Preliminary Draft
Comments Welcome

Security of Supply and Forward Price Premia: Evidence from the Natural Gas Industry

Severin Borenstein, Meghan Busse and Ryan Kellogg

February 2007

Abstract: Firms must often choose whether to procure essential inputs to production in either spot or forward markets. In this paper, we focus on how natural gas local distribution companies (LDCs) make this decision when procuring gas in wholesale markets. A significant feature of this regulated industry is that the costs to the LDCs of insufficient procurement of gas far exceed the costs of over-procurement, which leads LDCs to be concerned about "security of supply." We argue that because spot markets for wholesale gas can be thin, LDCs will be willing to pay forward price premia to guarantee access to gas at times when supply is expected to be constrained. Using price data on 117 local gas spot markets and five major interstate gas pipelines, we find evidence that forward prices do exceed expected spot prices at such times. We further present institutional and empirical evidence that a more standard price risk aversion phenomenon is not the cause of these forward premia.

Borenstein: University of California Energy Institute (www.ucei.org); Haas School of Business, University of California, Berkeley (faculty.haas.berkeley.edu/borenste); and National Bureau of Economic Research (www.nber.org). Email: borenste@haas.berkeley.edu. **Busse:** University of California Energy Institute and Haas School of Business, University of California, Berkeley. Email: meghan@haas.berkeley.edu. **Kellogg:** University of California Energy Institute and Department of Agricultural and Resource Economics, University of California, Berkeley. Email: kellogg@are.berkeley.edu. We are grateful to U.C. Berkeley's Coleman Fung Risk Management Research Center for financial support.

1. Introduction

In buying inputs, firms often face a choice between buying on a spot market or through a forward transaction. Each kind of transaction has advantages, and which market the firm chooses will often depend on the characteristics of the input it is buying. For example, if the input is a standard commodity, the spot market will be more attractive than if the input involves important "specificities" that must be negotiated between the firm and its supplier. Similarly, if production of the input requires a long lead time, that will favor using the spot market more than if productioncan be ramped up quickly. On the other hand, if the buyer is unsure of the quantity it will need, it might choose to purchase in the spot market, after it has a better idea of its own demand.

In the natural gas industry, end users of gas, primarily local distribution companies (LDCs), speak of "security of supply" as another important consideration in the decision whether to procure their inputs—wholesale natural gas—in a spot or in a forward market. Security of supply refers to an LDC's desire to have supplies of gas on hand in case it needs them, rather than relying on the spot market to procure gas should the firm need more than a typical amount.

Such a preference arises because LDCs face asymmetric penalties for having too much versus insufficient gas inputs to meet demand. The consequence of having insufficient supplies of gas is the curtailment of gas supply to downstream customers. Curtailments are likely to cause regulatory scrutiny and political fallout, both of which the LDC will very much want to avoid. In contrast, the costs of having too much supply are likely to be much smaller. These costs could include either explicit inventory holding costs or the opportunity cost of not selling excess holdings of gas; however, there is little advantage to an LDC of selling off excess gas, even if there are profits to be earned by doing so. This is because rate-of-return regulation will pass most of these profits on to ratepayers. Thus, an LDC is in a position of facing a large penalty for having insufficient supply of natural gas inputs, but a small penalty for having excess supply.

We suspect that such considerations, which we also refer to as "quantity risk," may exist in other industries, even if they are not generally referred to as "security of supply" concerns. For instance, consider the hypothetical example of a manager of an auto plant who has on hand more tires than he thinks he is likely to need given the range of output levels the plant might be called on to produce in the next month. Should he try to rid the plant of some of the excess supply of tires if he can do so at a profit? If he makes a profit on the sale of the tires, it is likely to lead to a small reward. If, however, the plant is called upon to produce an unusually large number of cars, and the production line has to be shut down because there aren't enough tires, then the manager is likely to suffer substantial negative consequences. In this case, it is organizational incentives rather than regulatory incentives that are the root cause of the security of supply concern. Yet, in both cases, the penalties for having insufficient supply are much larger than the penalty for having excess supply.

The presence of asymmetric penalties alone, however, will not necessarily lead a firm to pay a price premium to transact on the forward market instead of the spot market. If the spot market is sufficiently liquid, then the firm will be able to rely on the spot market even with security of supply concerns. For example, consider the case of a long-distance trucking firm. If the firm were to have insufficient fuel, it would cripple the firm's ability to perform. However, long-distance trucking firms can rely on a dense, geographically diverse, and very reliable spot market for fuel in the form of gas stations and truck stops. They do not need to make forward purchases of fuel in order to avoid security of supply risk.

Thus, if the input spot market is sufficiently thick, we expect that it would be a good substitute for the forward market, even for firms that have security of supply concerns. In such a situation, we would not expect firms to pay a price premium in order to transact in the forward market. Instead, we would expect the forward price to equal the expected spot price. However, if spot markets are thin, meaning that a firm seeking to purchase inputs on the spot market may find itself unable to do so despite a very high reservation value, then we may find that the forward price exceeds the expected spot price.

In this paper, we argue that the combination of security of supply concerns and illiquid spot markets in the natural gas industry can generate forward price premia paid to guarantee access to natural gas supply. We develop a model of natural gas forward and spot markets which predicts that at times when gas is readily available, the forward price will equal the expected spot price. However, when gas supply is "tight" and the gas pipeline transportation network is capacity constrained, the model predicts that the forward price will exceed the expected spot price. Using a dataset of forward and spot prices in local gas markets and on five interstate pipelines, we present empirical evidence consistent with this model. We also argue that the institutional setting and some of the empirical findings are not conducive to a more conventional price risk aversion explanation of the observed forward premia.

In what follows, we first describe in Section 2 the relevant institutional details of the natural gas industry. In Section 3 we present a theoretical model of forward and spot wholesale gas procurement which draws from the institutional setting. Section 4 discusses empirical tests of this model using price data from local natural gas markets, and Section

5 does the same with price data from markets for gas transportation rights. Section 6 presents additional evidence from forward trading volume data, and Section 7 concludes.

2. The Natural Gas Industry

2.1. Market Structure

The wholesale natural gas industry can be broadly classified into three main segments: production, transportation, and local distribution. Natural gas producers range in size from large, integrated oil and gas firms to very small firms owning only a few wells; production is generally considered to be competitive.[1] The vast majority of U.S. natural gas production occurs in a belt running northwest to southeast from the Rocky Mountains to the Gulf of Mexico; however, the demand centers for natural gas are primarily located in the northeast, upper midwest, and west coast. Thus, a network of interstate gas transmission pipelines has been developed to transport gas to where it is needed. Each pipeline is a distinct private legal entity, though some entities are jointly owned by a single holding company.

Interstate natural gas pipelines are regulated by the Federal Energy Regulatory Commission (FERC). According to FERC regulation since 1992, pipelines are not permitted to buy and sell natural gas; they may only act as the transporter of gas on behalf of their customers, known as shippers, who hold title to the gas itself. Shippers may be upstream gas producers, downstream gas consumers, or gas marketers.

Pipelines sell rights to most of their capacity to shippers under long-term contracts which can be as long as 30 years in length. In 2002, 78% of all subscribed capacity was contracted for longer than one year, and 36% was contracted for longer than five years (NPC, 2003). The maximum price, known as the "reservation charge", that shippers can be charged for this property right is set by FERC. [2] There is a secondary market, called the capacity release market, in which a shipper who holds long-term transportation rights can sell part or all of its transportation rights to another shipper for a fixed period of time, usually one calendar month. The price for this transfer is negotiated beteween the "releasing" and "awarded" shippers, but is capped at the maximum reservation charge set

[1] Wellhead natural gas price deregulation began in 1978 with the Natural Gas Policy Act. Prices were fully de-controlled in 1993 under the Natural Gas Wellhead Decontrol Act.

[2] The charge is set at a level that will recover capital costs plus a reasonable rate of return on investment.

by FERC.[3] Shippers can also buy "interruptible" service from the pipelines directly on a short-term basis. However, the service is not guaranteed, and the pipelines may have to deny access to an interruptible customer if the holders of firm capacity rights utilize their full capacity.

Via these transportation arrangements, gas is delivered to wholesale end-users, which are often regulated local distribution companies (LDCs) responsible for gas distribution to ratepaying customers.[4] The LDCs are regulated by state Public Utilities Commissions (PUCs) with two primary objectives: (1) to ensure reliable gas supply so that customers will not be interrupted, even during peak demand periods, and (2) to minimize customer rates while allowing the utility to achieve a reasonable rate of return. The prices LDCs charge to their customers are regulated through traditional cost-of-service regulation, under which the wholesale cost of gas supply is passed through to ratepayers with no markup.

2.2. An LDC's Gas Procurement Decision

Given the market structure described above, a load-serving LDC has numerous gas procurement options at its disposal. Broadly speaking, it must make its decision along two dimensions: (1) whether to take ownership of wholesale gas "downstream" at the location of its customers or at an "upstream" producing location, and (2) how far in advance to procure. We now consider the institutional aspects of these decisions in turn.

Should the LDC elect to procure gas downstream, its procurement process is relatively simple—it must contract with a gas supplier for delivery in its local area. In this case, it does not need to contract for capacity on an interstate pipeline: it is the responsibility of the gas supplier to ensure transportation to the point of sale. Alternatively, if the LDC arranges for gas delivery in an upstream producing zone, it is then also responsible for contracting sufficient transportation to ship the gas to its local area. In either case, coordination is necessary amongst the three market participants—the LDC, the gas seller, and the pipeline—to ensure that the appropriate quantities are delivered at the appropriate times. The transactions costs involved in achieving this coordination, which include the search costs of finding a suitable counterparty for each transaction[5], can be signifi-

[3] Pursuant to FERC Order 637, the rate cap on the capacity release market was waived from 1 September 2000 - 30 September 2002. While this does not impact our primary analysis, it does factor into some supplementary results described in Section 6.

[4] Some merchant electric generators and large industrial firms also purchase gas directly in wholesale markets.

[5] Apart from the NYMEX natural gas futures market, local natural gas and pipeline capacity markets

cant, particularly in high demand periods when pipelines are operating at their capacity constraints and interruptible transportation service is generally unavailable.

LDCs also have several options for the timing of their purchases. Day-ahead spot markets for gas currently operate in approximately 100 locations across the U.S., and are liquid to varying degrees: while the day-ahead market at Henry Hub in Louisiana (the delivery point for gas traded under a NYMEX monthly natural gas futures contract) is regarded as deeply liquid, industry participants have told us that at many locations, firms may sometimes be unable to purchase gas on the spot market because liquidity is limited by the coordination problems referenced above. That is, in the spot market, an LDC faces a risk that it will be unable to procure gas when it needs it.

As an alternative to participating in day-ahead spot markets, an LDC can also procure gas and gas pipeline capacity in advance on a monthly basis in the "bidweek" market and the capacity release market. "Bidweek" refers to the last five trading days of each month. It is during this period that firms contract for natural gas delivery in local gas markets for the coming month. Our interviews with industry participants have indicated that both the bidweek market and the capacity release market can exhibit a lack of liquidity similar to that of the spot market. However, an LDC that successfully purchases gas in these forward markets eliminates the possibility that it may need to rely on an uncertain spot market for its gas supply.

3. A Theory of Forward and Spot Prices with Quantity Risk

As we have explained, wholesale buyers can procure gas in a downstream spot market or a downstream forward market, or can procure pipeline capacity to bring gas from an upstream producing location using either a spot or forward market (a buyer of pipeline capacity, of course, must also line up gas supplies at the upstream location). We present the model in the context of the local gas market, but point out how it applies to either the gas or the pipeline capacity market.

3.1. Quantity Risk

In an efficient market with many risk-neutral buyers and sellers of gas, one would expect the forward price to be equal to the expected spot price. If a market is found in which forward prices systematically differ from spot, the difference is frequently attributed to risk aversion over price, either *buyers* being systematically more averse to price risk – causing

are bilateral trading systems without a central market-maker.

a forward premium – or *sellers* being more price risk averse – causing a forward discount. We return to price risk aversion later in this section and argue that it is not a plausible explanation in this market. Instead, we present an alternate theory of a forward premium that does not rely upon price risk aversion and is consistent with the factors that industry participants have spoken about in our interviews with them.

Industry participants have uniformly emphasized the importance of "security of supply," *i.e.,* the risk that an LDC will not be able to obtain the quantity of gas that its customers require. Most LDCs hold limited gas inventories, so they must balance the supply they purchase with end-use customer demand over short periods. In a well-functioning wholesale market with many participants, supply shortages would cause the price to rise until the shortage is eliminated. Price increases at the retail level when supplies are tight would also decrease end-use quantity demanded. However, neither of these mechanisms works well in the gas market.

At the retail level, end-use demand is not rationed through price increases when supply is short. Regulated retail prices usually do not change more frequently than once a month. More importantly, retail prices are set based on the LDC's *average* cost of procurement, which averages in longer term contract prices, not its marginal cost, so retail price adjustments do not accurately reflect changes in the contemporaneous market price. As a result, end-use demand elasticity is minimal. In those parts of the country where gas is a major fuel for electricity generation, there may be more elasticity due to demand for use in generation, but the economics of electricity generation dampen that price responsiveness as well.[6]

In the wholesale market, the regulatory process that governs most gas buyers creates a barrier to efficient adjustment when supplies are short. Because regulated utilities procuring gas for end use or for electricity generation pass through the cost of gas acquisition, including the gains or losses from buying or selling pipeline capacity, these companies have less incentive to respond to potentially profitable trading opportunities. In particular, they are likely to eschew opportunities to sell local gas reserves or firm pipeline capacity rights unless they are extremely confident that they will not need the rights themselves to bring in gas for their own end-use customers. The profits from such sales would be nearly completely passed through to customers. But the cost and regulatory scrutiny borne by

[6] The reason is that, over a fairly large range of gas prices, gas generation is more costly than baseload coal, nuclear, or hydroelectric units, and less costly than oil-fired units. Furthermore, gas generation plays a smaller role in the winter, when we are most likely to see gas supply constraints. This is not to say that gas generation cannot be curtailed in a shortage situation, but that such curtailments are the sort of quantity risk that utilities try to avoid.

the LDC would be significant if it had to curtail customers because it ran short of gas after it sold some of its local gas reserves or pipeline capacity rights.

Thus, a retail provider of gas faces the concern that when end-use demand spikes it will be unable to quickly procure the necessary supplies and will instead be forced to curtail the usage of some of its customers. This is what retailers seem to fear when they argue, as did one utility executive to us, that cutting off customers involuntarily is "suicide" because of the expected reaction from regulators.

3.2. Supply and Demand Model

To illustrate the implications of these incentives, consider a model in which several wholesale gas buyers—LDCs who need to meet end-use customer demand—are present at the downstream end of a pipeline. The LDCs hold "reserves," the term we will now use to refer to the company's contracted access to gas in either local storage or through firm contracts for pipeline delivery, and can use forward and spot gas markets to true up their reserve position with their anticipated customer demand in a given period. Customer demand varies exogenously due to weather and other shocks and is extremely inelastic in the short run.

When the period of delivery arrives, the exogenously-determined customer demand and the storability of excess supply combine to yield a convex LDC demand function. The marginal value of supplies is very high for units up to the customer demand that the LDC faces and flattens out rapidly for units beyond that level, since additional quantities can be put into storage. At the delivery period the LDC faces very asymmetric costs of failing to match its procured supply with customer demand—much greater costs if it has too little supply than if it has too much. If, however, it could always be sure that it could adjust its total procurement quantity to exactly match its customers' demand through use of the spot market, this asymmetry would be of no consequence.

In reality, we have been told by many market participants that the local spot market for natural gas is often not sufficiently liquid that LDCs can be certain that they can buy or sell the quantities they wish to trade. For this reason, LDCs procure much of their supplies in the forward markets, usually between one and six months in advance. Consider the situation that the LDC faces when it attempts to buy its requirements before retail demand shocks are known. To begin with, we consider the extreme case in which no adjustments are feasible after demand shocks are reveled; *i.e.,* no spot market transactions.

Assume that an LDC knows its inverse demand $D(q, w)$, and assume that the probability distribution of the weather and other demand shocks that could occur, $f(w)$, is common knowledge. Prior to w being revealed, the LDC knows that it could have one of many different demand functions based on different draws of w, as illustrated in figure 1. For the reasons just discussed, assume $D_q(q, w) < 0$, $D_{qq}(q, w) > 0$, and $D_w(q, w) > 0$, where the subscripts denote derivatives.

Finally, we assume that if $D(q_1, w_1) = D(q_2, w_2)$ and $w_1 < w_2$, then $D_q(q_1, w_1) \geq D_q(q_2, w_2)$. That is, for a given price, a positive demand shock (a larger w), causes demand to be not only greater, but also less elastic. This assumption is not entirely innocuous, but seems realistic. More importantly, the point of this analysis is to show the results are theoretically plausible, not that they need always hold. The empirical analysis will speak to their importance.

We assume that the LDC is risk neutral with respect to price. Thus, recognizing that it can make no adjustments in a spot market, the LDC would purchase or sell gas in the forward market according to the demand function $\bar{D}(q)$ that would be the weighted average demand function based on the probability distribution of $f(w)$, $\bar{D}(q) = \mathbf{E}_w[D(q, w)]$. That is, $\bar{D}(q)$ is the probability-weighted average of the marginal valuation of additional gas at a given q over all possible demand states. Since $\bar{D}(q)$ is a weighted average of convex functions, with non-negative weights, it follows that $\bar{D}(q)$ is convex.

Result 1: For any forward market price p, an LDC will choose to purchase a larger quantity in the forward market than it would in expectation if it faced that same p in the spot market (*i.e.,* after w is revealed).

That is, the quantity that equates the expected marginal value of gas with p before w is revealed is greater than the expected quantity that equates the marginal value of gas with p after w is revealed. The result follows immediately from the convexity of $D(q, w)$ in q and the assumption that a positive weather shock decreases elasticity at a given price. Essentially, the value of holding extra supplies of gas before w is revealed is large because the cost of a shortage is much greater than the cost of having excess supplies at delivery time.

While $\bar{D}(q)$ would be an LDC's demand for gas if it had no reserves going into the forward market, each LDC actually starts with an endowment of gas reserves, q_{res}, for which it has already contracted. Figure 2 presents the net demand position of an LDC, where it would supply quantity to the market if p is above its marginal forward valuation at its current reserve level, denoted by $p_0 = \bar{D}(q_{res})$, and it would purchase from the market if $p < p_0$. The supply and demand functions of all market participants aggregate into the market supply and demand. Because of the enhanced value of reserves in the forward market, result 1 indicates that each participant's forward market valuation curve lies to the right of and above its expected spot market valuation. As a result, in all but one special case, the forward market equilibrium price is above the expected spot market equilibrium price that would result if transactions occurred following the revelation of w.

Result 2: If the demand shocks faced by the LDCs in the market are not perfectly correlated, then the forward price, p_f, will be greater than the expected spot price $\bar{p}$.

Proof: TBA.

In reality, of course, a spot market does exist and some transactions occur after w is revealed. There is some risk, however, that a firm will not be able to access the spot market for some trades due to lack of market liquidity. Assume that the probability a firm will be unable to access the spot market is ϕ.

If the firm could access the spot market with certainty ($\phi = 0$), then, because it is risk neutral with respect to price, it would be unwilling to buy in the forward market at any price above the expected spot price, which we will call $\bar{p}$, or sell at any price below the expected spot price. Its forward market demand curve would be flat at $\bar{p}$. If, however, a firm has some risk of failing to make a spot purchase, that failure has an expected cost in lost value equal to the value of the incremental reserves, $\bar{D}(q)$. Thus, a firm expecting a spot price of $\bar{p}$ will have a willingness to pay in the forward market of $\hat{D}(q) = (1-\phi)\cdot\bar{p}+\phi\cdot\bar{D}(q)$. The liquidity risk in the spot market means that the firm's marginal willingness to buy or sell in the forward market is a weighted average of its marginal forward valuation of reserves (with no access to a spot market) and the expected spot price. This is illustrated in figure 3. Access to a spot market with probability $1 - \phi$ reduces the magnitude of the forward valuation premium, but does not eliminate the effect identified in Result 2.

Since the forward premium is a function of the convexity of $D(q, w)$, it will be more pronounced at times when valuations are more convex. In the gas market, this corresponds to times when the market is relatively tight, and a significant share of market participants face possible shortages that would necessitate reductions in end-use demand through non-market mechanisms, causing significantly higher marginal valuations of gas. In contrast, if all LDCs have comfortable reserve margins and the distribution of possible demand outcomes includes little risk of a sufficiently large positive shock, then the marginal valuation in the market reflects storage costs and changes fairly little with quantity. This is illustrated in figures 4a and 4b, which contrast a tight-market and a loose-market scenario. Even in the loose-market scenario, there may be some inframarginal demand with a high marginal valuation of supplies—because some LDC needs to buy gas in order to meet its end-use demand—but the *marginal* supply and demand are both nearly flat around the equilibrium price. Thus, the forward market demand curve may be significantly above the expected spot price over some range, but not at the equilibrium quantity, so the forward premium disappears.

It is worth noting that the asymmetry between buyers and sellers creates a forward premium when markets are tight due to concerns about security of supply, but does not create a forward *discount* when markets are loose due to concerns about "security of demand." It is indeed the case that in loose markets a reserve seller might be unable to find a buyer, but the loss from failing to complete a transaction will be minimal. While an LDC that is short of reserves may be desperate to procure supplies in order to prevent customer curtailments and thus has a very high $D(q, w)$ for the reserves it wishes to buy, an LDC that has excess reserves will not be desperate to sell. Its downside risk is limited to simply forgoing potential revenue from selling the reserves at the market price. In most cases, it could simply carry the reserves forward to the next period or, in the case of transmission, allow the transmission right to expire unused. As a result, the supply curve is relatively elastic, implying small surplus gains from making the sale. In addition, because sellers are also mostly regulated LDCs who pass through their net gas procurement costs to end-use customers, the financial risk for the LDC itself is limited even further.

The forward premium in a tight-market scenario at first appears to create an opportunity for speculators to engage in a risky arbitrage that is profitable in expectation. However, the same liquidity risk that creates the premium also undermines the opportunity for profitable arbitrage. A speculator would take advantage of the forward premium by selling gas in the forward market, and then buying those supplies from the spot market for delivery to the buyer from the forward market. Like other market participants, however, the speculator would face the liquidity risk that it could not obtain the supplies from the spot market that it had committed to deliver. In the event of a failure to deliver, the speculator would be liable for the damage imposed on the buyer of the forward gas, which would presumably reflect the buyer's realized $D(q, w)$. Thus, the speculator would be in no better position to take advantage of this premium than other market participants.

While this model results in a forward premium for gas, it does not yet capture one of the institutional aspects of the market that has been expressed or confirmed to us by all of the participants we have interviewed. That is, the regulatory oversight in the retail natural gas market results in especially harsh treatment of an LDC that *sells* supplies it had in reserve and then finds itself in a shortage situation. To be clear, this is an additional cost beyond $D(q, w)$, the value of the marginal quantity, that is incurred only if the firm was a seller in the forward market and then found it necessary to engage in some retail quantity reduction due to a shortage of supplies and an inability to purchase in the spot market. Such punitive regulatory response does not occur if the LDC made a "good faith effort" to procure the necessary supplies in the forward and spot market, but it is likely to be significant if the LDC is perceived as having attempted to "profiteer" by selling needed reserves in the forward market in order to take advantage of a high forward price.

We represent this feature of the market with a penalty function that raises an LDC's cost of selling reserves in the forward market if there is a risk it will then have to procure supplies in the spot market in order to meet its own demand. The expected penalty, and associated cost of selling in the forward market, is greater if the LDC is more likely to have to buy quantities back in the spot market. This likelihood is a decreasing function of its starting reserve position and an increasing function of the quantity of gas it sells in the forward market. We illustrate two scenarios in figure 5, where in each case $\hat{D}(q)$ represents the willingness to buy and sell in the forward market absent the regulatory penalty and $\tilde{D}(q)$ is the willingness to buy and sell in the forward market inclusive of the regulatory penalty. In 5a, the LDC has plentiful supplies and could sell substantial quantities in the forward market before the expected value of the penalty would become non-negligible. In 5b, however, the LDC is in a tighter reserve position and more likely to need to purchase supplies in the spot market. In this case, selling even a small quantity in the forward market creates a significant probability that the LDC will incur regulatory punishment. Note that this creates a discontinuity in the marginal value of reserves at the LDC's starting reserve position.

The obvious effect of this regulatory penalty is to reinforce the convexity of the forward marginal value function and result in a larger forward price premium when supplies are expected to be tight. A second effect is on volume in the forward market. The model without such a regulatory penalty has no prediction for the volume of trade in the forward market. However, the asymmetry of the regulatory penalty—among all the LDCs with insufficient supplies in the spot market, it affects only those who sold gas in the forward market—acts to discourage efficient trade in the forward market when markets are tight. A sufficiently large penalty could eliminate all trade in the forward market at times when LDCs perceive a real risk of a severe demand shock and tight gas market.

One might wonder why the regulator would impose a penalty threat that causes the LDC to forego efficient trades in the forward market and instead engage in what might be termed "cover your ass" behavior. In the appendix, we show that, in a principal-agent framework in which the principal is uncertain of the agent's ability to trade efficiently on its behalf, an optimizing principal may want to impose a penalty on *ex post* unfavorable outcomes of an action, but not of an inaction. The principal may prefer to do so even though the policy would in equilibrium lead to excessive "ass covering." This is not to say that the regulatory penalty LDCs claimed they face is necessarily optimal, but that it is also not necessarily irrational.

3.3. Price Risk Aversion

Thus far, we have assumed that the market participants are risk neutral with respect to price. However, price risk aversion offers a competing explanation as to why forward transaction prices may differ from expected spot prices. If buyers and sellers are not risk neutral, then the price risk they face will cause them to be willing to give up some expected surplus in a transaction in order to reduce that risk. If buyers are systematically more willing to pay to reduce price uncertainty than are sellers, then the observed forward transaction prices will be greater on average than the associated spot prices.[7]

Risk aversion over prices, however, is not as compelling an explanation for a forward price premium in these markets as it might at first appear. First, recall that buyers and sellers in downstream gas markets and pipeline capacity markets are generally the same firms. A holder of a certain quantity of long-term gas delivery contracts or long-term pipeline capacity rights will sometimes find itself in need of more, and thus buy on the local gas or pipeline capacity release market, and at other times find itself with excess supply and thus sell on the gas or capacity release markets. It is difficult to argue that buyers are systematically more risk averse to price volatility than sellers when the same set of players occupy both sides of the market at various times—in our final capacity release dataset, at least 79% of the sellers are LDCs.[8] In addition, the LDCs that are buying and selling in this market are allowed to pass through gas acquisition costs to their customers, so we would not expect them to exhibit much price risk aversion.

Empirically, a price risk aversion explanation also has somewhat different implications than the quantity risk model we have proposed. The first test of our quantity risk theory does not allow us to rule out price risk aversion: one would also expect to see a greater forward premium when the market is expected to be tight because that is when prices

---

[7] See Dusak (1973), Bodie and Rosansky (1980) and Carter, Rausser and Schmitz (1983) for a more thorough discussion of risk premia in efficient and liquid commodity markets.

[8] We say "at least" because among the other 21% of the sellers are gas marketing companies, who may be selling on behalf of their client LDC—thus the "true" fraction of sellers in this market that are LDCs may be even higher than 79%.

are volatile. However, a second test does distinguish between the theories empirically. Our quantity risk theory with the regulatory penalty just discussed predicts that trading volume will be low when the market is tight. A price risk aversion explanation would predict the opposite. When prices are volatile the gains from trade between more and less price risk averse traders are greatest, so a price risk aversion story will predict trade volume to be highest at such times.

4. Evidence from markets for natural gas

The theoretical model presented above implies that forward price premia for natural gas should develop when markets are tight, and should increase with market tightness. Here, we test this implication by examining panel data on forward and spot natural gas prices across North America. We then follow this in Section 5 with an examination of data regarding forward and spot markets for natural gas transportation.

4.1 Data

All data used in this project were obtained from Platts' GASdat product, and consist of observations covering the three markets of interest in this paper: the day ahead (spot) market, the forward month (bidweek) market, and the capacity release market. The spot and bidweek markets are both markets for gas at specific locations. Platts obtains pricing data via surveys of trades made at each pricing location, and reports the average location-specific spot price at daily intervals. Bidweek data occur on a monthly basis and represent the volume-weighted average price of all surveyed trades that occur at each location during bidweek, which is the last five trading days of each month.[9]

To make the daily spot data compatible with the monthly bidweek data, we average the daily spot observations within each location and month to obtain an average spot price for the month. We proceed at this level of aggregation for the remainder of the paper. There exist 8,433 location-months for which both spot and bidweek prices exist,[10] spread over 117 locations, across which the duration of coverage varies. For example, while data for Henry Hub in Louisiana span 1991 to 2005, data at the Carthage Hub in northeast Texas are only available for 1997-2002. These variations in data availability occur because

---

[9] Platts will sometimes use the median of reported prices if it finds that one high-volume transaction skews its bidweek sample. Unfortunately, there are no indicators in the data to determine which observations are computed in this way.

[10] The count of 8,433 location-months includes only those observations for which rolling regression spot price predictions can ultimately be generated, as discussed later. Without this restriction, there are 11,468 location-months.

trading activity in some locations varies over time, and Platts does not record observations when there are an insufficient number of trades to allow it to determine the average price.

Summary statistics for the spot and bidweek prices are shown in table 1. Both data series are highly right-skewed, as indicated by the excess of the mean over the median prices, and by the large observed maximum prices. The summary statistics of the spot and bidweek prices are very similar, and the difference in means of 4.6 cents is not statistically significant.[11] Thus, on average, there is no forward premium or discount in prices for natural gas. This result, however, does not speak to whether forward premia develop when the market is expected to be tight—a topic to which we now turn.

4.2 Empirical framework

Our working hypothesis is that the forward price premium in the market for natural gas should rise with the expected tightness of the market. A natural measure of expected market tightness is the expected spot price, since natural gas is most valuable precisely when access to it is difficult. We therefore test our hypothesis by estimating the parameters of the following equation.

$$BidWeek_{it} - Spot_{it} = \beta_0 + \beta_1 \ E[Spot_{it}] + \mu_i + g(t) + \epsilon_{it}, \qquad [2]$$

Here, $BidWeek_{it}$ and $Spot_{it}$ are the forward and spot prices for gas at location i at month t, and $E[Spot_{it}]$ is the expectation at the beginning of month $t-1$ (before the bidweek or spot markets occur) of the spot price at location i at month t. The μ_i are location fixed effects, and $g(t)$ is a fourth order polynomial in time that controls for the secular upward trend in natural gas prices. If security of supply concerns do indeed make gas buyers quantity risk averse, then β_1 will be positive, meaning that bidweek prices exceed realized spot prices on average when expected spot prices are high.

Estimation of equation 2 requires a measure of $E[Spot_{it}]$. Because we do not observe market participants' expectations of spot prices for a particular month, we construct these expectations using historical spot prices and the model below.

$$E[\ln(Spot_{it})] = a_0 + a_1 \ \ln(Spot_{i,t-12}) + a_2 \left[\ln(Spot_{i,t-2}) - \ln(Spot_{i,t-14})\right] + d_i \qquad [3]$$

[11] Statistical significance was tested for via a paired t-test with standard errors clustered on year-month: the t-statistic is 0.90.

In words, we model the expected spot price for location i in month t as a function of the spot price for the same calendar month of the last year, the change in spot price between fourteen months and two months before month t, and a location fixed effect. We use the two-to-fourteen month difference, rather than one to thirteen, in order to avoid using data from month $t-1$.[12] The model is run in logarithms rather than absolute price levels because the price data are right skewed and never negative.

To use equation 3 to predict expected spot prices, we must first estimate the equation's parameters—a_0, a_1, a_2, and the d_i—by running the regression specified in equation 4, below. This equation includes an unobserved orthogonal disturbance ν_{it} to account for information revealed between month $t-2$ and month t.

$$\ln(Spot_{it}) = \alpha_0 + \alpha_1 \ln(Spot_{i,t-12}) + \alpha_2 \left[\ln(Spot_{i,t-2}) - \ln(Spot_{i,t-14})\right] + \delta_i + \nu_{it} \qquad [4]$$

In the process of estimating equation 4, and then generating forecasts using equation 3, we take care to avoid using any future information in our forecasts. That is, when we forecast the expected spot price for month t using equation 3, we only use parameters estimated using information from prior to month t. This means that we do not use our entire sample of spot price information to produce estimates of the α's and δ_i's from equation 4, and then apply these estimated parameters to generate a full time series of expected prices from equation 3.

Instead, we estimate equation 4 using a "rolling regression" approach. Rather than estimate a single set of a's, we estimate a different coefficient vector $\hat{\alpha}_t = (\hat{\alpha}_{0t}, \hat{\alpha}_{1t}, \hat{\alpha}_{2t})$ for each month t using data only up to month t. These coefficients are then substituted for the corresponding a's in equation 3 to generate expected spot prices for month $t+2$. While this approach ensures that our spot price prediction for any month t does not include any information revealed after t, it does come with the cost that there are few data with which to estimate equation 4 in the early part of our sample. To avoid generating estimates based upon only a handful of points, we predict spot prices only for locations and months for which we have at least 38 months of spot price history. This is the shortest history that guarantees the use of at least 24 observations in the estimation of equation 4, when one of the regressors has a 14 month lag.

[12] Because we ultimately aim to use the predicted spot price for month t to explain bidweek prices in month $t-1$, we wish to avoid incorporating any information made available over the course of month $t-1$ in the prediction.

4.3 Results

We first report results from the estimation of equation 4—the first stage of developing *ex-ante* expectations of spot prices. The first column of table 2 reports results using the full sample of spot price data. Our two primary determinants of spot price are the one year lagged spot price and the "trend" between the 14 month lagged price and the 2 month lagged price (all in logs). The estimated coefficients of these variables are positive, as expected, and strongly significant with either clustered or OLS standard errors. Lagged prices clearly carry useful information with which to predict current prices.

Table 2 also reports the results of rolling regressions of equation 4; the estimated parameters of these regressions are ultimately used to generate expected spot prices using equation 3.[13] We ran 151 rolling regressions in order to generate parameters for the estimation of spot margins from March 1993 to November 2005. The estimated coefficients on the one year lagged price and the trend are positive for each individual regression.

Results from the estimation of our main specification—equation 2—are reported in column I of table 3. We estimate that the forward premium of bidweek prices over spot prices increases systematically with the expected spot price: a \$1.00 rise in the expected spot price is expected to cause a \$0.12 rise in the forward premium. This effect is statistically significant with either clustered or OLS standard errors, though the significance level is only 10% when clustered on year-month—indicative of strong cross-sectional correlation in the residuals and expected prices. These results are robust to the addition of interactions between the location fixed effects and the time polynomial to the specification, as shown in column II.

We further investigate these results by unpooling the specification, so that the effect is estimated separately for locations in the northeastern U.S. We do so because the Northeast, a natural gas consumption region that is distant from any major gas producing area, is highly susceptible to weather-driven demand shocks that can test the constraints of the transportation and storage infrastructure. Security of supply concerns should therefore be particularly salient in the Northeast, a common intuition that is supported by the unpooled

[13] Because equation 3 only generates the expected logarithm of the spot price, but equation 2 requires an expected absolute price, we calculate $E[Spot_{it}] = \exp(E[\ln(Spot_{it})] + \sigma^2/2)$, where σ^2 is the variance of the residuals from the estimation of equation 4. This calculation relies on the homoskedasticity and normality of the disturbance ν_{it}. To test the calculation's accuracy, we use it to generate expected absolute prices via the full-sample parameters from table 2, and then regress actual absolute prices on these expected absolute prices. The results show that the expected absolute prices are valid predictions of actual prices: the estimated slope coefficient of this regression is statistically indistinguishable from one, and the estimated constant is statistically indistinguishable from zero.

results presented in column III of table 3. A $1.00 increase in the expected spot price in the Northeast drives an $0.18 increase in the forward premium there, on average, whereas the same expected spot price increase elsewhere only increases the forward premium by $0.11. Similar effects are found when location-time interactions are added to the specification, as indicated in column IV. As with the pooled estimation, however, clustering the standard errors on year-month reduces the statistical significance of these results, in some cases to below the 10% level.

Consistent with the model of security of supply, data from natural gas markets indicate the presence of forward price premia at times and locations when the market is expected to be tight. We now turn to exploring whether similar evidence is found in the market for natural gas transportation, asking whether prices for forward transmission rights exhibit premia over spot valuations of transportation in tight markets.

5. Evidence from markets for natural gas transportation

5.1 Data

In addition to data on the prices for *natural gas* at locations around the country, the Platts dataset also contains information and pricing for the *capacity release* market, the explicit market for pipeline transportation rights. In these data, each observation represents a transaction in which what is sold are capacity rights for a fixed time period on a specific pipeline route. Unlike the spot and bidweek datasets described above, these data are not based on surveys; they are a comprehensive set of all releases that have occurred on each pipeline since Platts began tracking them.[14] Each observation provides information regarding price, the transaction date, the duration of the release, and the route along which capacity is released (many pipelines have multiple routes either because they have a branched structure or because there exist receipt and delivery points at intermediate locations along the line).

Many capacity release observations are dropped from the dataset for institutional reasons. We drop releases that occur between companies that are affiliates, as well as releases that are recallable. In a recallable release the releasing firm has the right to recall its capacity from the awarded firm after a contractually agreed notice period, which can be as brief as 24 hours. Because we do not observe the specific contractual terms governing recall, we drop all observations for which any form of recall rights are granted. Further,

[14] The year in which observations begin varies by pipeline, but is generally in the late 1990s.

an unusually large share of transactions (greater than 50% on most pipelines) are priced at the FERC's maximum tariff rate for the released capacity, because doing so allows the parties to renew the transaction after the expiry of the initial contract term without making the capacity available for bidding by other firms. Because such releases are effectively bundles of both current and future rights, we drop these observations as well. For further standardization, we also drop the small number of releases that are for terms other than one calendar month, and those releases for which the transaction occurs more than one month ahead of the effective month of the release.

To facilitate comparisons to spot data, we define a capacity release *product* as a pipeline route-month combination for which capacity may be released. For example, capacity on Transcontinental Pipeline's route from Louisiana to New York for October 2003 constitutes a product. For each product for which the necessary data are available, we compute a *spot margin* as the spot price at the downstream node of the route during that month minus the sum of the spot price at the upstream node of the route during that month and a small variable cost of pipeline transportation (which must be paid to the pipeline by any user of capacity rights). This margin represents the profit that a firm would make, absent a capacity payment, by shipping gas on the released route over the contract term, and is therefore an ex-post measure of the value of a capacity release for that product. Each product is associated with exactly one spot margin; however, for any given product we may observe zero, one, or multiple capacity release transactions, each with a potentially different capacity release price.

The final capacity release dataset contains 459 transactions that can be matched to spot margins (and to predicted spot margins, as discussed below), spread over 150 products on 14 routes on five major interstate pipelines: Columbia Gulf, Florida Gas Transmission, Texas Eastern, Texas Gas Transmission, and Transcontinental.[15] The geographic distribution of these pipelines is indicated in the map in Figure 6, while a summary of the number of observations by pipe is given in table 4. While all five pipelines receive gas in the Texas/Lousiana production basin, their delivery points vary from New York to Appalachia to Florida.

Summary statistics for the capacity release prices (CR prices) and spot margins observed for the 14 routes in our data are indicated in table 5. The distributions of CR prices and spot margins are significantly right-skewed, consistent with a gas pipeline infrastructure that has spare capacity at most times, but occasionally does become constrained. It

[15] While capacity release data were available for several other major pipelines, no more than thirty transactions were observed for these lines, and they were therefore dropped.

is during these constrained periods that possession of firm transportation rights can be extremely valuable. In addition, the mean CR price is statistically indistinguishable from the mean spot margin (with any construction of standard errors), consistent with a lack of forward price premium under most market conditions.

5.2 Empirical framework

Our empirical strategy for the natural gas transportation market mirrors that used earlier in the market for natural gas itself. We test the theoretical prediction that, when markets are expected to be tight, forward purchases of firm transportation should carry a price premium over the spot valuation of that transportation, on average.

In our tests involving markets for natural gas, we used expected spot prices as our measure of expected market tightness. In parallel fashion, we now use expected spot margins, and test for security of supply by estimating equation 5, which is adapted from equation 2.

$$CapRelease_{ijt} - SpotMargin_{it} = \beta_0 + \beta_1 \; E[SpotMargin_{it}] + \mu_i + \epsilon_{ijt}, \qquad [5]$$

Here, $CapRelease_{ijt}$ is the price paid in transaction j for a capacity release on route i during month t (the price is set in month $t-1$). $SpotMargin_{it}$ is the realized spot margin over route i during month t, and $E[SpotMargin_{it}]$ is the expectation of this spot margin at the beginning of month $t-1$. The μ_i are route fixed-effects. The base specification does not include a polynomial in year-month, because margins do not follow a secular trend; however, we do include this polynomial in specification checks. As with equation 2, the theoretical model of security of supply predicts that β_1 will be positive, so that CR prices will carry a price premium over spot margins on average when spot margins are expected to be high.

To construct expected spot margins, we re-use the procedures developed for constructing expected spot prices. Specifically, we predict spot margins using equation 6, which in turn relies on estimates from rolling regressions of equation 7. We use absolute spot margins here rather their logarithms because margins can be negative.

$$\begin{aligned} E[SpotMargin_{it}] =& a_0 + a_1 \, SpotMargin_{i,t-12} \\ & + a_2 \, [SpotMargin_{i,t-2} - SpotMargin_{i,t-14}] + d_i \end{aligned} \qquad [6]$$

$$
\begin{aligned}
SpotMargin_{it} =& \alpha_0 + \alpha_1\, SpotMargin_{i,t-12} \\
& + \alpha_2 \left[SpotMargin_{i,t-2} - SpotMargin_{i,t-14} \right] + \delta_i + \nu_{it}
\end{aligned} \qquad [7]
$$

5.3 Results

As with our discussion of results in local natural gas markets, we begin here with an exposition of results regarding our estimation of spot margins. The first column of table 6 shows the results of the estimation of equation 7 using the full sample of spot margin data. The estimated coefficient on the one-year lagged spot margin is positive and statistically significant over all constructions of standard errors, while the estimated coefficient on the "trend" is positive but not statistically significant when the standard errors are clustered on year-month. Furthermore, over the 67 rolling regressions on the spot margin data, both coefficients are positive on average, though the coefficient on the trend is negative for 25 of the 67 regressions. This lack of precision in estimating the trend coefficient likely reflects the small size of the spot margin dataset relative to the spot price dataset discussed earlier, as well as a relatively low month-to-month persistence of spot margin shocks.[16]

Column I of table 7 reports the results from the estimation of the primary specification, equation 5. The data support the existence of security of supply concerns in natural gas transportation: the premium of CR prices over spot margins increases with the expected spot margin. A \$1.00 increase in expected spot margin is expected to increase the forward premium by \$0.21, an effect that is statistically significant across all three constructions of the standard error. The increase in standard error caused by clustering on year-month here is less dramatic than was the case in the regressions using data on markets for natural gas itself, reflecting spatial heterogeneity in pipeline capacity constraints. The statistical significance of the estimate is robust to the addition of a polynomial in year-months and interactions of this polynomial with route fixed effects, as shown in columns II and III.

In parallel with our discussion of the gas price results in Section 4.3, we unpool our sample here to separately estimate equation 5 for routes serving the Northeast (all such routes occur on Texas Eastern Pipeline and Transcontinental Pipeline). The results presented in columns IV through VI of table 7 demonstrate that evidence supporting security of supply concerns in gas transportation markets is confined to the Northeast. This reflects the

[16] While the results reported in table 7 rely on predicted spot margins that were generated using the trend variable, we verify the robustness of these results to the use of predicted spot margins that were generated from rolling regressions that exclude the trend. In fact, the use of such predicted margins enhances the statistical significance of the results relative to that shown in table 7.

relative incidence of capacity constraints in transportation to the Northeast, as compared to transporation to Florida or Appalachia via the other pipelines in our sample.

6. Supplementary evidence from capacity release markets

While the primary empirical prediction that we took away from our theoretical discussion in Section 3 was that forward prices should exceed spot prices when capacity is expected to be scarce, a secondary prediction is that the number of transactions should also decrease when capacity is expected to be scarce. This occurs because the number of potential sellers decreases as fewer firms are *ex ante* certain that they have adequate gas to meet their maximum possible demand. Recall that a price risk aversion explanation for forward premia in natural gas markets predicts the opposite result: when prices are volatile the gains from trade between more and less price risk averse traders are greatest, so trade volume should be highest at these times.

A preliminary graphical analysis provides useful evidence supporting this secondary prediction from the security of supply model. Figure 7 presents a plot of expected spot margins and the number of observed capacity releases on the five pipelines in our sample, averaged by month. Fewer capacity releases are observed in the winter, when expected spot margins are generally high, than in the summer. This is more consistent with a security of supply model in which firms are relucant to make forward sales of capacity, than with a price risk aversion model.

To investigate the occurrence of capacity releases more formally, table 8 presents the results from estimating a Poisson count model of the *number* of capacity release transactions observed for a particular product (route-month) as a function of the expected spot margin and month fixed effects (route fixed effects are also included but unreported in the table). That is, we investigate whether increases in expected spot margins drive decreases in capacity release transactions, even after conditioning on month. Results are reported as incidence rate ratios (IRR's)[17] and the associated confidence intervals, with standard errors clustered on year-month. This clustering yields wider confidence intervals than either OLS standard errors or clustering on route, and also takes over- or under-dispersion into account. The first column reports results for all products for which we observe capacity

[17] Each incidence rate ratio (IRR) should be interpreted as the proportional change in the rate of occurrence of capacity release transactions as the result of a one unit increase in the associated explanatory variable. The “null” effect is an IRR of 1; an IRR of more than one indicates that the explanatory variable tends to increase the rate of the event being counted, and an IRR of less than one indicates the opposite. For example, an IRR of 0.4 indicates that a one unit increase in the independent variable causes the expected count to decrease by 60%. The IRR has a log-normal distribution, which is why we report confidence intervals instead of standard errors.

releases (150 of these, as noted in table 4), as well as those products for which we do not observe a release, yet occur on a route for which releases have been observed in prior periods (259 of these, for a total of 409 products). A comparison of the IRR's estimated for the month fixed effects indicates that the months of April-November have approximately 3-4 times as many transactions as do the months of December-March, consistent with Figure 7. Furthermore, the estimated IRR on the rolling regression spot margin prediction indicates that a \$1.00 increase in the expected spot margin will cause a product to have only 32.6% of the number of transactions as an otherwise identical product. Though the 95% confidence interval around this point estimate includes one, the results from this first regression provide further evidence that capacity release transactions are less frequent when expected spot margins are high.

To explore this result further, we break our sample into two periods: (1) September 2000 - September 2002, and (2) all other times. During September 2000 - September 2002, the FERC waived rate caps on the capacity release market. Intuition suggests that these caps act as a barrier to releases when markets are tight and expected spot margins are higher than than the caps: potential sellers of capacity have no incentive to sell because they cannot obtain market value for their capacity. This hypothesized difference in market behavior is apparent in the second and third sets of results presented in table 8. During the waiver period, a \$1.00 increase in the expected spot margin caused, on average, a 58% reduction in the number of capacity releases, while the effect was a 92% reduction outside the waiver period. Furthermore, separating the sample in this way increases the precision of the estimated IRR's on the expected margin: the 95% confidence interval around this estimate within the waiver period does not include one, and, outside the waiver period, extends only to 1.05. These estimated effects, which exist even after stripping out seasonal variation, are inconsistent with a risk aversion model, but consistent with a model in which security of supply concerns are important, and there exist barriers to forward sales of capacity.

7. Conclusion

The natural gas industry (and we suspect other industries as well) involves firms that face security of supply concerns regarding the procurement of essential inputs. When spot markets for these inputs are thin—as is often the case in natural gas—then firms may be willing to pay a price premium to purchase inputs on forward rather than spot markets, so as to guard against stockouts. Because the natural gas regulatory environment constrains the forward supply of gas when demand is expected to be high, we show theoretically that forward prices will tend to exceed expected spot prices and volume in the forward market will tend to decline at such times. Using a dataset of forward and spot prices for natural gas and gas pipeline capacity, we find evidence that, while at most times there does not exist a forward price premium, forward prices do significantly exceed expected spot prices at times when natural gas capacity is expected to be constrained. Furthermore, trading volume at these times tends to decline. The implication is that, even in an industry where firms are not price risk averse, regulation and spot market illiquidity can combine to generate forward price premia for a critical input where the costs of shortages is much greater than the cost of excess inventories.

References

Bodie, Zvi and Victor I. Rosansky. "Risk and Return in Commodity Futures," *Financial Analysis Journal*, **36**(May/June 1980), pp. 27-39.

Carlton, Dennis W. "Contracts, Price Rigidity, and Market Equilibrium," *The Journal of Political Economy*, **87**(5) (Oct., 1979), pp. 1034-1062.

Carlton, Dennis W. "The Rigidity of Prices," *The American Economic Review*, **76**(4) (Sep., 1986), pp. 637-658.

Carter, Colin A., Gordon C. Rausser and Andrew Schmitz. "Efficient Asset Portfolios and the Theory of Normal Backwardation," *The Journal of Political Economy*, **91**(2) (April 1983), pp. 319-331.

Chambers, Marcus J. and Roy E. Bailey. "A Theory of Commodity Price Fluctuations," *The Journal of Political Economy*, **104**(5) (Oct., 1996), pp. 924-957.

Cremer, Helmuth, Farid Gasmi, and Jean-Jacques Laffont. "Access to Pipelines in Competitive Gas Markets," *Journal of Regulatory Economics*, **24**(1) (2003), pp. 5-33.

Deaton, Angus and Guy Laroque. "Competitive Storage and Commodity Price Dynamics," *The Journal of Political Economy*, **104**(5) (Oct., 1996), pp. 896-923.

Deaton, Angus and Guy Laroque. "On the Behaviour of Commodity Prices," *The Review of Economic Studies*, **59**(1) (Jan., 1992), pp. 1-23.

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

Finoff, David, Curtis Cramer, and Sherrill Shaffer. "The Financial and Operational Impacts of FERC Order 636 on the Interstate Natural Gas Pipeline Industry," *Journal of Regulatory Economics*, **25**(3) (2004), pp. 243-270.

Williams, Jeffrey C. and Brian D. Wright. *Storage and Commodity Markets.* New York: Cambridge University Press, 1991.

Figure 1

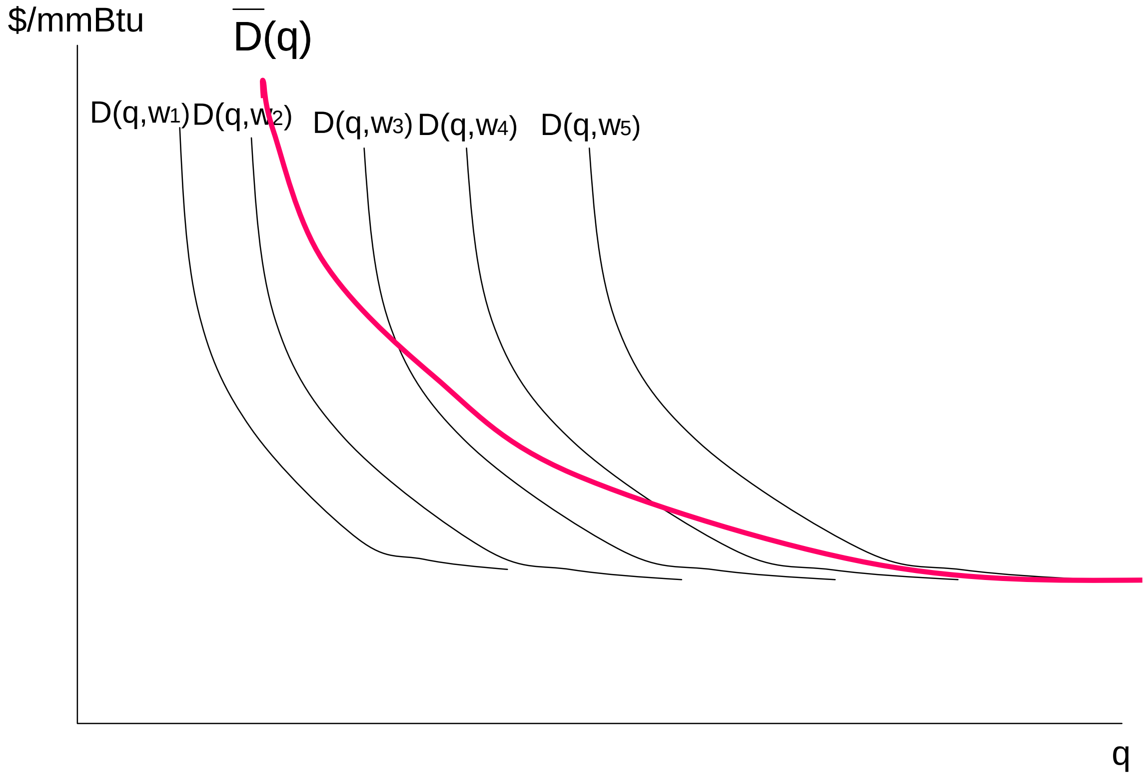

Figure 2

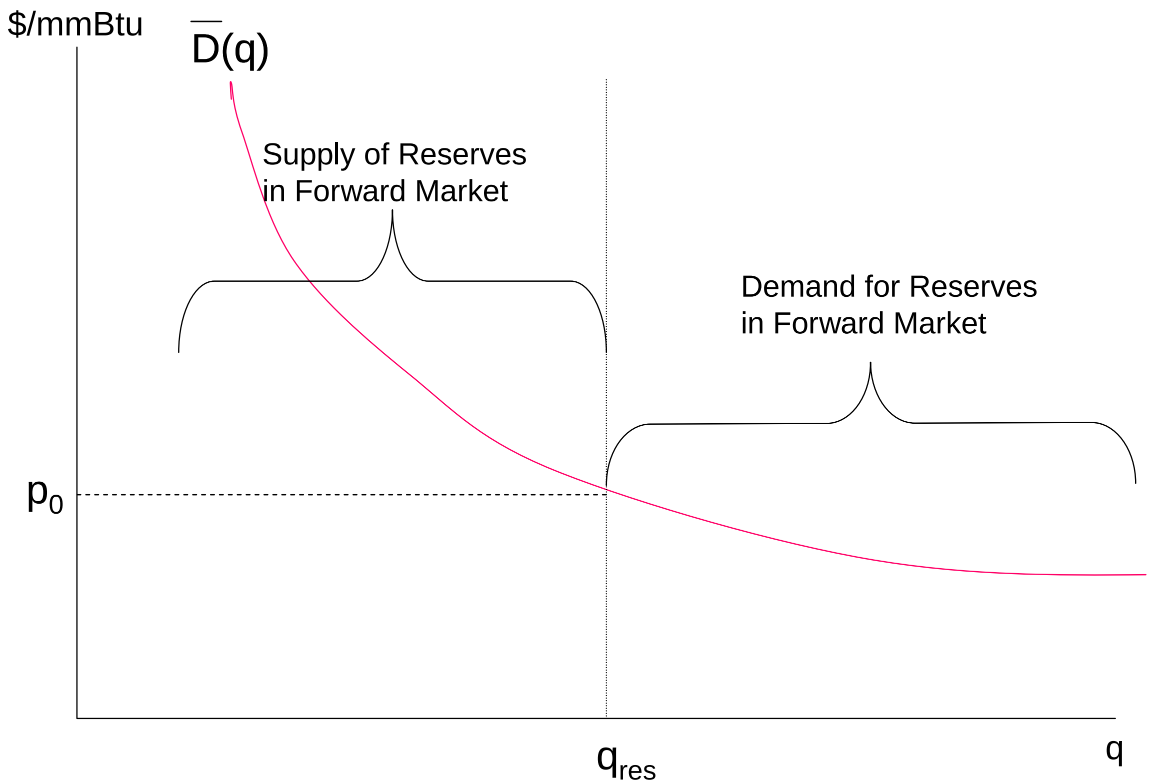

Figure 3

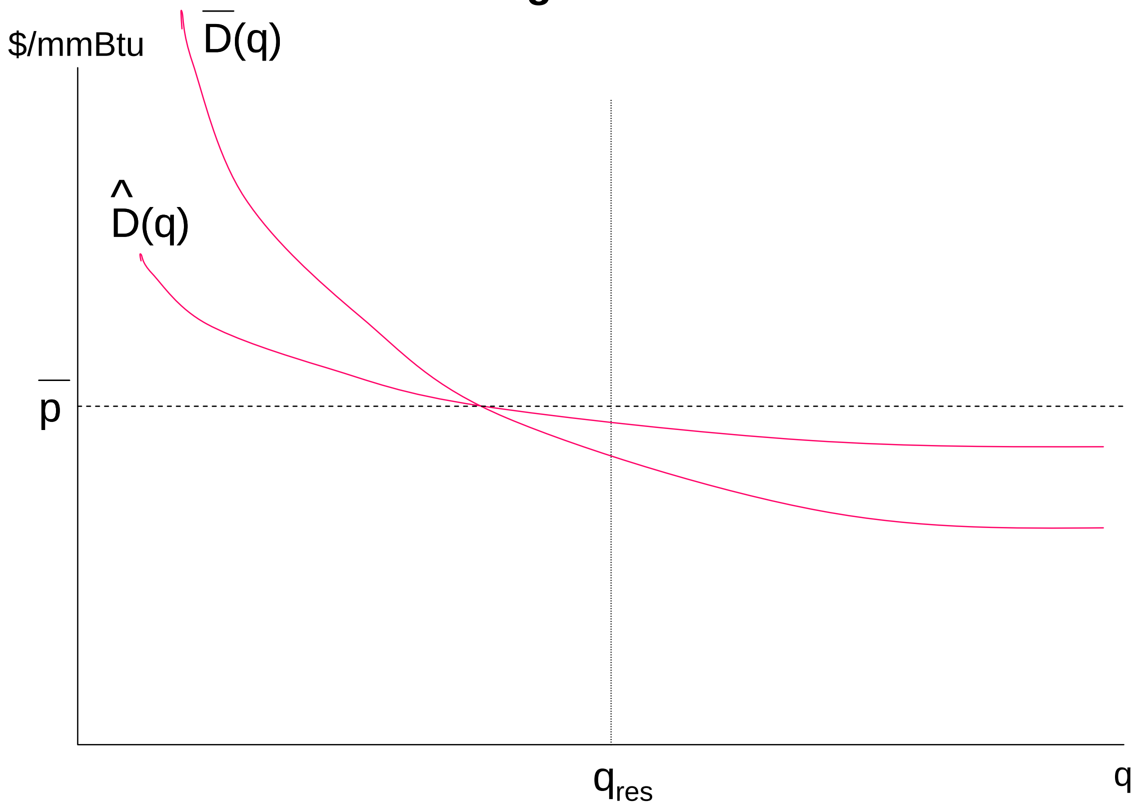

Figure 4

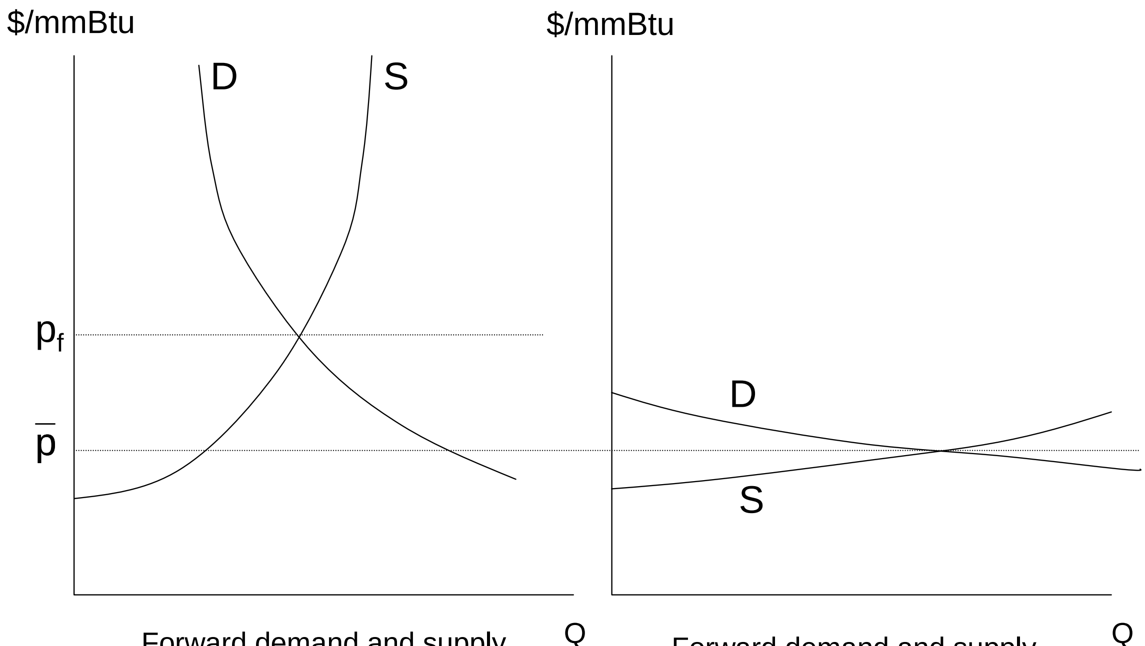

Forward demand and supply in a tight market

Forward demand and supply in a loose market

Figure 5

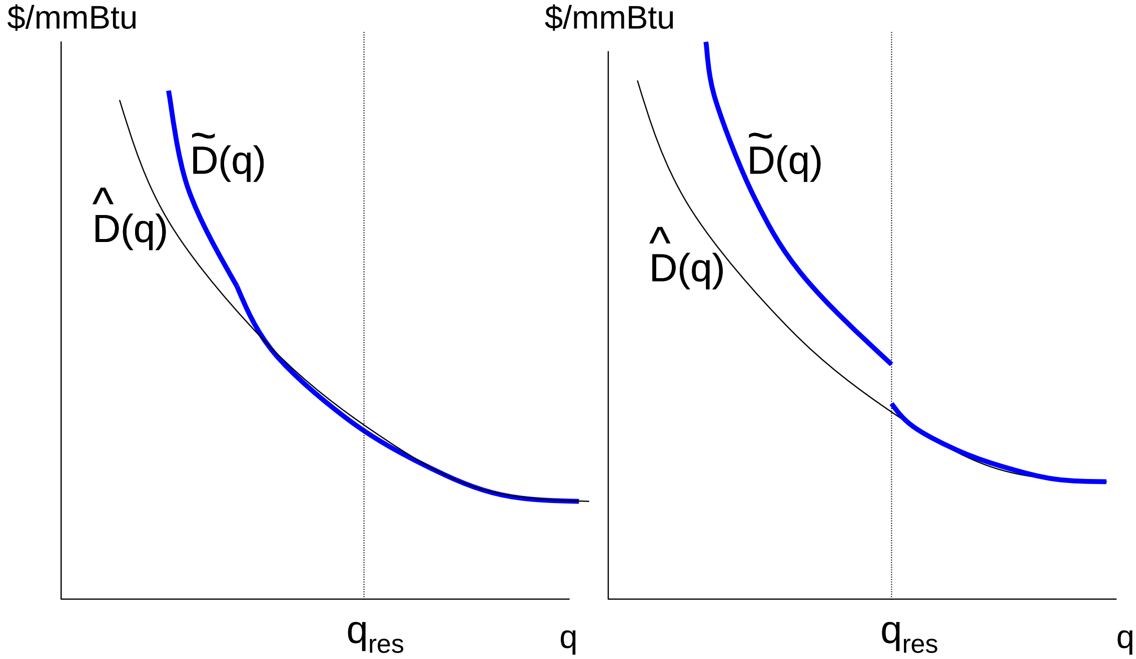

LDC with Extra Reserves — LDC with Possible Reserve Shortage

Table 1: Spot and Bidweek Price Summary Statistics

	Number of observations	Number of locations	Median	Mean	Std Dev	Min	Max
			Prices in $/mmbtu				
Spot Prices	8433	117	2.67	3.44	1.97	1.19	23.96
Bidweek Prices	8433	117	2.62	3.39	1.92	1.18	19.76

Table 2: Full Sample and Rolling Regression Results for Determinants of ln(Spot Price)

	Full Sample Results†	Mean Coef Over Valid RR Predictions (151 regressions)	Std Deviation of Coef Over Valid RR Predictions (151 regressions)
ln(12 month lagged price)	0.6615 (0.0067) (0.0138) (0.0554)	0.3197	0.1448
ln(2 month lag) minus ln(14 month lag)	0.6452 (0.0062) (0.0064) (0.0562)	0.5090	0.1085
N	11468	-	-
R^2	0.6808	-	-

Regressions include location fixed effects

†Standard errors listed from left to right are: OLS, clustered on location, and clustered on year-month

Table 3: Determinants of Bidweek Price Minus Spot Price

	I	II		III	IV
RR Prediction	0.1170 (0.0089) (0.0106) (0.0655)	0.1263 (0.0095) (0.0095) (0.0660)	RR Prediction, Northeast	0.1789 (0.0239) (0.0190) (0.1113)	0.2081 (0.0307) (0.0304) (0.1473)
			RR Prediction, Other	0.1096 (0.0093) (0.0104) (0.0635)	0.1176 (0.0100) (0.0086) (0.0612)
Location fixed effects	Y	Y	Location fixed effects	Y	Y
4th order polynomial in year-month	Y	Y	4th order polynomial in year-month	Y	Y
Location-time polynomial interactions	N	Y	Year-location interactions	N	Y
N	8433	8433	N	8433	8433
R^2	0.0395	0.0678	R^2	0.0404	0.0687

Standard errors listed from left to right are: OLS, clustered on location, and clustered on year-month

Figure 6

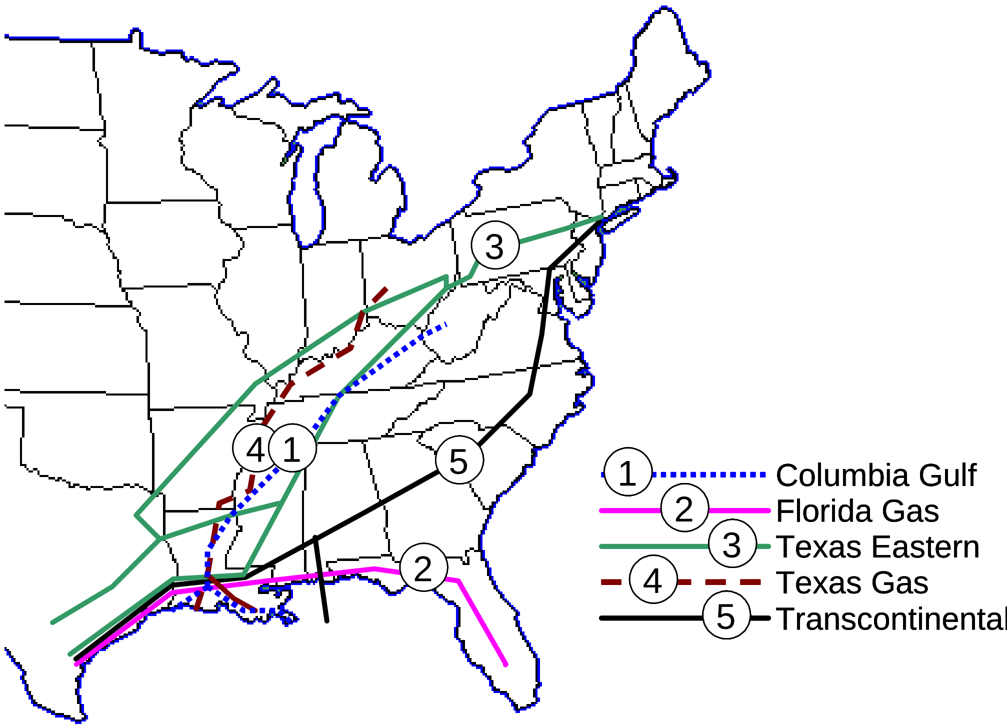

Table 4: Number of Observations per Pipe

Pipe Number	Pipeline	# of Routes	# of Cap Release Products	# of Cap Releases
1	Coumbia Gulf	2	47	209
2	Florida Gas	3	21	28
3	Texas Eastern	5	39	81
4	Texas Gas	2	15	32
5	Transcontinental	2	28	109
Total		14	150	459

Table 5: Summary Statistics of Transportation Data

		Prices in $/mmbtu				
	N	Median	Mean	Std Dev	Min	Max
Capacity Release Prices	459	0.050	0.082	0.115	0.000	0.816
Spot Margins	459	0.062	0.078	0.115	-0.115	1.986

Table 6: Full Sample and Rolling Regression Results for Determinants of Spot Margin

	Full Sample Results†	Mean Coef Over Valid RR Predictions (67 regressions)	Std Deviation of Coef Over Valid RR Predictions (67 regressions)
12 month lagged margin	0.4429 (0.0278) (0.0385) (0.1697)	0.3589	0.1676
2 month lag minus 14 month lag	0.0758 (0.0293) (0.0161) (0.0575)	0.0339	0.0634
N	1125	-	-
R^2	0.2988	-	-

Regressions include route fixed effects

†Standard errors listed from left to right are: OLS, clustered on location, and clustered on year-month

Table 7: Determinants of CR Price Minus Spot Margin

	I	II	III	IV	V	VI
RR Prediction	0.2149 (0.0606) (0.0391) (0.0548)	0.2551 (0.067) (0.0249) (0.0380)	0.2605 (0.0592) (0.0432) (0.0521)			
RR Prediction, Northeast Routes				0.2491 (0.0623) (0.0037) (0.0486)	0.2612 (0.0619) (0.0299) (0.0421)	0.2703 (0.0600) (0.0493) (0.0584)
RR Prediction, Other Routes				-0.3022 (0.2424) (0.2866) (0.3166)	0.1156 (0.2774) (0.2190) (0.3339)	-0.0927 (0.3603) (0.0968) (0.2850)
Location fixed effects	Y	Y	Y	Y	Y	Y
4th order polynomial in year-month	N	Y	Y	N	Y	Y
Location-time polynomial interactions	N	N	Y	N	N	Y
N	459	459	459	459	459	459
R^2	0.1552	0.1832	0.3049	0.1644	0.1837	0.3065

Standard errors listed from left to right are: OLS, clustered on route, and clustered on year-month

Figure 7

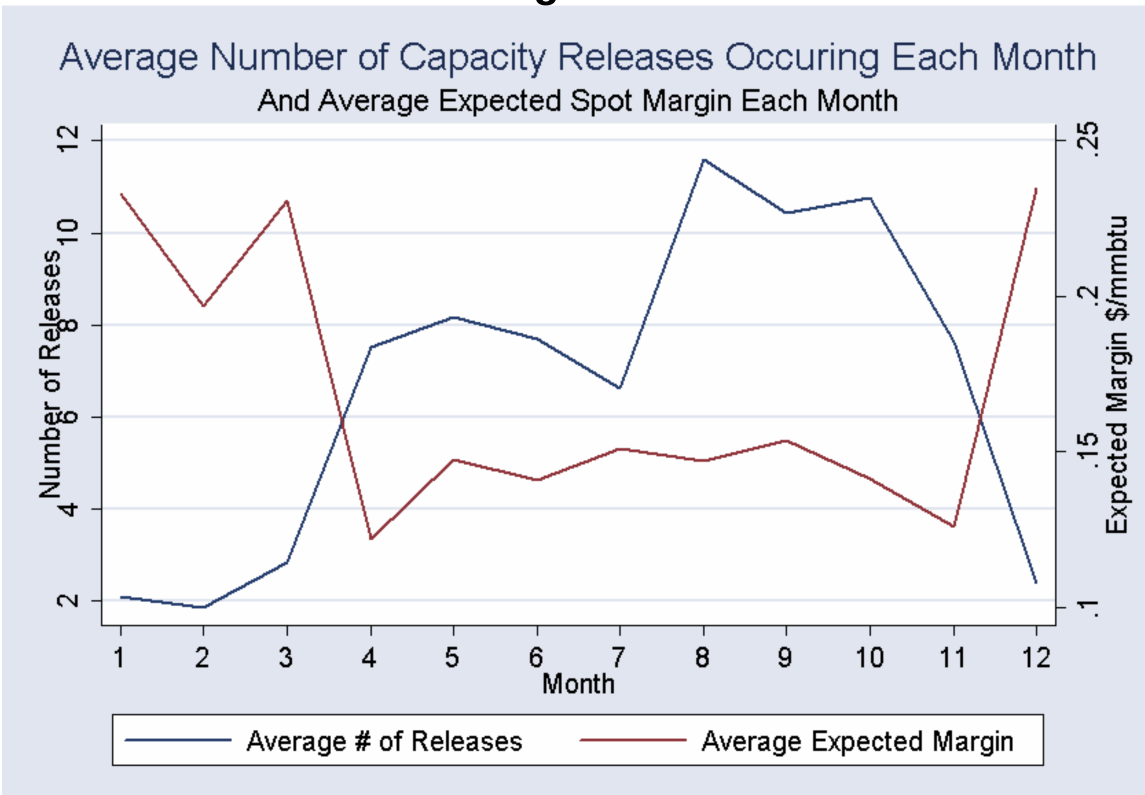

Table 8: Determinants of Number of Releases Per Month: Poisson Regressions

	Full Sample			Within Waiver Period			Outside Waiver Period		
	Incidence Rate Ratio	95% Confidence Interval		Incidence Rate Ratio	95% Confidence Interval		Incidence Rate Ratio	95% Confidence Interval	
RR Prediction	0.3256	0.0617	1.7172	0.4216	0.2197	0.8089	0.0818	0.0063	1.0562
January		omitted			omitted			omitted	
February	0.9137	0.4847	1.7224	1.1293	0.7315	1.7433	0.7727	0.2951	2.0232
March	1.5060	0.7290	3.1109	1.3533	0.8967	2.0423	1.6528	0.5540	4.9312
April	3.5727	2.1141	6.0375	2.4559	1.5918	3.7891	4.2039	2.2053	8.0139
May	3.8580	2.5174	5.9123	4.3149	2.8453	6.5437	3.7947	2.0315	7.0882
June	3.0884	1.8641	5.1168	3.2539	2.1787	4.8598	2.7424	1.4255	5.2759
July	3.2074	1.4643	7.0254	0.4605	0.3031	0.6998	4.9305	2.3684	10.2644
August	3.8203	2.1106	6.9148	3.2683	1.6972	6.2938	4.2918	1.9553	9.4203
September	3.4320	2.0262	5.8131	4.4498	2.9079	6.8094	2.1670	1.1870	3.9561
October	3.9174	2.3395	6.5597	5.1398	2.5583	10.3263	3.5603	1.8136	6.9890
November	2.8359	1.5723	5.1151	4.3912	2.6363	7.3142	2.1708	1.0092	4.6695
December	0.9942	0.3913	2.5264	2.0244	1.3434	3.0505	0.6944	0.2360	2.0431
		N: 409			N: 178			N: 231	

Regressions include route fixed effects. Standard errors are clustered on year-month.

Appendix: Principal-Agent Model of Procurement

Here, we offer a model and numerical example in which a principal would want to fire an agent that held an insufficient level of reserves relative to ex post needs, but only if that agent had sold reserves and come up short. The primary features of the model correspond to those discussed in the paper: the principal's cost of being short of reserves is higher than the cost of holding excess reserves, and the agents are heterogeneous—some are more skilled at forecasting demand than others. The intuition behind the model is tied to the concept of "plausible deniability," in that the principal will not fire an agent who can plausibly claim that he held insufficient reserves only because there was no additional supply available for purchase in the market.

Set-up

The timing of the model is as follows:

Period 0: The principal and the agent both observe the initial level of reserves, R_0.

Period 1: The agent receives a signal, S, about the probability distribution of the ultimate level of need, N, and decides on his target level of reserves, R^*.

Period 2: If $R_0 > R^*$, the agent sells reserves until he is left with R^*. If $R_0 < R^*$, the agent attempts to buy reserves, which may or may not be successful. The final level of reserves is R_2. After this, the principal observes $R_2 - R_0$, i.e., whether the agent bought or sold reserves.

Period 3: The level of need, N, is realized. The principal observes whether R_2 is greater than or less than N, and rewards or punishes the agent accordingly.

The level of need, N, is distributed uniformly on one of two intervals, either $[\mu_1 - r, \mu_1 + r]$ or $[\mu_2 - r, \mu_2 + r]$. The probability that N is distributed on either interval is equal to 1/2. All agents are aware of these ex-ante distributions.

Each agent may be either of the good type or the bad type. An agent is good with probability q and bad with probability $1 - q$. The good agent receives a perfect signal of which of the two possible distributions for N is correct; that is, he knows whether N is distributed on $[\mu_1 - r, \mu_1 + r]$ or $[\mu_2 - r, \mu_2 + r]$. The bad agent receives only a noisy signal of which distribution is the true one—specifically, the signal indicates which of the two distributions is correct with probability $s \geq 0.5$. If $s = 0.5$, the signal is useless, and if $s = 1.0$, a bad agent is equivalent to a good agent.

The agents know their types and are rational. In particular, the bad agent knows that his signal may be incorrect, and takes this into account in his decisions.

In period 2, if the agent wishes to sell, he can do so. If he wishes to buy, he may or may not be able to do so due to a thin marketplace. We model this by supposing that Y, a random variable, is the maximum amount that the agent will be able to buy if he decides to buy. Y is distributed according to the cdf $\Pr(Y \leq y) = 1 - e^{-\tau y}$, $y \geq 0$ so that the probability the agent will be able to buy at least y is given by $e^{-\tau y}$. The probability that the agent will be able to buy the quantity sufficient to reach his target level of reserves, R^*, is therefore $e^{-\tau(R^* - R_0)}$.

The principal's payoff function is given by:

$$
\begin{aligned}
U_P &= -F(R_2 - N) \;\text{ if } R_2 > N && (1)\\
&= -G(N - R_2) \;\text{ if } R_2 < N && (2)
\end{aligned}
$$

where $F(\cdot)$ and $G(\cdot)$ are both positive and monotonically increasing, $F(0) = G(0) = 0$, and $G(x) > F(x)\,\forall x$. This means that the principal's payoff function is maximized (at 0) when $R_2 = N$, and that shortfalls and excesses have asymmetric costs: a shortfall in reserves of any amount is more costly to the principal than excess reserves of the same amount.

While we model this problem as a single three period game, we recognize that a more comprehensive approach would consider a repeated game. In such a framework, the agents' valuations of future rounds would be modeled endogenously, and we would model the fact that, when the principal and agent interact repeatedly, the principal may gradually learn about the agent's type even if the agent does not sell reserves and come up short. Here, we abstract from the repeated game by compressing the model into one period, yet retain the features that cause an agent who sells reserves and stocks out to be fired.

The mechanism

Let the principal's transfers to the agent be defined as follows, in which W is a fixed wage (or more precisely, a wage in excess of what the agent could earn through outside employment), and α is a constant between zero and one:

$$
A_P = W - \alpha F(R_2 - N) \;\text{ if } R_2 > N \qquad (3)
$$

$$= \quad W - \alpha G(N - R_2) \text{ if } R_2 < N \text{ and } R_2 \geq R_0 \tag{4}$$

$$= \quad -\alpha G(N - R_2) \text{ if } R_2 < N \text{ and } R_2 < R_0 \tag{5}$$

Under this mechanism, the agent is fired (and is no longer paid the fixed wage) if he is short of reserves after having sold reserves, but is not fired if he is short following a purchase of reserves or no change in reserve levels. In this sense, W can be thought of as the value of all future wages, to be received by the agent only if he is not fired.

Whether this feature is attractive to the principal will hinge upon whether useful information is conveyed when an agent sells reserves and then stocks out. If such an observation implies that it is highly probable that the observed agent is of the bad type, then the expected value of firing the agent and drawing a new one from the pool will be positive. This expected value must be sufficiently large that it outweighs the expected cost to the principal of distorting the agent's behavior when selling reserves—the threat of being fired will cause the agent to target a higher reserve level when selling reserves than when buying reserves. Further, we also allow for a fixed cost of firing, C.

The example below shows that, with appropriate choices for parameter values, the good agent will never sell reserves below the maximum possible demand, which is equal to either μ_1+r or μ_2+r, depending on the known distribution of N. However, the bad agent, who does not know which distribution is correct, will find it value-maximizing to hold fewer reserves than the maximum possible demand of $\mu_2 + r$. He thereby accepts a small risk of being fired so that he may avoid large payments for holding excess reserves. The bad agent will therefore, on occasion, sell reserves below the maximum possible demand, and may be short after having sold. This difference in behavior allows the principal to perfectly separate the two types. We show that, given this separation, firing a bad agent and replacing him with a new agent from the pool can increase the principal's expected payoff. We also show that the principal will not want to fire the agent if he is observed with an insufficient level of reserves following a purchase, as such behavior does not send a clear signal of type.

Setup of numerical example

We will use a specific example with the following values:

$\mu_1 = 9$, $\mu_2 = 10$, and $r = 2$ (these together imply that $N \sim U[7, 11]$ half the time and $N \sim U[8, 12]$ the other half)

$q = \frac{1}{2}$ (the probabilities of drawing either a good or bad agent from the pool are equal)

$s = \frac{1}{2}$ (the bad agent's signal is useless)

R_0 has a discrete distribution and is equal to 8 with probability $\frac{1}{2}$ and equal to 14 with probability $\frac{1}{2}$.

$\tau = 0.05$

$F(R_2 - N) = 2(R_2 - N)$ (linear cost of holding excess reserves)

$G(N - R_2) = 48(N - R_2)$ (linear cost of being short)

$\alpha = \frac{1}{2}$ (half of the principal's payoff is passed through to the agent)

$W = 4$

The optimal choice of R^*

We begin by determining the optimal choice of R^*, from the principal's point of view, under the information sets of both the good agent and the bad agent. That is, we calculate what the principal would direct the agent to procure, had he the ability to do so. Suppose first that it is known that $N \sim U[8, 12]$. The principal's choice of R^* is the solution to the following maximization problem:

$$\max_{R_2} \int_8^{12} \frac{1}{4}[-2(R_2 - N) \cdot I_{R_2 > N} - 48(N - R_2) \cdot I_{R_2 < N}]dN$$

Here, the $I_{R_2>N}$ and $I_{R_2<N}$ are indicator functions for $R_2 > N$ and $R_2 < N$, and the 1/4 is present because this is the value of the uniform pdf of N on its support interval [8, 12]. This maximization is equivalent to:

$$\max_{R_2} \int_8^{R_2} -2(R_2 - N)dN - 48 \int_{R_2}^{12} (N - R_2)dN$$

Taking the first order condition via Leibniz's rule yields:

$$\int_8^{R^*} -dN + 24 \int_{R^*}^{12} dN = 0$$

The solution to this is $R^* = 11.84$. That is, if the principal knew that $N \sim U[8, 12]$ and chose reserve levels himself, he would choose to hold a reserve level equal to 11.84. While the set-up of this model does not permit the principal to choose reserves directly, note that a variant of the mechanism above, without the fixed wage that can be lost upon being fired, provides the agent with incentives that perfectly reflect the principal's payoffs. That is, the good type of agent, when incentivized with such a mechanism, will also choose to hold a reserve level equal to 11.84 when $N \sim U[8,12]$. In a similar fashion, the good type of agent will hold $R^* = 10.84$ when $N \sim U[7,11]$.

Now consider the choice of R^* that would maximize the principal's payoffs if the information available were that of the bad agent. In this case, the principal must choose R^* to maximize his payoff given that N may be distributed on either of two intervals. This maximization problem is given by:

$$\max_{R_2} \left\{ \begin{array}{c} \int_7^{11} \frac{1}{8}[-2(R_2 - N) \cdot I_{R_2>N} - 48(N - R_2) \cdot I_{R_2<N}]dN \\ +\int_8^{12} \frac{1}{8}[-2(R_2 - N) \cdot I_{R_2>N} - 48(N - R_2) \cdot I_{R_2<N}]dN \end{array} \right\}$$

Given an initial guess that $R^* > 11$, this is equivalent to the following, in which $\int_7^{11} -48(N - R_2) \cdot I_{R_2<N} dN$ equals zero:

$$\max_{R_2} \int_7^{11} -2(R_2 - N)dN - \int_8^{R_2} 2(R_2 - N)dN - 48 \int_{R_2}^{12} (N - R_2)dN$$

Taking the FOC via Leibniz's rule as before yields that $R^* = 11.68$. That is, if the principal had only the bad agent's information set, he would like to direct the agent to hold a reserve level equal to 11.68. If the principal were to offer a bad agent a payoff mechanism similar to that above, but without the threat of being fired, the bad agent would also choose this level of reserves.

Solution to the example when the agents are sellers

We now turn our attention to the target levels of reserves the agents will choose given the mechanism above (including the threat of being fired). We first focus on the case in which $R_0 = 14$. In this case, the initial level of reserves is higher than the maximum possible level of need under either distribution of N. Therefore, both types of agents will be sellers of reserves. We aim to show two results: first, that the good agent will target a reserve level equal to the maximum possible demand, given his signal, and, second, that the bad agent will target a reserve level below the maximum possible demand given his noisy signal, namely 12.

The good agent, who knows the true distribution of N, will choose a level of reserves that will balance the risk of being short (and therefore fired) against the cost of holding excess reserves. Suppose that the agent receives a signal that $N \sim U[8,12]$. The maximization problem that he solves is the following:

$$\max_{R_2} \int_8^{12} \frac{1}{4}[-(R_2 - N) \cdot I_{R_2>N} - [24(N - R_2) + 4] \cdot I_{R_2<N}]dN$$

Solving this using Leibniz's rule yields that $R^* = 12$—the maximum level of demand. The cost of being fired is sufficiently high that the risk incurred by holding fewer reserves than the maximum level perfectly balances the cost of holding excess inventory (for a fixed wage lower than $W = 4$, the good agent will hold fewer than 12 units of reserves). Similarly, when the signal is $N \sim U[7, 11]$, the agent will sell reserves until reaching $R^* = 11$. These results imply that the principal will never observe the good agent sell reserves and then fail to meet realized demand.

The bad agent solves:

$$\max_{R_2} \left\{ \begin{array}{c} \int_7^{11} \frac{1}{4}[-(R_2 - N) \cdot I_{R_2>N} - [24(N - R_2) + 4] \cdot I_{R_2<N}]dN \\ + \int_8^{12} \frac{1}{4}[-(R_2 - N) \cdot I_{R_2>N} - [24(N - R_2) + 4] \cdot I_{R_2<N}]dN \end{array} \right\}$$

The solution is $R^* = 11.84$. This implies that, if the $N \sim U[8, 12]$ state of the world is realized, then it is possible for the realized value of N to be high enough (namely, greater than 11.84) that the bad agent is short in Period 3, and therefore fired. As an aside, note that this result holds even when the bad agent's signal is somewhat, though not fully informative—the bad agent may be short for $s \in (0.5, 1)$.

Thus, the good agent will never expose himself to being fired by selling reserves that he might need, but the bad agent will. Therefore, whenever the principal observes that the agent sells reserves but stocks out, he has a perfect signal that the agent is of the bad type.

Solution to the example when the agents are buyers

The second case to consider is that of $R_0 = 8$, which implies that the agents must buy reserves in order to meet demand. When the agents purchase reserves, the firing penalty does not apply, so they will target reserve levels equal to those that are optimal for the principal. That is, the good agents will target $R^* = 10.84$ or $R^* = 11.84$, depending on their signal, and the bad agents will target $R^* = 11.68$. However, there is a risk that the agents will not be able to achieve their target: their probability of success in purchasing any $y = R_2 - R_0$ is given by $e^{-0.05y}$. Both types of agent may therefore be short of reserves when $R_0 = 8$. When the principal observes the agent purchase reserves, then come up short, he must therefore use Bayes' rule to determine the likelihood that the agent of the bad type. The analysis below shows that this likelihood is not much greater than the underlying probability of drawing a bad type from the pool, equal to $q = 0.5$.

We first determine the probability that a good agent will purchase reserves, and still be short. Taking the case of $N \sim U[8, 12]$, the probability that the agent will successfully procure $R^* = 11.84$ is given by $e^{-0.05(11.84-8)} = 0.8253$. The probability density function of R_2 is therefore given by:

$$f(R_2) = \begin{cases} 0.05\exp(-0.05(R_2 - 8)) & \text{if } R_2 \in [8, 11.84) \\ 0.8253 & \text{if } R_2 = 11.84 \end{cases}$$

Given that $N \sim U[8, 12]$, the probability that the agent will be short conditional on a given R_2 is the probability that the realized value of N will be less than R_2, given by $3 - \frac{R_2}{4}$. The unconditional probability that the good agent will be short is then determined by the following integral, which integrates the conditional probability over the pdf of R_2:

$$\Pr(R_2 < N|goodagent) = \int_8^{11.84} (3 - \frac{R_2}{4}) \cdot 0.05e^{-0.05(R_2-8)} dR_2 + 0.8253(3 - \frac{11.84}{4})$$

This evaluates to a probability of 0.1265. When $N \sim U[7, 11]$, the good agent targets $R^* = 10.84$, and the probability of a stockout is 0.0881. Thus, a good agent will be short of reserves following a purchase with an average probability of 0.1073.

The bad agent always targets $R^* = 11.68$ when buying reserves. The agent will be able to procure reserves sufficient to reach this level with probability 0.8319, and the overall probability of being short is given by:

$$\frac{1}{2}\int_8^{11.68} (3 - \frac{R_2}{4}) \cdot 0.05e^{-0.05(R_2-8)} dR_2 + 0.8319(3 - \frac{11.68}{4})$$
$$+\frac{1}{2}\int_8^{11} (2.75 - \frac{R_2}{4}) \cdot 0.05e^{-0.05(R_2-8)} dR_2$$

The first integral represents the probability of stockout if $N \sim U[8, 12]$, and the second represents the stockout probability if $N \sim U[7, 11]$. This evaluates to a probability of 0.1399.

With these probabilities in hand, we are now in a position to use Bayes' rule to find the probability that, having observed a stockout following a purchase, the agent is bad. We have:

$$\begin{aligned} \Pr(badagent|R_2 < N) &= \frac{\Pr(badagent) \cdot \Pr(R_2 < N|badagent)}{\Pr(R_2 < N)} \\ &= \frac{1/2 \cdot 0.1399}{1/2 \cdot (0.1073 + 0.1399)} \\ &= 0.5659 \end{aligned}$$

Observing a purchase followed by a stockout causes the principal to believe that the agent is bad with a 56.59% probability—a small increase over his prior of 50%.

Payoffs from the mechanism, and the incentive to fire

With the agents' actions now determined under all possible scenarios, we can calculate the expected payoffs to the principal. As an example of such a calculation, the expected payoff to the principal from having a good agent with a low signal who is a seller of reserves (so that $R_2 = R^* = 10.84$) is given by:

$$E[Payoff] = \frac{1}{4}[\int_7^{10.84} -2(10.84 - N)dN - 48\int_{10.84}^{11}(N - 10.84)dN] = -3.84$$

The case of a good agent with a low signal who is a buyer of reserves requires a double integral, to account for the probability distribution of R_2:

$$\begin{aligned} E[Payoff] &= \int_8^{10.84} \frac{1}{4}[\int_7^{R_2} -2(R_2 - N)dN - 48\int_{R_2}^{11}(N - R_2)dN] \\ &\quad *0.05\exp(-0.05(R_2 - 8))dR_2 - 0.8676 * 3.84 \\ &= -6.14 \end{aligned}$$

Proceeding with integrals such as these for all cases yields the following table of results for the principal's payoffs:

	high signal good agent		low signal good agent		bad agent		good - bad
	R^*	Payoff	R^*	Payoff	R^*	Payoff	
$R_0 = 14$, no firing	11.84	-3.84	10.84	-3.84	11.68	-4.68	0.84
$R_0 = 14$, firing	12.00	-4.00	11.00	-4.00	11.84	-4.76	0.76
$R_0 = 8$ (buying)	11.84	-9.46	10.84	-6.14	11.68	-8.52	0.72
Fire - don't fire		-0.08		-0.08		-0.02	

The first row of the table shows the target R^* for each type of agent, as well as the expected loss to the principal, when the agents are sellers of reserves and the mechanism of transfers from the principal to the agent does not include the threat of being fired. The second row also shows values for the case when the agents are sellers, but indicates behavior when the agents may be fired for selling and then stocking out. The third row illustrates the agents' behavior and the principal's losses when the agents are buyers of reserves (this behavior is invariant with regards to whether the mechanism includes firing or not, since the mechanism calls for firing the agent only if the agent has a stockout after having sold reserves).

The final row provides the expected change in the one-shot payoff to the principal when firing is included in the mechanism. Because the threat of being fired distorts the agents' incentives when they are sellers of reserves, this mechanism causes a small decrease in payoff. However, the expected payoff of firing a bad agent and replacing him with a new agent from the agent pool must also be considered. The last column of the table indicates the expected difference in payoff between the good agent and the bad agent—the good agent yields substantially greater expected profits to the principal: when the firing mechanism is in place, the additional expected payoff of a good agent is equal to 0.74 (this is the average of 0.72 and 0.76). Given the 50/50 chance of drawing a good agent from the pool, firing the bad agent yields an expected benefit of 0.37, greater than the small distortionary cost of the firing mechanism.

Also note that, given a small fixed cost of firing, the principal will not choose to fire the agent if the agent is observed with a stockout following a purchase of reserves. Recall that such an observation implies that the agent is bad with a probability of only 0.5659. Thus, if the principal fires the agent, there is only a $0.5 \cdot 0.5659 = 0.2830$ probability of actually replacing a bad agent with a good agent, and a $0.5 \cdot 0.4341 = 0.2170$ probability of replacing a good agent with a bad agent! The expected payoff of carrying out the firing is only 0.05. Therefore, as long as the fixed cost of firing, C, is greater than 0.05 but less than 0.37, it is beneficial for the principal to fire the agent when he observes a stockout following a sale, but not when he observes a stockout following a purchase.