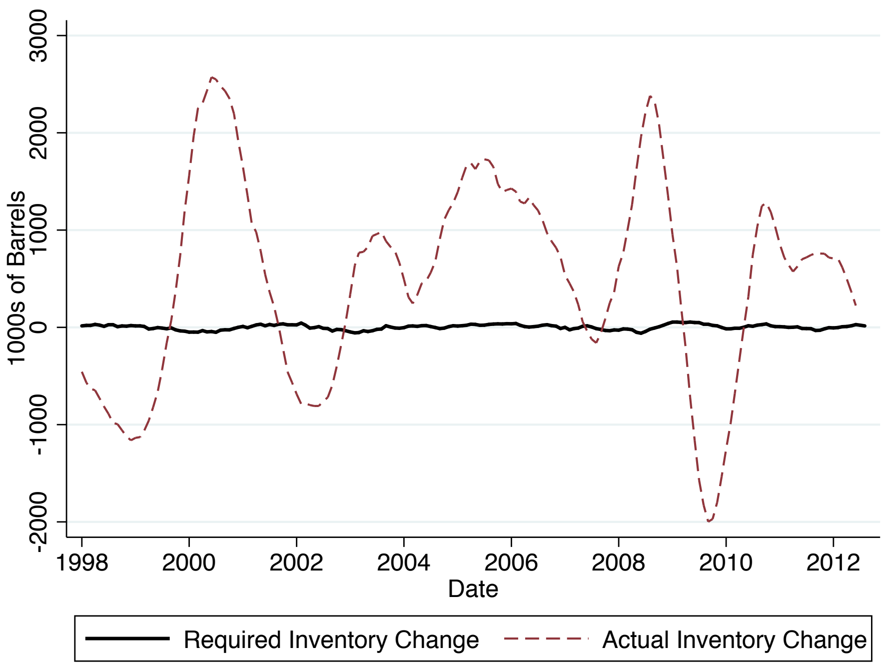

Figure 25: Actual Inventory Changes and Implied Inventory Changes if Prices Changes are Due to Speculation: Twelve-Month Periods using 0.30 for Supply and -0.30 for Demand