Accounting for Fuel Price Risk When Comparing Renewable to Gas-Fired Generation: The Role of Forward Natural Gas Prices

*Mark Bolinger, Ryan Wiser, and William Golove**

Unlike natural gas-fired generation, renewable generation (e.g., from wind, solar, and geothermal power) is largely immune to fuel price risk. If ratepayers are rational and value long-term price stability, then – contrary to common practice – any comparison of the levelized cost of renewable to gas-fired generation should be based on a hedged gas price input, rather than an uncertain gas price forecast. This paper compares natural gas prices that can be locked in through futures, swaps, and physical supply contracts to contemporaneous long-term forecasts of spot gas prices. We find that from 2000 -2003, forward gas prices for terms of 2-10 years have been considerably higher than most contemporaneous long-term gas price forecasts. This difference is striking, and implies that comparisons between renewable and gas-fired generation based on these forecasts over this period have arguably yielded results that are biased in favor of gas-fired generation.

This article, which is adapted from a longer Berkeley Lab report that can be accessed at http://eetd.lbl.gov/ea/EMS/reports/53587.pdf, was funded by the Assistant Secretary of Energy Efficiency and Renewable Energy of the U.S. Department of Energy under Contract No. DE-ACO3-76SF00098.

*Lawrence Berkeley National Laboratory, MS 90-4000, One Cyclotron Road, Berkeley, CA 94720, USA. E-mail: MABolinger@lbl.gov.

INTRODUCTION

The cost of generating electricity from renewable resources has declined dramatically over the past decade. As revealed by recent long-term power purchase agreements, the cost of wind power in particular now ranges from 2.5¢-4¢/kWh (including the impact of federal tax incentives) throughout much of the United States. At such levels, wind power is potentially competitive with new gas-fired generation, and utilities, regulators, and resource planners are beginning to compare wind to gas-fired generation on the basis of economics alone. A critical component of any such comparison, however, should be an assessment and quantification of the relative risks associated with renewable energy – both positive (e.g., reduced environmental compliance risk) and negative (e.g., resource variability).[1] This paper analyzes just one such risk: the ability of renewable generation to mitigate natural gas fuel price risk.

Against the backdrop of increasing – and increasingly volatile – natural gas prices, renewable energy resources, which by their nature are immune to natural gas fuel price risk, provide a potential economic benefit. Unlike many contracts for natural gas-fired electricity generation, renewable generation is typically sold under *fixed-price* contracts.[2] Assuming that electricity consumers are rational (and therefore prefer the less risky of two otherwise identical expected cash flows), a retail electricity supplier that is looking to expand its resource portfolio (or a policymaker interested in evaluating different resource options) should compare the cost of fixed-price renewable generation to the *hedged* or *guaranteed* cost of natural gas-fired generation, rather than to *projected* costs based on *uncertain* gas price forecasts.[3] Nonetheless, utilities and others often compare the costs of renewable to gas-fired generation using as their fuel price input long-term gas price forecasts that are inherently uncertain, rather than long-term natural gas forward prices that can actually be locked in. This practice raises the critical question of how these two price streams – i.e., forwards and forecasts – compare.

Building on earlier work relating to fuel price risk (e.g., Awerbuch 1993, 1994, 2003; Kahn and Stoft 1993; Humphreys and McClain 2002), in this paper we compare prices that can be locked in through long-term traditional gas-based hedging instruments (e.g., futures, swaps, and fixed-price physical supply contracts) to contemporaneous forecasts of spot natural gas prices, with the purpose of identifying any systematic differences between the two. Although our data are quite limited, we find that from November 2000 through 2003, forward gas prices for terms of 2-10 years have been considerably higher than most natural gas spot price forecasts. This difference is striking, and implies that resource planning and modeling exercises based on these forecasts over this period have yielded results that are biased (with respect to fuel price risk) in favor of gas-fired generation.

---

[1] For example, Bachrach et al. (2003) catalog the relative risks of renewable and gas-fired generation, and how those risks are treated in long-term electricity contracts. Hoff (1997) provides an analytical framework for evaluating a number of different risk types, while Cavanagh et al. (1993) and Repetto and Henderson (2003) specifically examine environmental compliance risk in the electricity industry. Brooks et al. (2003), Electrotek (2003), Dragoon (2003), and Hirst (2002) each estimate the cost of integrating significant amounts of intermittent wind power into different utility grid systems.

[2] While our analysis and results are presented in the context of comparing renewable to gas-fired generation, they are equally applicable to comparisons of other stable-priced forms of generation (e.g., coal or nuclear power) or demand reduction (e.g., energy efficiency) to variable-price gas-fired generation.

[3] Again, although this article focuses exclusively on fuel price risk, the cost of fuel (and its impact on total generation costs) is only one of many important considerations involved in any resource comparison. For example, the relative dispatchability of generating resources – regardless of levelized costs – may be of prime importance in some instances.

Although from a policy and analytic standpoint these empirical results are sufficiently interesting on their own, we spend the majority of this article examining several potential explanations for our empirical findings. Specifically, within the context of the extensive (though inconclusive) treatment of this subject within the literature, we first discuss the possibility that hedging is not costless.[4] We then consider the possibility that the spot price forecasts we employ have been biased downwards, or that other data or sampling problems are driving the empirical premiums. We find none of these potential explanations to be either fully satisfying or easily refutable. Regardless of the explanation, however, the implications of our analysis remain the same: when comparing fixed-price renewable to variable-price gas-fired generation, forward gas prices are the most appropriate fuel price input if long-term price stability is valued.

BACKGROUND

For better or worse, natural gas has become the fuel of choice for new power plants being built across the United States. According to the Energy Information Administration (EIA), natural gas combined-cycle and combustion turbine power plants accounted for 96% (138 GW out of 144 GW total) of the total generating capacity added in the U.S. between 1999 and 2002 (EIA 2003). Looking ahead, the EIA expects that gas-fired technology will account for 61% of the 355 GW of new generating capacity projected to come on-line in the US through 2025, increasing the nationwide market share of gas-fired generation from 18% in 2002 to 22% in 2025 (EIA 2004). While the data are specific to the US, natural gas-fired generation is making similar advances in other countries as well.

With increasing competition for dwindling domestic natural gas supplies, it is likely that gas prices will be as or more volatile than they have been in the past. Figure 1 shows natural gas futures prices (from the "first-nearby" futures contract – i.e., the contract that is closest to expiration at each point in time) on a daily basis going back to the inception of trading on the New York Mercantile Exchange (NYMEX) in April 1990. While the "twin peaks" of December 2000 and February 2003 clearly dominate the graph and make the rest of the price history look comparatively tame, many of the "lesser" price spikes during the early 1990s represent doublings or more in price.

[4] This potential explanation – which boils down to whether or not forward prices are unbiased estimators of future spot prices – has been widely debated in the literature and is interesting in its own right, independent of its implications for resource comparisons. Accordingly, we devote a disproportionate amount of space to this topic.

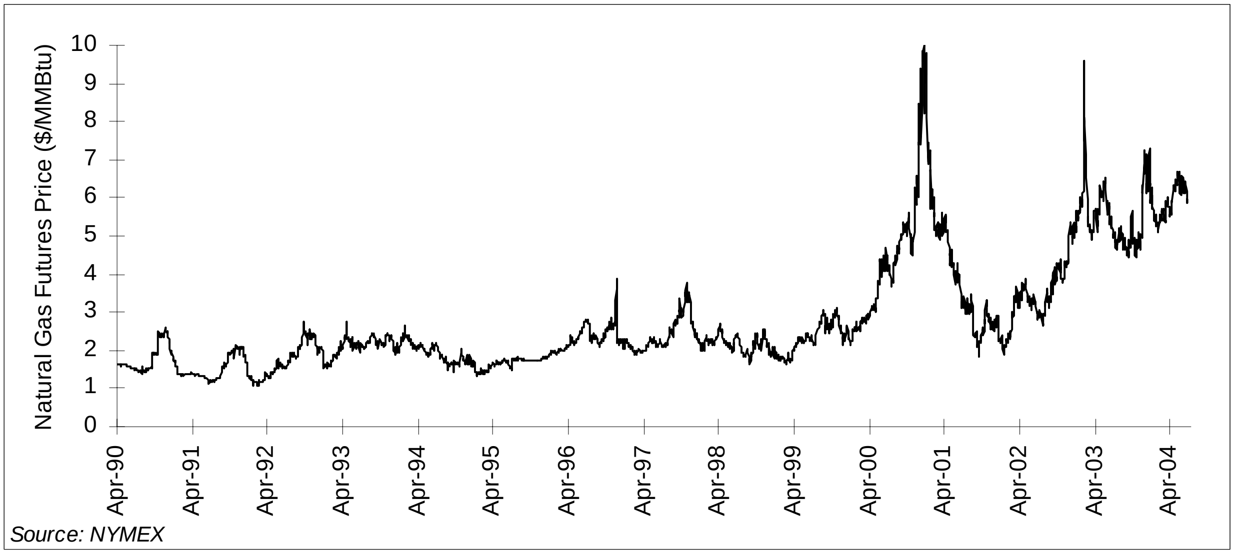

Figure 1. NYMEX Natural Gas Futures Prices (First-Nearby Contract)

The path of natural gas prices depicted in Figure 1 is particularly troubling considering that at mid-2004 spot prices of around $6/MMBtu, the cost of natural gas accounts for *more than 70%* the levelized cost of energy from a new combined cycle gas turbine, and *more than 90%* of its operating costs (EIA 2004). Moreover, gas-fired plants are often the marginal units that set the market-clearing price for *all* generators in a competitive wholesale market, allowing natural gas price volatility to flow directly through to wholesale electricity price volatility. Clearly, gas price volatility is a major contributor to wholesale electricity price volatility, and poses a major risk to both buyers and sellers of gas-fired generation.

Electricity can be bought and sold either (1) on the spot market, (2) through contracts that are indexed to (i.e., vary with) the price of the fuel input,[5] (3) through tolling agreements (whereby the power purchaser delivers fuel to the generator and takes delivery of the resulting power that is produced, having effectively "rented" the use of the generation plant), or (4) through fixed-price contracts. Natural gas-fired generation is commonly sold through all four of these contract types (Bachrach et al. 2003), with gas price risk falling on the power purchaser in the first three, and the generator in the final type. Renewable generation, on the other hand, is typically sold through long-term fixed-price contracts (perhaps indexed to inflation), and generally imposes no gas price risk on either the buyer or seller. In order to achieve a fuel price risk profile similar to that of fixed-price renewable generation, either the buyer (under spot, indexed, and tolling contracts) or seller (under fixed-price contracts) of gas-fired generation must hedge away natural gas price risk.

To hedge natural gas price risk, a retail electricity supplier can either purchase renewable generation (which is immune to gas price risk), purchase variable-priced gas-fired generation and choose among a number of gas-based financial and physical hedging instruments, or – if available – purchase *fixed-price* gas-fired generation (in which case the *generator* may wish to

[5] Though fuel price indices are most common, electricity contracts may instead be linked to other price indices. For example, an aluminum smelter wishing to stabilize its profit margin may seek an electricity contract that is indexed to the price of aluminum (i.e., the company's output). For the purposes of this article, we will assume that indexed contracts are linked to the price of the fuel input, natural gas.

hedge using financial or physical instruments).[6] Financial gas-based hedges include futures (or, more generically, forwards), swaps, options on futures, or some combination or derivation thereof (e.g., collars). Physical hedges include long-term fixed-price gas supply contracts and natural gas storage.

To fairly evaluate fixed-price renewable and variable-price gas-fired contracts on an apples-to-apples basis with respect to fuel price risk, we must look to those gas-based hedging instruments that provide a hedged payout pattern similar to that of renewables – i.e., flat and symmetrical, immune to both gas price increases and decreases. Bolinger et al. (2003) demonstrate that such instruments include gas futures, swaps, and fixed-price physical supply contracts, but not options or storage. The prices that can be locked in through these instruments are therefore the appropriate fuel price input to modeling and planning studies that compare – either explicitly or implicitly – renewable to gas-fired generation.

As shown in Bolinger et al. (2003), however, utilities and others conducting such analyses tend to rely primarily on uncertain long-term forecasts of spot natural gas prices, rather than on prices that can be locked in through futures, swap, or fixed-price physical supply contracts (i.e., "forward prices"). This practice raises a critical question: how do the prices contained in uncertain long-term gas price forecasts compare to actual forward prices that can be locked in?

If the two price streams closely match one another, then one might conclude that forecast-based resource acquisition, planning, and modeling exercises are implicitly accounting for the price stability benefits of renewable relative to gas-fired generation, approximating an apples-to-apples comparison. If, however, forward prices systematically differ from long-term spot price forecasts (e.g., if there is a cost to hedging, or if the forecasts are out of tune with market expectations), then the use of such forecasts in resource acquisition, planning, and modeling exercises will arguably yield results that are either (1) biased in favor of renewable generation (if forwards < forecasts), or (2) biased in favor of natural gas-fired generation (if forwards > forecasts).

EMPIRICAL FINDINGS OF A PREMIUM

EIA Forecast Comparisons

We first investigate the extent to which natural gas forward prices match price forecasts by comparing the prices of futures, swap, and fixed-price physical gas supply contracts to "reference case" gas price forecasts from the EIA. The data necessary to conduct our analysis are simple: a forward gas price and a gas price forecast, ideally generated at the same time. While long-term gas price forecasts are relatively easy to come by (e.g., the EIA forecasts are publicly available and updated every year), long-term forward prices – and in particular those of sufficient duration to be of interest – present a greater challenge. The NYMEX gas futures strip extends out six years (but is liquid for much less than that) – a period that is only about one-third as long as the typical term of a power purchase agreement for renewable energy (which commonly extend 15-25 years). Forward gas contracts in excess of 6 years are traded infrequently and, when traded, are traded bilaterally "over the counter" (i.e., not on an organized

[6] Similarly, as noted earlier, investments in energy efficiency (e.g., through demand-side management), or even coal or nuclear power (with fuel costs that are quite stable compared to the cost of natural gas), may provide an equivalent natural gas price hedge. For example, Humphreys and McClain (1998) use modern portfolio theory to demonstrate how a shift toward coal-fired generation could reduce wholesale electricity price volatility.

exchange), and are therefore rarely documented in the public domain. In addition, we must further restrict our sample to those forward prices that were traded or posted at roughly the same time as the generation of a long-term gas price forecast (this timing issue will be discussed later).

Thus, despite efforts to obtain a larger sample, our analysis is limited to comparisons based on publicly available data that we collected from November 2000-2003, and for forward price terms not exceeding 10 years. Specifically, our limited sample of natural gas forward contracts and price forecasts includes:

- 2-, 5-, and 10-year natural gas swaps offered by Enron in early November 2000 and 2001, compared to reference case natural gas price forecasts from the Energy Information Administration's (EIA) *Annual Energy Outlook 2001* and *2002*, respectively;
- the six-year NYMEX natural gas futures strip (averaged each year) from early November 2002, compared to the reference case gas price forecast contained in *Annual Energy Outlook 2003*;
- a seven-year physical gas supply contract between Williams and the California Department of Water Resources (DWR) signed in early November 2002, again compared to the reference case gas price forecast contained in *Annual Energy Outlook 2003*, and;
- the six-year NYMEX natural gas futures strip from mid-October 2003, compared to the reference case gas price forecast contained in *Annual Energy Outlook 2004*.

In each case, we evaluate these forward prices against the EIA's reference case forecast of natural gas prices delivered to electricity generators (in nominal dollars[7]), which is generated in the fall of each year and presented in the *Annual Energy Outlook* (AEO) series released towards the end of the year. Although the use of EIA reference case gas price forecasts for this purpose is somewhat controversial (as discussed later), we use them nonetheless because they are publicly available, have been widely vetted, and most importantly are commonly adopted by the EIA and others as a "base case" price scenario in policy evaluations, modeling exercises, and even resource acquisition decisions.

To make a proper comparison between the EIA spot price forecasts (which represent the average price of gas delivered to electricity generators nationwide) and the Enron swap and NYMEX futures prices (which are indexed or deliverable to the Henry Hub in Louisiana), each year we subtract \$0.38/MMBtu from the EIA spot price forecast to account for the average cost of transportation from the Henry Hub to electricity generators nationwide.[8] We make the same adjustment in order to place the EIA forecast on the same terms as the Williams/DWR contract, which will result in delivery of natural gas to the Southern California border (i.e., we assume that prices at the Southern California border will, on average, match prices at the Henry Hub).[9]

---

[7] Unless otherwise noted, all forecasts are expressed in *nominal* (as opposed to *real* or constant dollar) terms in order to be comparable to forward prices, which are also expressed in nominal terms. For example, we inflate the EIA gas price forecasts – which are expressed in real terms – to nominal terms using the EIA's own inflation projections.

[8] This adjustment is necessary because the EIA does not provide a forecast of Henry Hub spot prices. To estimate the average cost of transportation from the Henry Hub to electricity generators nationwide, we compared historic first-nearby NYMEX futures prices (which are indexed to the Henry Hub) to delivered (to electricity generators) spot prices on a monthly basis from April 1990 through December 2002 (n=153 months).[8] This comparison revealed an average transportation margin of \$0.38/MMBtu, with a 95% confidence interval that ranges from \$0.33 to \$0.43/MMBtu.

[9] The locational basis differential between Southern California and the Henry Hub has varied over time in both sign and magnitude. While it was strongly positive (i.e., prices in Southern California were sharply higher than at the Henry Hub) during the California energy crisis of 2000 and 2001, the basis has been consistently flat or slightly negative since late 2001. As such, we simply assume here that the basis between Southern California and the Henry Hub will be flat (i.e., zero) over the duration of the Williams/DWR contract. This assumption appears to be

Each of these comparisons reveals that forward natural gas prices have traded above EIA reference case price forecasts during this four-year period, sometimes significantly so. Figure 2 consolidates the resulting levelized premiums (in terms of $/MMBtu and ¢/kWh, assuming a heat rate of 7,000 Btu/kWh) from each of these comparisons into a single graph.[10] As shown, the magnitude of the empirically derived premiums varies somewhat from year to year, contract to contract, and by contract term, ranging from $0.5-$0.8/MMBtu, or 0.4-0.6¢/kWh assuming a highly-efficient gas-fired power plant.

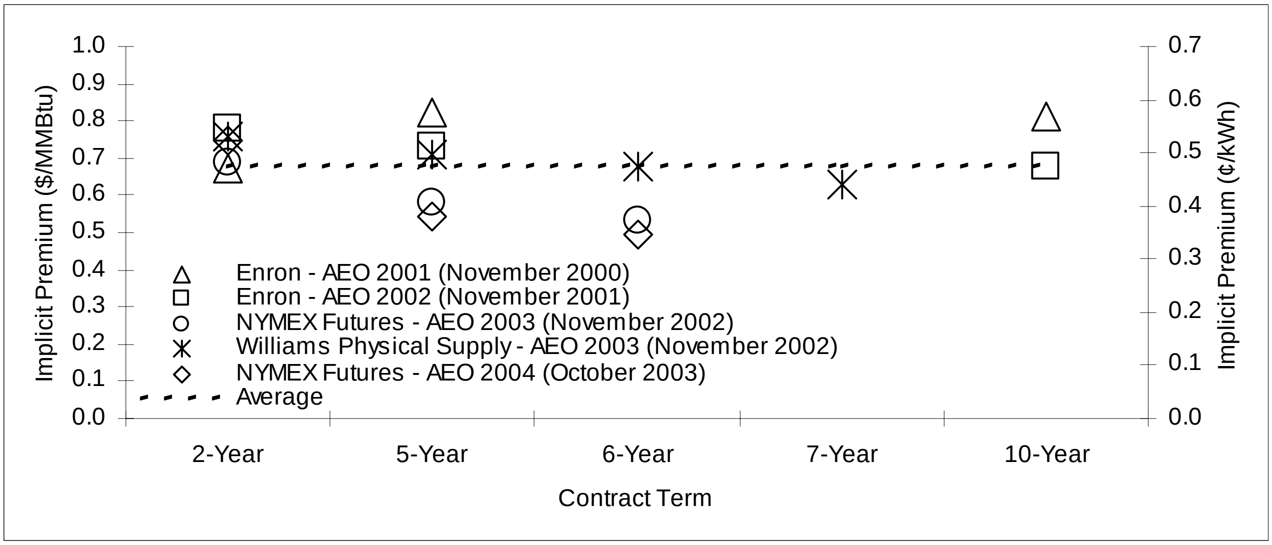

Figure 2. Implied Premiums in $/MMBtu and ¢/kWh (assuming 7,000 Btu/kWh)

While the relatively tight range of premiums is somewhat remarkable, the limitations in our data sample mean that one cannot easily extrapolate these findings beyond the four periods for which we have data, or to contract terms longer than those examined. It is, however, at least apparent that utilities and others who have conducted resource planning and modeling studies based on EIA reference case gas price forecasts from November 2000-2003 have produced "biased" results that favor variable-price gas-fired over fixed-price renewable generation, potentially to the tune of ~0.4-0.6¢/kWh on a levelized basis. This is because an apples-to-apples comparison of renewable to gas-fired generation must be based on fuel prices that can be locked in with certainty,[11] which, as shown in Figure 2, have traded at a premium over uncertain fuel price forecasts during this period.

---

consistent with, and perhaps even conservative relative to, historical post-crisis basis values (including at the time the contract was signed), as well as prices that can currently be locked in for the next 30 months through the forward market. Thus, we compare the Williams/DWR contract prices to the same adjusted EIA forecast – i.e., reference case prices delivered to generators, less $0.38/MMBtu – as we used for the Enron and NYMEX comparisons.

[10] We derived the premiums by levelizing the first 2, 5, 6, 7, and 10 years of the EIA forecasts (using a discount rate of 10%) and subtracting the resulting levelized forecast price from the corresponding forward prices. Because in this case levelizing involves taking the present value of a price stream and amortizing it forward *at the same discount rate*, the calculation is relatively insensitive to the level of the discount rate chosen. For example, using a discount rate of 5% barely changed the results.

[11] Of course, even gas prices that are locked in through long-term contracts still carry default risk – for example, the risk that gas suppliers will default on their obligation to deliver gas in the event of a sustained price increase or some other unforeseen circumstance. Renewable generators, which are typically not dependent on a fuel input supplied by others, need not bear this particular risk.

Other Forecasts

The EIA's forecasts are by no means the only long-term gas price forecasts available to market participants. In order to assess how the premiums presented in Figure 2 would change had we compared natural gas forward prices to some forecast other than the EIA's, we reviewed a number of other long-term gas price forecasts, sourced from the EIA's own forecast comparisons (contained in each year's AEO), as well as from various utility integrated resource plans.[12] As shown in Bolinger et al. (2003), with few exceptions, the EIA reference case forecast has generally been higher – and often substantially so – than most other forecasts generated from 2000-2003 and used by utilities and others. These findings suggest that the premiums presented in Figure 2 would be *even larger* when comparing forward prices to some of the other commonly used gas price forecasts. Utilities and others that have used these other (i.e., non-EIA) forecasts to compare fixed-price renewable to variable-price gas-fired generation from 2000-2003 have therefore arguably obtained results that are *even more* "biased" (with respect to fuel price risk) in favor of gas-fired generation than those resulting from EIA-based reference case comparisons.

For example, had we compared the November 2001 Enron 10-year natural gas swap to the gas price forecast contained in Idaho Power's *2002 Integrated Resource Plan*,[13] we would have observed a 10-year levelized premium of $1.29/MMBtu – i.e., *nearly twice as large* as the $0.68/MMBtu benchmarked against the EIA reference case forecast. This translates to a 0.9¢/kWh premium at an aggressive heat rate of 7,000 Btu/kWh; had Idaho Power used forward rather than forecast data in its integrated resource plan, comparisons between renewable and gas-fired generation may have looked significantly different. Though not as large in magnitude as the Idaho Power example, most other forecast comparisons yielded similar results (for details, see Bolinger et al. 2003).

POTENTIAL EXPLANATIONS FOR EMPIRICAL PREMIUMS

The differences between natural gas forwards and gas price forecasts from November 2000-2003 revealed in the previous section, and the implications such differences hold for resource comparisons, are significant. We are keenly aware, however, that these empirical results are based on limited historical data over a 4-year period characterized by extreme price volatility. In considering whether it is possible to extrapolate our findings, or the implications thereof, into the future, it is important to try to understand some of the causes of our empirical findings. At least three explanations may partially or wholly account for the sizable differences between natural gas forwards and forecasts: (1) hedging is not costless, and the observed premiums represent the cost of hedging; (2) the forecasts are out of tune with market expectations, so the observed empirical premiums reflect forecast bias; and (3) other data or sampling issues are driving the premium. Below we assess each of these three broad possibilities.

[12] Specifically, we examined gas price forecasts contained in the most recent integrated resource plans from Avista (DRI-WEFA), Idaho Power (WEFA, November 2001), Pacificorp (PIRA, March 2002), Portland General Electric (EIA, November 2001), and Puget Sound Energy (PIRA, January 2003).

[13] The Idaho Power forecast is based on a November 2001 WEFA forecast, and is therefore compared to the gas price forecast contained in AEO 2002 (for the Mountain Region).

Hedging May Not Be Costless

This potential explanation may have two components or causes: (1) forward prices are upwardly biased predictors of future spot prices, and (2) transaction costs are significant. The first notion, that forward prices are biased predictors of future spot prices, has been extensively debated in the literature ever since Keynes first introduced the idea of *normal backwardation* in his 1930 *Treatise on Money*.

Specifically, Keynes (1930) argued that *hedgers* who use futures markets to mitigate commodity price risk must compensate *speculators* for the "insurance" that they provide.[14] A futures market that is dominated by short hedgers – i.e., a market with more natural sellers than buyers of future contracts – results in what is known as *positive net hedging pressure*, and requires speculators to step in and purchase the excess futures contracts that are for sale in order to clear the market. To entice speculators to provide this "insurance" or liquidity, short hedgers must, according to Keynes (1930), be willing to sell futures contracts to speculators at prices that are *lower* than the expected spot price, thereby enabling speculators to earn a positive return simply by buying the contracts and holding them to expiration (when spot and futures prices presumably converge). The opposite occurs in a market dominated by long hedgers, resulting in *negative net hedging pressure*: to clear the market, speculators must *sell* futures contracts, and are compensated in the form of futures prices that are *higher* than the expected spot price, again enabling them to earn a positive return simply by holding the short position to expiration.

Over the years, a number of studies have attempted – with mixed results – to empirically confirm or refute the existence of risk premiums in futures prices by examining the returns to speculators.[15] Much of this work, however, has been based on a strict application of Keynes' theory, which, among other things, assumes that hedgers – characterized mainly as producers who are natural short hedgers – will be net short futures and that speculators will therefore be net long futures. As explained above, this results in *positive net hedging pressure*, and futures contract prices that should, in theory, be below expected spot market prices (i.e., Keynes' *normal backwardation*). Under this constraint, one need only test for a positive return to holding futures contracts to see whether speculators have earned a risk premium. This is the general approach taken by Telser (1958) and Gray (1961), who find no evidence of a risk premium in agricultural commodities such as corn, wheat, and cotton.

As noted by Chang (1985), however, a number of researchers have relaxed this constraint in recognition that both producers *and* consumers hedge, and that net hedging pressure may therefore not always be positive (i.e., in aggregate, short hedges may not always outnumber long hedges). Such researchers have generally incorporated information about the aggregate net

---

[14] *Hedgers* can be thought of as commodity producers (e.g., farmers) or consumers (e.g., manufacturers) who have a natural underlying position in a commodity that can be hedged by selling or buying futures contracts. A "short" hedge transaction is one in which the hedger currently owns the underlying commodity and will be selling it in the future (e.g., post-harvest), and so today sells futures contracts to lock in the future sales price of the commodity and thereby secure a profit margin. Conversely, a "long" hedge transaction is one in which the hedger does not currently own the commodity but will be buying it in the future (e.g., as orders for finished goods increase), and so today buys futures contracts to lock in the future purchase price of the commodity and thereby secure a profit margin. In contrast to hedgers, *speculators* have no natural underlying position in the commodity, and trade in the futures market purely to make a profit. If a short hedger cannot find a long hedger to trade with, a speculator will step in to buy the short hedger's futures contracts, in the hopes of being able to re-sell them at a profit. In this way, speculators provide liquidity to the market.

[15] Due to the historical nature of the debate, which began in the 1930s, most of these studies deal primarily with agricultural commodities. Studies focusing specifically on natural gas and wholesale electricity futures markets, which did not come into being until the 1990s, are discussed later.

position of hedgers into their analysis, checking for positive futures returns when net hedging pressure is positive, and negative futures returns when net hedging pressure is negative. Though results remain inconclusive, the relaxation of the assumption that speculators are always net long has generally led to results, as reported in Chang (1985) and Hull (1999), that are more supportive of Keynes' notion of a risk premium (positive or negative, depending on whether net hedging pressure is negative or positive, respectively) embedded in futures prices. For example, Houthaker (1957), Cootner (1960), Carter et al. (1983), Chang (1985), and de Roon et al. (2000) all relax the assumption that speculators will be net long futures and find empirical evidence of risk premiums embedded in futures prices.[16]

Still other researchers have searched for risk premiums in commodity futures prices from within the framework of the Capital Asset Pricing Model (CAPM). Under CAPM, it is not the variability of prices per se, but rather the correlation of price variability with changes in total wealth that impact the futures premium. Dusak (1973) examined wheat, corn, and soybean futures within the CAPM framework, and found no evidence of either non-zero systematic risk or non-zero futures returns.[17] Hirshleifer (1988), meanwhile, developed a model whereby the risk premium consists of two terms – a systematic risk term (i.e., related to CAPM) and a term due solely to residual risk (i.e., related to net hedging pressure).

Below we examine in more detail the possibility that either net hedging pressure or CAPM-related systematic price risk is responsible for the empirical premiums observed in the natural gas market from November 2000-2003, as presented earlier. We then consider the possibility – still within the context that there is a cost to hedging – that high transaction costs are driving at least a portion of the observed empirical premiums.

Net Hedging Pressure in the Natural Gas Market

A number of studies have examined the efficiency of the natural gas market specifically (i.e., whether gas futures prices are unbiased predictors of subsequent spot prices), and like the broader literature, have found mixed results. Analyzing just the first few years of the NYMEX natural gas futures market (which began in April 1990), Herbert (1993) finds the market to be inefficient (futures prices > realized spot prices, suggesting a positive risk premium). With a few more years of price history to examine, Walls (1995) finds the gas futures market to be generally efficient (futures prices = realized spot prices, suggesting no risk premium). More recently, Buchanan et al. (2001) find evidence of positive returns to large natural gas speculators, and Chinn et al. (2001) find that natural gas future prices are a biased and poor predictor of future spot prices.[18] Of course, all of these studies analyze short-term futures contracts, and it is not clear that their results can be extrapolated to the longer-dated futures of interest here.

To test whether the positive premiums (natural gas forward prices that are higher than

---

[16] We again note that none of these studies focus on energy commodities, though de Roon et al. (2000) do include crude and heating oil within a broad array of 20 commodities.

[17] Note that these findings do not rule out CAPM as a potentially useful tool for this purpose, since under CAPM, one would expect zero systematic risk to lead to no risk premium.

[18] In addition, though not directly related to natural gas, there has recently been some interesting work on this issue in the wholesale electricity markets. Bessembinder and Lemmon (2002) develop an equilibrium model that allows net hedging pressure to be either positive or negative, depending on both the variance and distribution (i.e., skewness) of wholesale electricity prices. In testing their model, they – along with Pirrong and Jermakyan (2001) and Longstaff and Wang (2003) – find evidence of significant positive risk premiums in the PJM day-ahead wholesale electricity markets (i.e., day-ahead forward prices were priced at significantly higher levels than realized spot prices).

contemporaneous forecasts of spot prices) observed earlier could be the direct result of negative net hedging pressure, Figure 3 depicts net hedging pressure in the natural gas market from 1999 through 2003.[19] Though largely positive over this period (suggesting that futures prices should generally be *lower* than expected spot prices), net hedging pressure has clearly varied quite a bit, and has been negative at times. Of particular interest to this study is the fact that during three of the four periods in which we sampled forward prices (November 2001, November 2002, and October 2003, as depicted by the shaded bars), net hedging pressure was either neutral or slightly negative (i.e., suggesting futures prices should match or be higher than expected spot prices). In the fourth sample period – November 2000 – net hedging pressure was clearly positive, though not nearly to the degree seen earlier in that and the previous year.

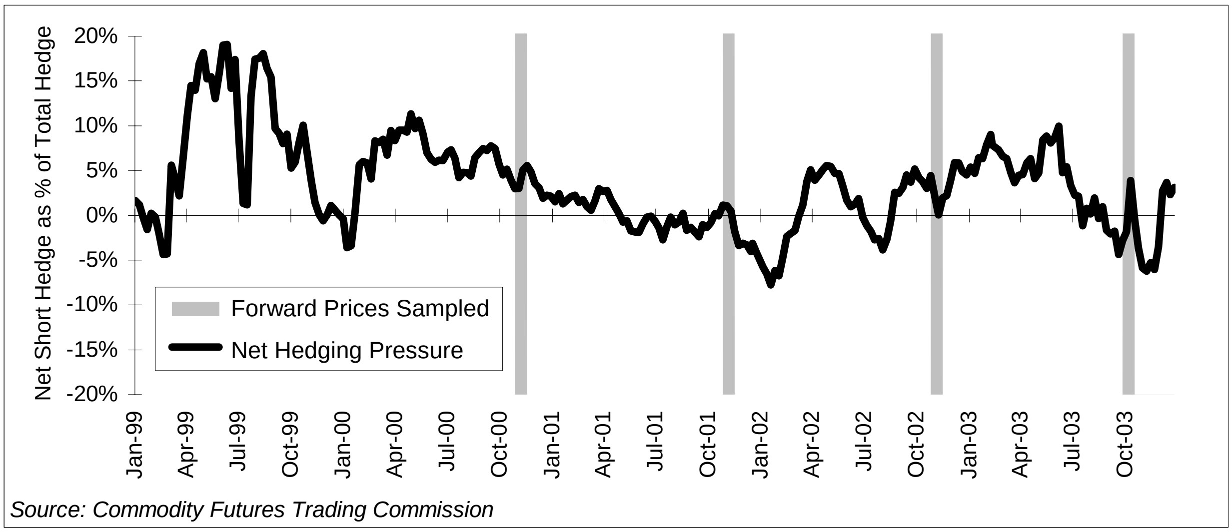

Figure 3. Net Hedging Pressure in the Natural Gas Futures Market, 1999-2003

Thus, while Figure 3 does not support the idea that negative net hedging pressure is directly responsible for the positive premiums observed earlier in this article, it does show a notable lack of *positive* net hedging pressure during our sample period. In other words, net hedging pressure appears to provide little support either for *or* against our specific findings of significant positive premiums in the natural gas market from November 2000 – 2003. More generally, however, the fact that net hedging pressure is not *uniformly* negative suggests that one should not expect – at least on the basis of net hedging pressure alone – such positive premiums to exist at all times.

Systematic Risk in Natural Gas Prices

Setting aside net hedging pressure and the returns of speculators, what if producers benefited from price volatility, while consumers were hurt by it? In this case, producers would require compensation (i.e., a premium) for being locked into long-term fixed-price contracts, and

[19] The data come from the weekly *Commitments of Traders* reports published by the Commodity Futures Trading Commission (CFTC). Traders are classified as "commercial" or "non-commercial" depending on whether their futures positions in a given commodity are used for hedging or speculative purposes, respectively. Net hedging pressure is defined as the difference between the number of outstanding short and long contracts divided by the sum of outstanding short and long contracts (among commercial traders). Commercial traders (i.e., hedgers) regularly account for 60-75% of the open interest in the natural gas futures market.

consumers would be willing to pay such compensation. Economic theory provides some support for this very scenario in the form of the Capital Asset Pricing Model (CAPM).[20]

While CAPM was originally derived as a financial tool to be applied to investment portfolios, its basic tenet – that an asset's risk depends on the correlation of its revenue stream with variability in the asset-holder's overall wealth – can be applied much more broadly, for example in evaluating investments in physical assets such as power plants (Awerbuch 1993, 1994; Kahn & Stoft 1993). Specifically, in the context of natural gas-fired generation, one can think about the correlation between a gas consumer's overall wealth (as proxied by the economy or, more specifically, the stock market) and natural gas prices. If gas prices, and therefore consumer expenditures on gas, rise as the stock market declines (e.g., because rising gas prices hurt the economy), then natural gas is said to have a negative "beta,"[21] and is risky to gas consumers and beneficial to gas producers. In other words, at the same time as gas consumers and producers feel the pinch of a weak stock market, expenditures on natural gas also rise, compounding overall wealth depletion among consumers while providing some consolation to producers.[22]

In this specific case, where gas with a negative beta is risky to consumers and beneficial to producers, consumers have an incentive to hedge natural gas price risk, while producers do not. Intuitively, it follows that even if both consumers and producers share identical expectations of future spot gas prices, producers will still require – and consumers will be willing to pay – a premium over expected spot prices in order to lock in those prices today. Both Pindyck (2001) and Hull (1999) mathematically demonstrate this to be the case: when beta is negative, futures prices should, at least theoretically, trade at a premium to expected spot prices.

Thus, if the beta of natural gas is indeed negative, this theory might explain our empirical observations of a positive premium embedded in contract prices. One can test this notion by regressing natural gas price changes against stock market returns. Below we survey past efforts to quantify the beta of natural gas, report results from our own analysis, and then reconcile our regression results with our empirical findings.

Past Estimates of Beta

Literature from the early 1990s supports the existence of a negative beta for natural gas. Kahn and Stoft (1993) regressed spot wellhead gas prices against the S&P 500 using annual data from 1980 through the first 6 months of 1992 and arrived at an estimate of beta of -0.78 (±0.27 standard error). Awerbuch has written several papers advocating the use of risk-adjusted discount rates for evaluating investments in generation assets; in them he usually cites a natural gas beta ranging from -1.25 to -0.5 (Awerbuch 1993, 1994). Awerbuch (1994) also cites another study from 1993 (by Talbot) as having found a natural gas beta of -0.45.

More recent literature surrounding the beta of energy commodities has focused on crude

[20] For a good introduction to CAPM, see Brealey and Myers (1991).

[21] In its original application to the stock market, beta represents the risk premium of a particular stock, and is related in a linear fashion to that stock's market risk (i.e., beta = expected risk premium on stock / expected risk premium on entire market). Stocks that carry the same market risk as the entire stock market (i.e., stocks whose returns are perfectly correlated with those of the broad market) have a beta of 1, while stocks that are perfectly uncorrelated with the market have a beta of 0. Similarly, stocks that are riskier than the market as a whole have betas > 1, while stocks that are negatively correlated with the market have betas < 0. While *assets* with a negative beta are desirable for diversification purposes, *liabilities* with a negative beta are undesirable for the same reason. In the case of natural gas, the producer holds the asset (and benefits from a negative beta) while the consumer is faced with a liability (and is hurt by a negative beta).

[22] Of course, *sustained* gas price increases may ultimately *not* benefit gas producers, as higher gas prices will induce substitution away from gas towards cheaper fuels.

oil, whose price should be at least moderately correlated with natural gas prices. Sadorsky (1999) and Papapetrou (2001) both conclude that there is a negative correlation between changes in oil prices and stock returns (i.e., that oil exhibits a negative beta). Sauter and Awerbuch (2002) provide an extensive survey of the literature regarding oil's impact on both the economy and stock markets, while Awerbuch (2003) estimates that the historical betas of coal, gas, and oil in the European Union are on the order of -0.4, -0.1, and -0.05, respectively.[23]

On the other hand, Pindyck (2001) – though not specifically investigating beta – notes (without citation) that estimates of beta for crude oil have been in the range of +0.5 to +1.0, and qualitatively explains why one should expect to see positive betas – strong economic growth leads to higher prices for industrial commodities. While qualitatively plausible, this intuition that strong economic growth puts upward pressure on commodity prices (resulting in a positive beta) is no more or less believable than the opposing view that high commodity prices put a damper on economic growth (resulting in a negative beta) – both interpretations merely represent successive phases of the economic cycle.[24]

Our Estimate of Beta

Seeking an updated estimate of the beta of natural gas, we regressed historical percentage changes in natural gas prices delivered to electricity generators (from the EIA) against historical percentage changes in the S&P 500 index. We choose to work with *delivered* (rather than *wellhead*) prices in an attempt to capture the full risk facing gas-fired generators: volatility in both wellhead prices and transportation costs.[25] Like Kahn and Stoft (1993), we first attempted to use monthly data (going back to January 1979), but were unable to correct for seasonality despite employing several different approaches. As a result, we too retreated to using annual averages, which remove seasonality yet also mask intra-year movements and greatly restrict sample size. We corrected for autocorrelation using the Hildreth-Lu procedure.

Figure 4 graphically displays our estimate of beta over time. The dashed line represents a "cumulative" estimate of beta, resulting from progressively longer-term regressions as one moves forward in time.[26] The gray shaded area represents a 90% confidence interval around our central estimate of cumulative beta. Meanwhile, to illustrate shorter-term variations, the solid line represents a rolling 10-year estimate of beta.[27]

As shown, our cumulative estimate of beta (the dashed line) is, not surprisingly, fairly stable over time, typically ranging from -0.2 to -0.4, yet coming to rest in 2002 at -0.1 – lower in

---

[23] It is somewhat counterintuitive that the beta of coal is found to be more negative than the beta of gas, given that the price of natural gas is generally perceived to be much more volatile than the price of coal. Price regulation in the European Union may partially explain this apparent distortion.

[24] This ambiguity brings to light two problems that arise in estimating beta using CAPM – defining what constitutes "the market" (i.e., the stock market or the broader economy), and using coincident regressions to describe relationships that are dependent on lead/lag cycles. For example, Pindyck's intuition that economic growth puts upward pressure on industrial commodity prices implies a positive beta with *the economy*. Yet if the stock market tends to anticipate economic cycles, then the stock market may begin to decline just as economic growth is pushing commodity prices higher, implying a negative beta with *the stock market*. Moreover, if the stock market's lead time over the economy (generally thought to be from 9-12 months) undergoes a long-term secular shift (e.g., as access to information improves), historical estimates of beta may not hold much meaning for the future.

[25] As demonstrated during the California energy crisis of 2000/2001, locational basis risk can be substantial: while Henry Hub gas prices peaked at around $10/MMBtu, prices delivered to the Southern California Border exceeded $50/MMBtu.

[26] This is essentially a rolling regression with a fixed starting point; i.e., the first estimate of cumulative beta shown (in 1989) results from a 10-year regression, while the 1990 estimate is from an 11-year regression, the 1991 estimate is from a 12-year regression, and so on building up to a 23-year regression in 2002.

[27] This line is simply the result of a 10-year rolling regression; i.e., each year looks only at the past 10 years.

magnitude than estimates from the early 1990s,[28] but still slightly negative. Even so, the 90% confidence interval, while skewed to the negative side of zero, is fairly wide and does not rule out the possibility of a positive beta, particularly from 1996 on. In fact, it is clear from both the confidence interval and the rolling 10-year estimate of beta that Awerbuch and others who looked at gas betas in the early 1990's were doing so at perhaps the optimal moment to conclude a negative beta. Since that time, the confidence interval has widened considerably – the opposite of what one would expect as sample size increases – and the rolling 10-year beta has oscillated between negative and positive territory. Thus, while the cumulative beta shown in Figure 4 appears to have historically been negative, it would be unwise to conclude that this will always be the case.

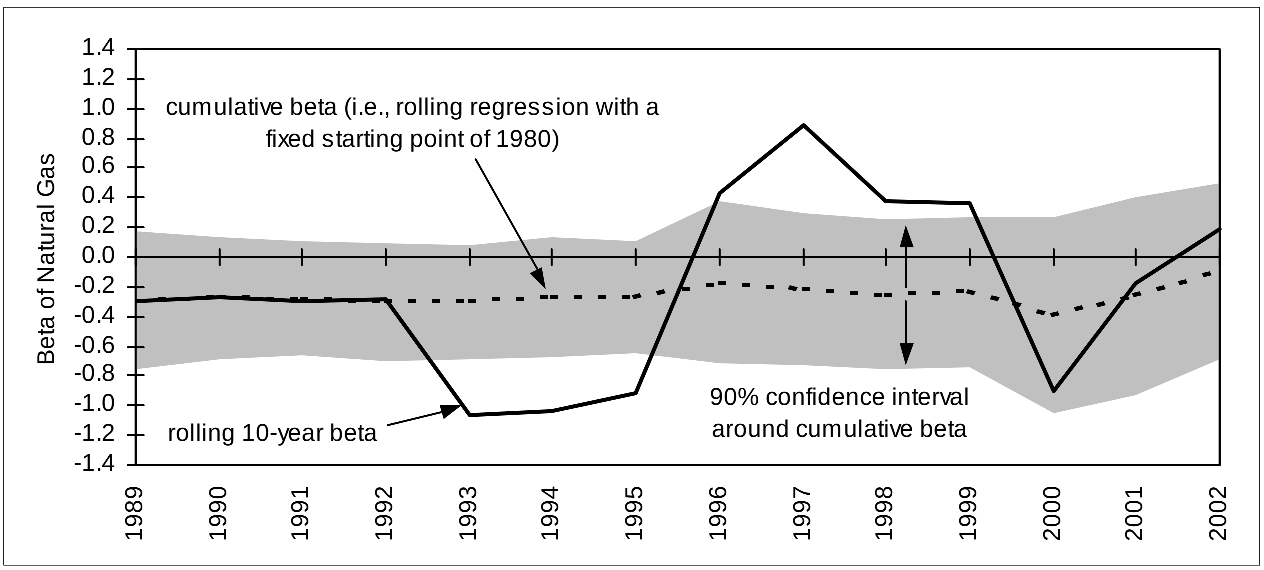

Figure 4. Estimate Of Beta Of Natural Gas Delivered To Electricity Generators

Reconciling Our Estimate of Beta with Our Empirical Premiums

Using our sample of forward prices and EIA reference case gas price forecasts, it is possible to back into empirical estimates of the beta for natural gas. These are the specific betas that would be required to fully explain our empirically derived discrepancy between natural gas forward and forecast prices within the context of CAPM. To do this, one must assume that the forward prices are "riskless" (i.e., known in advance and able to be locked in), while the price streams represented by the EIA gas forecasts are "risky" (i.e., merely a forecast and bound to be wrong). One then calculates the present value of both price streams – the forward market price stream using the known "riskless" discount rate (i.e., the U.S. Treasury bill yield at the time), and the EIA forecast price stream using whatever discount rate results in the same present value as the discounted forward market price stream. The difference between the resulting empirically derived risk-adjusted discount rate and the known "riskless" discount rate is then divided by the "market risk premium" – i.e., the historic out-performance of risky assets (stocks) over riskless assets (T-bills) – to yield beta.[29]

[28] Our cumulative estimate of beta through 1992 is less than half that estimated by Kahn and Stoft (1993) due to different data sources as well as our use of delivered prices instead of wellhead prices.

[29] Since by definition $R_{risk\text{-}adjusted} = R_{risk\text{-}free}$ + ß * Market Risk Premium, then ß = ($R_{risk\text{-}adjusted}$ - $R_{risk\text{-}free}$) / Market Risk Premium.

Performing this exercise using our sample of forward market prices and EIA reference case forecasts, and data on the historic returns of stocks and T-bills going back to 1926 from Ibbotson (2002),[30] we arrive at the various estimates of beta presented in Table 1.

Table 1. Empirical Estimates of the Beta of Natural Gas

	2-Year	5-Year	6-Year	7-Year	10-Year
Enron-AEO2001	-1.80	-1.23	*N/A*	*N/A*	-0.66
Enron-AEO2002	-2.39	-0.98	*N/A*	*N/A*	-0.40
NYMEX-AEO2003	-1.88	-0.79	-0.59	*N/A*	*N/A*
Williams-AEO2003	-2.00	-0.91	-0.72	-0.56	*N/A*
NYMEX-AEO2004	-1.59	-0.57	-0.44	*N/A*	*N/A*

The empirical estimates of beta presented in Table 1 are, at least for the longer-term forward contracts, close to the regression estimates presented in Figure 4, which estimate betas of -0.40 through 2000, -0.26 through 2001, and -0.10 through 2002. This degree of similarity suggests that the fact that longer-term natural gas forwards were priced higher than contemporaneous EIA price forecasts from 2000-2003 may be explained by CAPM. It is also evident, however, that the shorter-term forward contracts yield progressively more negative estimates of beta (as low as –2.39 for the November 2001 2-year swap) as the contract term declines, which is largely inconsistent with CAPM.[31]

As theoretically appealing as CAPM may seem, we are nevertheless hesitant to place too much faith in CAPM as the sole explanation for our empirical findings, for several reasons. First, CAPM formally requires that each individual's portfolio be fully diversified so that only market risk remains. Kahn and Stoft (1993) note that this may not be the case for the management of gas-producing companies, whose careers and reputation (if not portfolios – witness Enron's retirement plan, which was heavily invested in Enron stock) are closely tied to the profitability of the firm, and who therefore may view gas price volatility as risky even if it is negatively correlated with the stock market. Second, if CAPM were a dominant influence, then one would expect past studies of the efficiency of the natural gas futures market to regularly demonstrate a systematic positive difference between futures and realized spot prices. As discussed earlier, this is not the case: some, but not all, of these studies do show evidence of a premium, but in some cases the sign of the premium varies with net hedging pressure, rather than always being positive as would be required under CAPM with a negative beta. Thus, while CAPM may be a contributor towards natural gas futures prices that are in excess of expected spot prices, other factors are likely at work as well.

Transaction Costs

Transaction costs provide yet another potential explanation in support of the idea that those hedging gas price risk will incur incremental costs relative to market expectations of future spot prices. In commodity markets, transaction costs are manifested in the "bid-offer spread": the spread between the price at which one is willing to buy (bid) and sell (offer) a product. To

---

[30] Ibottson (2002) calculates that the average compound annual return of T-bills and large stocks (similar to the S&P 500) from 1926 through 2001 is 3.8% and 10.7% respectively, which yields a "market risk premium" (i.e., the average annual return of stocks over bills) of 6.65% (i.e., (1+10.7%)/(1+3.8%)-1 = 6.65%).

[31] Perhaps over shorter terms (i.e., 2 years), the comparison of forward market prices to a long-term gas price forecast – which is by nature not very sensitive to changes in spot or short-term futures markets – is less valid.

execute a deal with minimal price risk (i.e., the risk that the market price rises (falls) while one is trying to buy (sell)), one must typically "cross" the bid-offer spread (i.e., pay the offering price (if buying) or accept the bid price (if selling)). Since the "true" market price lies somewhere in between the bid and the offer, crossing the bid-offer spread to execute a deal results in transaction costs being incurred (i.e., paying more, or receiving less, than the "true" market price). For analytical purposes, the size of the transaction cost of buying or selling in a market is typically considered to be half the size of the bid-offer spread in that market (under the assumption that the "true" market price lies half way in between the bid and the offer).

In liquid markets, transaction costs (i.e., bid-offer spreads) are typically very small, and of little concern. In less-liquid markets (or even thinly traded segments of otherwise liquid markets), however, bid-offer spreads can be quite wide, and can have a more significant impact on the cost of transactions.

How high are transaction costs in the long-term natural gas market? While we do not have sufficient data on long-dated forwards to say, we can at least examine bid-offer spreads in the "highly liquid" but shorter-term NYMEX futures market. Figure 5 depicts typical bid-offer spreads over the first 36 contracts, plotted against "open interest" for the entire strip of 72 contracts.[32] For a trade size of 30 contracts executed under normal market conditions,[33] bid-offer spreads in the first four futures contracts (i.e., representing delivery in each of the next four months) are immaterial. Moving beyond these first few very liquid contracts, however, the bid-offer spread increases ten-fold out to 36 months, as liquidity (proxied here by open interest) declines. Beyond these first 36 months (NYMEX gas futures are listed out 72 months), NYMEX gas futures are *very* thinly traded (open interest essentially drops to zero), making it costly or even difficult to complete a trade of this size (Waller 2002).

---

[32] "Open interest" represents the number of open or outstanding contracts to the exchange (i.e., contracts that have not been closed out either through an offsetting position or via delivery). Volume – i.e., the number of contracts that changed hands in a given day – is perhaps a better measure of liquidity, though it is also more sensitive to the particular day chosen (e.g., on the Friday before a 3-day holiday weekend, volume may be very light, but open interest will likely be relatively unchanged from normal levels).

[33] Trade size and market conditions are relevant because the bid-offer spread will increase with trade size as well as during periods of heightened price volatility. By way of reference, 30 contracts per month is roughly equivalent to the amount of gas needed to fuel an 85 MW combined-cycle gas turbine (heat rate of 7,000 Btu/kWh) operating at a 70% capacity factor. A gas plant that is twice as large (170 MW) would obviously require twice as many contracts (i.e., 60) to *fully* hedge gas price risk (or could use the original 30 contracts to hedge *half* of its risk).

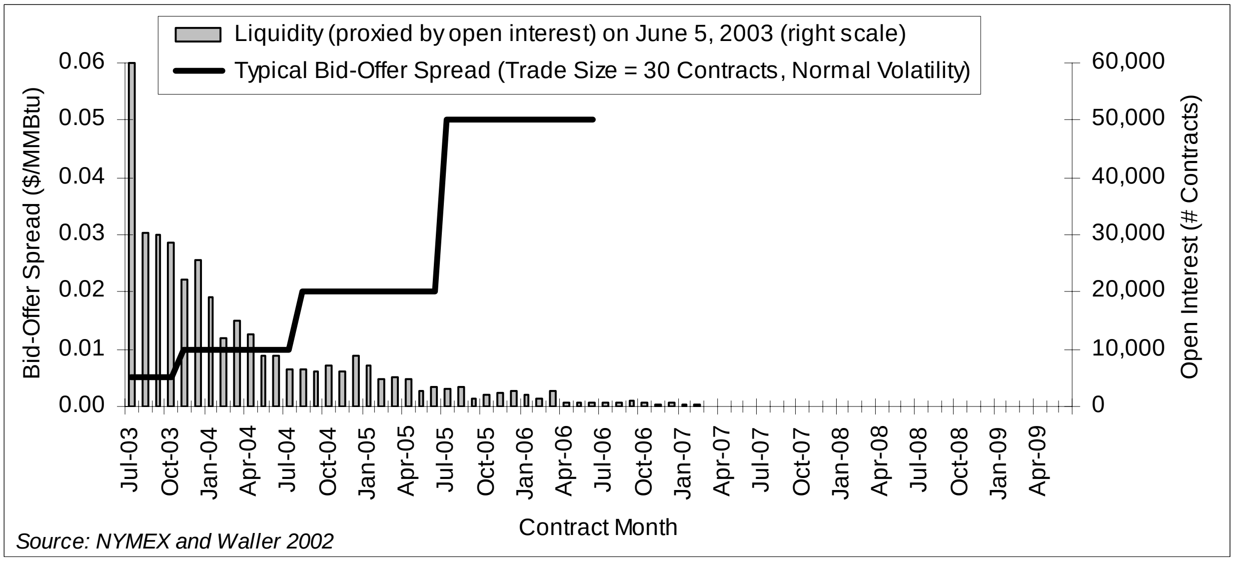

Figure 5. Typical Bid-Offer Spread and Open Interest for NYMEX Natural Gas Futures

Although the bid-offer spread shown in Figure 5 widens dramatically as liquidity declines along the strip, the absolute numbers shown (~1.5% of the contract price at its widest) are admittedly small. That said, the bid-offer spread curve in Figure 5 does not extend beyond 36 months because beyond this point there is very little liquidity on which to base a reliable estimate. In other words, beyond 36 months, there is not necessarily “a market” per se, and the cost of deal execution depends largely on the ability to locate a willing buyer or seller, which may require significant price concessions at times.

Given the steep slope of the bid-offer spread curve over the first 36 contracts, one would expect deals with longer maturities to have even higher transaction costs. Unfortunately, we do not have sufficient data to confirm this notion. When asked for an indication of a bid-offer spread on a *20-year* natural gas swap, however, one Enron gas trader estimated that *$0.50/MMBtu* would be a reasonable ballpark estimate. While the sheer magnitude of this spread is most likely an indication that no one trades natural gas swaps out that far, it nonetheless exemplifies the significant impact that transaction costs can have in illiquid markets. This impact has likely not been adequately captured by the existing net hedging pressure and efficient market literature cited earlier, which focuses almost exclusively on *short-term* natural gas futures contracts, rather than the long-term forwards in which we are interested. If the transaction costs of buying and selling long-term forward contracts are significant, then empirical studies of the returns earned by speculators on long-term forward contracts may yield results that are far more conclusive than those presented earlier for short-term contracts.

The Forecasts May Be Out of Tune with Market Expectations

Another possible explanation for our empirically derived premiums between gas forwards and forecasts is related to the forecasts themselves, which may have been biased downwards or otherwise inconsistent with market expectations of spot gas prices from November 2000-2003. Proponents of this explanation might argue that forward prices should represent the market’s view of future spot gas prices, and therefore that there is no incremental cost of hedging gas price volatility. Therefore, any real (i.e., not due to data problems) positive difference between gas forwards and gas price forecasts must represent a downward bias in the forecasts relative to the market's view of future spot prices.

For example, perhaps the U.S. natural gas industry has recently run into a serious short-term supply constraint that has shifted gas prices (and the market's expectations of future gas prices) into a new regime, and this paradigm has not been reflected in EIA reference case gas price forecasts generated from November 2000-2003; the EIA itself appears to both acknowledge and dispute such a scenario.[34] Furthermore, given that the EIA's reference case gas price forecasts were at the high end of the range relative to other common industry price forecasts over this period, this explanation directly takes issue not only with the EIA reference case forecast, but also with virtually every other publicly available gas price forecast from November 2000-2003.[35]

While this explanation is certainly plausible, particularly in light of the considerable market turmoil over this period, it cannot be directly supported or refuted by our data. The EIA's own evaluation of the accuracy of its past AEO forecasts has found that "As a general rule, the rate of increase in nominal energy prices has been overestimated rather than underestimated" (Sanchez 2003). Specifically, an examination of the accuracy of past gas price forecasts (from the past twenty-one AEOs) reveals that, except in the years 2000-2002, AEO gas price forecasts have consistently overestimated (not underestimated) the actual spot price of gas for the past eighteen years (Sanchez 2003). While these historic results may in no way be indicative of future trends, the weight of history nevertheless does not lend ready support to the view that the EIA's reference case natural gas price forecast is systematically biased *downwards* relative to the market's expectation of future natural gas prices.

Another possibility is that our use of such forecasts as proxies for the market's view of future spot gas prices is altogether inappropriate. For example, the EIA notes that its "reference case" forecast assumes that normal inventories and weather, as well as current laws and regulations, will hold throughout the forecast period, and therefore that the reference case forecast does not necessarily reflect what the EIA believes to be "most likely." In fact, the EIA does not assign probabilities to any of the forecasts it generates, so the "high economic growth case" forecast might be considered just as likely as the "reference case" or even "low economic growth case" forecast, for example. Furthermore, by assuming away weather and inventory variability, and possible changes in regulations – all of which have a major impact on prices – the EIA notes that it is not really forecasting *prices* at all, but rather long-term equilibrium *costs*.[36]

While the EIA reference case forecast may not be *designed* to represent EIA or market expectations of future gas prices, it deserves note that industry participants and energy analysts

[34] We noted with interest the following text accompanying the EIA's early release of AEO 2004: *"For almost 4 years, natural gas prices have remained at levels substantially higher than those of the 1990s. This has led to a reevaluation of expectations about future trends in natural gas markets, the economics of exploration and production, and the size of the natural gas resource. The Annual Energy Outlook 2004 (AEO2004) forecast reflects such revised expectations…."* This statement seems to suggest two things: (1) that the EIA believes that its AEO 2004 gas price forecast reflects upwardly revised market expectations concerning future gas prices, and (2) by extension, that the previous three years of AEO gas price forecasts may not have. It is interesting, therefore, that the NYMEX-AEO 2004 comparison shown earlier in Figure 2 reveals a positive price premium that is largely consistent with those found in the previous three years.

[35] While it is certainly possible that virtually all gas price forecasts have been biased downwards (relative to market expectations of future spot prices) over the 2000-2003 time period, especially given likely similarities between the various models behind each forecast, the "absolute" nature of this consideration is nonetheless worth noting. If true, this explanation would call into question the use of virtually all long-term gas price forecasts generated from 2000-2003 for *any* purpose.

[36] We note, however, that *Annual Energy Outlook* uses the term "price" instead of "cost" to describe its forecasts.

regularly adopt the EIA reference case projection as a "best estimate" of future energy outcomes; in fact, the EIA itself regularly uses its reference case forecast as the "base-case" forecast when evaluating the cost and impacts of energy policies. Furthermore, Bolinger et al. (2003) demonstrate that some utilities – one important segment of the gas market – are relying on EIA reference case forecasts as a "best estimate" of future gas prices for the purpose of long-term resource planning.

On a related note, some have argued that "reference case" gas price forecasts can best be thought of as *modes* (i.e., the single scenario that the forecaster believes to be the most likely) rather than *means* (i.e., a probability weighted average of all possible spot prices), and since market expectations are by definition *mean* expectations, reference case gas price forecasts cannot represent market expectations. While this argument does not hold for the EIA reference case forecast, which as described above may be neither a mean nor a mode, the implications of this argument are nonetheless worth noting. Specifically, since gas prices are generally believed to be lognormally distributed or positively skewed, the mean must lie above the mode, meaning that true market expectations must be higher than reference case gas-price forecasts (if those forecasts do indeed represent the mode). If this argument is accurate, it might explain some or all of the premium we have observed between forward prices and price forecasts. More importantly, however, this argument calls into question why utilities and others would ever place significant emphasis on the use of reference case gas price forecasts in modeling and planning exercises. By doing so, they would be *systematically underestimating* the market's expectations of future gas prices, thereby erroneously making gas-fired generation appear to be cheaper than it is likely to be on average.

Other Data Issues May Be Driving the Premium

A final possible explanation for our findings is simply that our analysis is plagued by data issues that prohibit a meaningful comparison between natural gas forward prices and price forecasts. While we have previously raised the possibility that forecast problems are at least partly responsible for the premium, it is also possible that the other half of the equation – the forward price – is upwardly biased. Or, even if the forward price and price forecast are unbiased, it is possible that potential changes in the market between when the forward prices were sampled and the forecasts were generated could account for some or all of the observed premiums.

The two main potential concerns with our forward price sample – and particularly prices from Enron and Williams, as the NYMEX is a regulated exchange[37] – are price manipulation (given the role that both Enron and Williams allegedly played in precipitating California's electricity crisis) and credit risk (given Enron's impending bankruptcy in November 2001). As shown in Bolinger et al. (2003), however, Enron 2-year swap prices from both November 2000 and 2001 are entirely consistent with the NYMEX strip at the time, implying that at least on the short end of the forward curve, Enron's prices were not biased upwards. Furthermore, any credit risk in the November 2001 price sample would lead to *lower*, not *higher* prices, as the buyer would require compensation for accepting Enron's credit risk.

[37] Though they are likely less susceptible to manipulation than are over-the-counter prices, NYMEX prices are not totally beyond question. As noted earlier, liquidity is extremely poor in the later years of the NYMEX gas futures strip, with several of the longest-dated contracts never having traded. The NYMEX has a somewhat involved procedure for calculating settlement prices for such illiquid contracts that involves considering the spread relationships between the contract in question and more actively traded contracts. Without considering the merits of this approach, it is worth noting that in the absence of actual trades (or at least strong bids and offers) it may not be possible to buy or sell a long-dated gas futures contract at or near the NYMEX settlement price.

Timing mismatch also does not appear to be a factor. Two to three weeks passed between when the AEO 2001 and AEO 2002 gas price forecasts were finalized and when we were able to sample Enron swap prices in Novembers 2000 and 2001, respectively. An examination of long-dated NYMEX gas futures contracts over this interim period shows very little price movement, implying that there were no fundamental shifts in market expectations between when the forecasts were finalized and when the forward prices were sampled. Meanwhile, our November 2002 and October 2003 forward price samples were essentially coincident with the finalization of the AEO 2003 and AEO 2004 gas price forecasts, respectively.

CONCLUSIONS

A comparison of long-term natural gas forward prices and EIA reference case natural gas price forecasts from November 2000-2003 reveals that long-term forward prices have traded above EIA price forecasts during this four-year period. As shown in Figure 2 earlier, the magnitude of the empirically derived premiums ranges from $0.5-$0.8/MMBtu, or 0.4-0.6¢/kWh assuming a highly efficient gas-fired power plant. Furthermore, in recent years the EIA reference case gas price forecasts have typically been higher – and often substantially so – than most other forecasts that are commonly used by utilities and others in resource acquisition, modeling, and planning exercises. Thus, the premiums observed relative to the EIA reference case forecasts would be *even larger* when comparing forward prices to some of the other forecasts commonly used in the electricity industry.

While one cannot easily extrapolate these findings beyond the time period of our data (2000-2003), or to contract terms longer than those examined (2-10 years), it is at least apparent that utilities and others who have conducted resource acquisition, planning, and modeling studies based on EIA reference case (as well as other) gas price forecasts from November 2000-2003 have arguably produced "biased" results that favor variable-price gas-fired over fixed-price renewable generation, potentially to the tune of ~0.4-0.6¢/kWh levelized. This is because if consumers are rational and value price stability, then the cost of fixed-price renewable generation should be compared to the *hedged* or *guaranteed* cost of natural gas-fired generation, rather than to *projected* costs based on *uncertain* gas price forecasts.

The cause of the premiums is less apparent. We explored three potential explanations for our empirical findings – that hedging is not costless, that the forecasts are biased, and that other data issues are plaguing our analysis – and found each of them to be neither fully satisfying nor easily refutable. Regardless of the explanation, however, the basic implication of our study for renewable energy (or any other energy source that is immune to natural gas fuel price risk) remains the same: *one should not blindly rely on gas price forecasts when comparing the levelized costs of fixed-price renewable to variable-price gas-fired generation contracts.* If there is a cost to hedging, gas price forecasts do not capture and account for it. Alternatively, if gas price forecasts are at risk of being biased or out of tune with the market, then one certainly would not want to use them as the basis for investment decisions or resource comparisons if a better source of data (i.e., forwards) existed. Accordingly, in virtually all cases, the most comprehensive way to compare the levelized costs of renewable and gas-fired generation would be to use forward natural gas price data as opposed to natural gas price forecasts.

Of course, given widespread credit concerns and the retrenchment of trading desks across the United States in the wake of Enron's collapse, data availability is likely to be an issue. In particular, beyond the six years of publicly available NYMEX futures prices, long-term forward gas price data may be hard to come by. Where forward price data from actual contracts are not publicly available, utilities and policymakers may still have access to (or be in a position to solicit) data that are not in the public domain. Alternatively, natural gas-fired generators may be willing to internalize any cost of hedging (or take on fuel price risk) and offer a long-term fixed-price electricity contract similar to that provided by most renewable generators, thereby obviating the need for a utility or regulator to collect forward gas price data for the purpose of addressing fuel price risk. Should these two suggestions fail, a second-best approach may be to start with the six-year NYMEX strip (appropriately adjusted for locational basis if necessary) and then transition to a long-term gas price forecast (or perhaps some average of multiple long-term forecasts) in the seventh year and thereafter. The transition between the NYMEX and forecast prices could be made by simply switching to the forecast price in year seven, or alternatively by applying the escalation rates implied in the long-term forecast to the year six NYMEX price.

Finally, as a last resort, if all else fails, utilities and others may wish to adjust forecast gas prices upwards to account for the fact that forward prices have, potentially for reasons discussed above, traded above price forecasts in recent years. While the analysis in this report suggests that an adjustment ranging from \$0.5-\$0.8/MMBtu (0.4-0.6¢/kWh at a heat rate of 7,000 Btu/kWh) is a reasonable starting point, we emphasize that these premiums were calculated with respect to EIA reference case price forecasts from 2000 to 2003, for terms ranging from 2-10 years. If using a different gas price forecast, a higher or lower adjustment may be warranted. Likewise, this historically derived premium may well vary in the future, and may vary with contract terms longer than 10 years. For these reasons, this approach is inferior to those previously mentioned, though it may still be superior to simply relying on forecast prices, which – at least in recent years – have been shown to be significantly below forward prices.

REFERENCES

Awerbuch, S. (1993). "The Surprising Role of Risk in Utility Integrated Resource Planning." *The Electricity Journal* 6 (3): 20-33.

________. (1994). "Risk-Adjusted IRP: It's Easy!!!" In *Proceedings of the NARUC-DOE Fifth National Conference on Integrated Resource Planning*, Kalispell, Mont., 228-269.

________. (2003). "Determining the Real Cost: Why Renewable Power is More Cost-Competitive Than Previously Believed." *Renewable Energy World*, 6(2), March-April 2003.

Bachrach, D., R. Wiser, M. Bolinger, W. Golove (2003). *Comparing the Risk Profiles of Renewable and Natural Gas Electricity Contracts: A Summary of the California Department of Water Resources Contracts.* LBNL-50965. Berkeley, Calif.: Lawrence Berkeley National Laboratory.

Bessembinder, H., and M. Lemmon (2002). "Equilibrium Pricing and Optimal Hedging in Electricity Forward Markets." *Journal of Finance*, 57 (3), 1347-1382.

Bolinger, M., R. Wiser, W. Golove (2003). *Accounting for Fuel Price Risk: Using Forward Natural Gas Prices Instead of Gas Price Forecasts to Compare Renewable to Gas-Fired Generation*. LBNL-53587. Berkeley, Calif.: Lawrence Berkeley National Laboratory.

Brealey, R.A. and S.C. Meyers (1991). *Principles of Corporate Finance*. San Francisco: McGraw-Hill, Inc.

Brooks, D., E. Lo, R. Zavadil, S. Santoso, J. Smith (2003). *Characterizing the Impacts of Significant Wind Generation Facilities on Bulk Power System Operations Planning: Xcel Energy – North Case Study, Final Report.* Prepared for The Utility Wind Interest Group.

Buchanan, W.K., P. Hodges, J. Theis (2001). "Which way the natural gas price: an attempt to predict the direction of natural gas spot price movements using trader positions." *Energy Economics* 23 (2001) 279-293.

Carter, C.A., G.C. Rausser, A. Schmitz (1983). "Efficient Asset Portfolios and the Theory of Normal Backwardation." *The Journal of Political Economy*, 91 (April 1983), 319-331.

Cavanagh, R., A. Gupta, D. Lashof, M. Tatsutani (1993). "Utilities and CO_2 Emissions: Who Bears the Risks of Future Regulations?" *The Electricity Journal*, 6(2): 64-75.

Chang, E.C. (1985). "Returns to Speculators and the Theory of Normal Backwardation." *The Journal of Finance*, 40 (1), 193-208.

Chinn, M.D., M. LeBlanc, O. Coibon (2001). *The Predictive Characteristics of Energy Futures: Recent Evidence for Crude Oil, Natural Gas, Gasoline and Heating Oil*. UCSC Economics Working Paper No. 490

Commodity Futures Trading Commission (2003). *Commitments of Traders Reports*. Available at www.cftc.gov, accessed July 28, 2003.

Cootner, P. (1960). "Returns to Speculators: Telser vs. Keynes." *The Journal of Political Economy*, 69 (August 1960), 396-404.

de Roon, F.A., T.E. Nijman, C. Veld (2000). "Hedging Pressure Effects in Futures Markets." *The Journal of Finance*, 55 (3), 1437-1456.

Dragoon, K. (2003). "PacifiCorp IRP Operating Impacts Study Update." *Presentation to the Utility Wind Interest Group Annual Meeting*, Denver, Colorado, April 16, 2003.

Dusak, K. (1973). "Futures Trading and Investor Returns: An Investigation of Commodity Market Risk Premiums." *The Journal of Political Economy*, 81 (6), 1387-1406.

Electrotek Concepts, Inc. (2003). *We Energies Energy System Operations Impacts of Wind Generation Integration Study*.

Energy Information Administration (EIA) (2000). *Annual Energy Outlook 2001.* DOE/EIA-0383(2001), December 2000, Washington, DC.

________________ (2001). *Annual Energy Outlook 2002.* DOE/EIA-0383(2002), December 2001, Washington, DC.

________________ (2003). *Annual Energy Outlook 2003.* DOE/EIA-0383(2003), January 2003, Washington, DC.

________________ (2004). *Annual Energy Outlook 2004.* DOE/EIA-0383(2004), January 2004, Washington, DC.

Gray, R.W. (1961). "The Search for a Risk Premium." *The Journal of Political Economy*, 69 (June 1961), 250-260.

Herbert, J.H. (1993). "The Relation of Monthly Spot to Futures Prices for Natural Gas." *Energy*, 18 (11), 1119-1124.

Hirshleifer, D. (1988). "Residual Risk, Trading Costs, and Commodity Futures Risk Premia." *The Review of Financial Studies*, Summer 1998.

Hirst, E. (2002). "Integrating Wind Energy With the BPA Power System: Preliminary Study." *Prepared for the Bonneville Power Administration*, September 2002.

Hoff, T.E. (1997). *Integrating Renewable Energy Technologies in the Electric Supply Industry: A Risk Management Approach*. NREL/SR-520-23089. Golden, Colorado: National Renewable Energy Laboratory.

Houthakker, H.S. (1957). "Can Speculators Forecast Prices?" *Review of Economics and Statistics*, 39 (May 1957), 143-151.

Hull, J.C. (1999). *Options, Futures, and Other Derivatives*. Upper Saddle River, New Jersey: Prentice Hall.
Humphreys, H.B. and K.T. McClain (1998). "Reducing the Impacts of Energy Price Volatility Through Dynamic Portfolio Selection." *The Energy Journal*, vol. 19 (3), pp. 107-131.
Ibbotson Associates (2002). *Stocks, Bonds, Bills, and Inflation: 2001 Yearbook*. Chicago: Ibbotson Associates.
Kahn, E. and S. Stoft (1993) (unpublished draft). *Analyzing Fuel Price Risks Under Competitive Bidding*. Berkeley, Calif.: Lawrence Berkeley National Laboratory.
Keynes, J.M. (1930). *A Treatise on Money*. London: Macmillan.
Longstaff, F.A. and A.W. Wang (2003). *An Empirical Analysis of the Risk Premia in Electricity Forward Prices*, June 4, 2003.
Papapetrou, E. (2001). "Oil Price Shocks, Stock Market, Economic Activity and Employment in Greece." *Energy Economics*, Vol. 23 (2001), pp. 511-532.
Pindyck, R. S. (2001). *The Dynamics of Commodity Spot and Futures Markets: A Primer*. Cambridge, Mass.: Massachusetts Institute of Technology.
Pirrong, C. and M. Jermakyan (2001). *The Price of Power: The Valuation of Power and Weather Derivatives*, October 12, 2001.
Repetto, R. and J. Henderson (2003). "Environmental Exposures in the US Electric Utility Industry." *Utilities Policy*, 11: 102-111.
Sadorsky, P. (1999). "Oil Price Shocks and Stock Market Activity." *Energy Economics*, Vol. 21 (1999), pp. 449-469.
Sanchez, E. (2003). "*Annual Energy Outlook* Forecast Evaluation." Accessed January 26, 2004 from www.eia.doe.gov/oiaf/analysispaper/forecast_eval.html.
Sauter, R. and S. Awerbuch (2002) (draft). *Oil Price Volatility and Economic Activity: A Survey and Literature Review*. IEA Research Paper, International Energy Agency, Paris.
Telser, L.G. (1958). "Futures Trading and the Storage of Cotton and Wheat." *The Journal of Political Economy*, 66 (3), 233-255.
Waller, R. (2002). E-mail from Director of Trading, Bridgewater Associates, citing NYMEX natural gas bid-offer spread information from Goldman Sachs. February 6, 2002.
Walls, W.D. (1995). "An Econometric Analysis of the Market for Natural Gas Futures." *The Energy Journal*, 16 (1), 71-83.