

**How Do Firms Exercise Unilateral Market Power?
Evidence from a Bid-Based Wholesale Electricity Market**
by

Shaun D. McRae
Department of Economics
Stanford University
Stanford, CA 94305-6072
sdmcræ@stanford.edu

Frank A. Wolak
Department of Economics
Stanford University
Stanford, CA 94305-6072
wolak@zia.stanford.edu

March 18, 2009
Preliminary Draft

Abstract

This paper uses the framework in Wolak (2003a,b and 2007) and data on half-hourly offer curves and market-clearing prices and quantities from the New Zealand wholesale electricity market over the period January 1, 2001 to June 30, 2007 to characterize how the four large suppliers in this imperfectly competitive industry exercise market power. To accomplish this we introduce half-hourly measures of the firm-level ability and incentive of an individual supplier to exercise unilateral market power that are derived from a simplified model of expected profit-maximizing offer behavior in a multi-unit auction market. We then show that half-hourly market-clearing prices are highly correlated with the half-hourly values of the firm-level and firm-average measures of both the ability and incentive of the four large suppliers in New Zealand to exercise market power. We then present evidence consistent with the view that this increasing relationship between the ability or incentive of individual suppliers to exercise market power and higher market-clearing prices is caused by the four large suppliers submitting higher offer prices when they have a greater ability or incentive to exercise unilateral market power. We show that after controlling for changes in input fossil fuel prices and other factors that impact the opportunity cost of producing electricity during that half hour, each of the four suppliers submits a higher offer price into the wholesale market when it has a greater ability or incentive to exercise unilateral market power. To strengthen the case that this increasing relationship between market prices and the ability and incentive of each of the suppliers to exercise unilateral market power is actually caused by the four large suppliers exercising unilateral market power by changing their offer prices in response to their ability and incentive to exercise market power, we also perform a test of the implications of the null hypothesis that the four large suppliers behave as if they had no ability to exercise market power. We find strong evidence against this null hypothesis and instead find that these hypothesis testing results are consistent with the perspective that these suppliers are exercising all available unilateral market power.

1. Introduction

Empirical examination of the implications of profit-maximizing firm behavior in imperfectly competitive markets is complicated by the fact that the primitives of the economic environment, such as market demand functions and firm-level cost functions, are not directly observable. Moreover, the researcher rarely knows the strategic variables that firms use to influence market prices or often even the details of how market prices are set. As a result, researchers rely on parametric models of market demand and firm-level cost functions and equilibrium models of strategic interaction such as non-cooperative quantity-setting or price-setting behavior to understand how firms behave in imperfectly competitive markets. Consequently, any conclusions about firm behavior or the extent of market power exercised are conditional on these functional form assumptions and the assumed model of strategic interaction between firms.

We pursue an alternative approach that relies on a data-rich environment where many of these economic primitives are observable and both the strategic variables that firms choose and the exact mechanism that translates these strategic variables into market-clearing prices are known. This economic environment allows us to examine many implications of expected profit-maximizing behavior in imperfectly competitive markets without relying on functional form assumptions for market demand or a specific model of strategic interaction among firms.

To understand the advantages of the approach we pursue, it is useful to review the traditional approach from the perspective of the rapidly expanding literature in what Bresnahan (1989) calls the new empirical industrial organization. This approach uses market-clearing prices and quantities and variables assumed to shift demand and production costs along with three economic and behavioral assumptions to recover estimates of the extent of market power exercised in an imperfectly competitive market.

The three main econometric and behavioral assumptions necessary for validity of the traditional approach are: (1) parametric functional forms for the market demand and firm-level or market-level variable cost functions, (2) a model of firm-level strategic interaction, such as monopoly, quantity-setting competition, or price-setting competition, and (3) profit-maximizing or expected profit-maximizing behavior. Using a cross-section of monopoly newspaper markets, Rosse (1970) was the first to demonstrate that the combination of these three assumptions can allow a researcher to recover the firm's marginal cost function from market-clearing prices and

quantities and demand and cost shifters. The results of this modeling effort can then be used to estimate the marginal cost of the highest cost unit of output produced by the firm. This marginal cost equals the market-clearing price if the firm were unable to exercise any market power. Consequently, the difference between the market price and this estimated marginal cost measures the extent of market power exercised.

Porter (1984) applied this basic approach to an oligopolistic industry—19th century railroads. He assumed that actual market outcomes are the result of non-cooperative quantity-setting behavior between market participants. Bresnahan (1981 and 1987) measures the extent of market power exercised in the United States automobile industry by specifying a parametric discrete choice model of automobile demand at the individual household level that he aggregates to obtain market demand curves for each product sold. He then assumes that industry outcomes are the result of non-cooperative (in Bresnahan (1987) cooperative) price-setting by the major automobile suppliers. By specifying functional forms for variable cost functions for each product, Bresnahan achieves an econometric model of aggregate demand and product-level cost functions that can be estimated from data on market-clearing prices and quantities to measure the extent of market power exercised in this industry.

Berry, Levinsohn and Pakes (1995) and Nevo (2001) extend this basic approach to measuring market power in a number of important directions, although both papers follow the same general approach as Bresnahan. A parametric model of demand at the consumer level is assumed and aggregated across individuals to obtain product-level market demand functions. The assumption of non-cooperative expected profit-maximizing price-setting behavior (Bertrand competition) and parametric assumptions for product-level cost functions are employed to estimate the parameters of the demand and cost functions that can be used to measure the extent of market power exercised. Goldberg (1995) follows a slightly different approach by estimating an econometric model of household-level demand. She then aggregates these household-level demand functions to obtain market demand functions. The assumption of non-cooperative price-setting expected profit-maximizing behavior by firms is then employed to estimate the parameters of the product level cost functions and quantify the extent of market power exercised.

All of these studies and many more recent ones employing these techniques rely on an assumed parametric model of demand and a model of competition among firms to derive an estimate of the extent of market power exercised from market-clearing price and quantity data.

As has been emphasized by a number of authors, most forcefully by Bulow and Pfleiderer (1983), the estimate of the extent of market power exercised depends on the functional form assumed for the market demand. The assumed model of competition can also exert a substantial influence on the estimate of the extent of market power exercised. The most stark example of this result is the difference in the estimated extent of market power exercised resulting from assuming price-setting Bertrand competition in a homogenous-product constant-marginal-cost oligopoly versus quantity-setting competition in this same economic environment. The former will produce marginal-cost pricing, whereas the latter will achieve prices with significant markups over the marginal cost of the highest cost unit produced that depends on the number of firms and elasticity of the market demand.

All of these studies quantifying the extent of market power exercised do not explicitly address the question of how firms exercise market power, specifically what factors determine the extent of market power that firms are able to exercise and the amount of market power they choose to exercise. Because the amount of market power exercised is identified from market-clearing prices and quantities (and demand and cost shifters) using the functional form assumed for demand and the assumed model of competition among firms, any conclusions about how firms exercise market power or what factors enhance their ability and incentive to exercise market power are conditional on these two assumptions. Taking the example of an oligopoly model with the market demand based on individual discrete choices among products differentiated by observed and unobserved characteristics, a conclusion that certain values of the product characteristics are associated with larger amounts of market power being exercised is conditional on the assumed functional form for the underlying demand curves and assumptions about the form of strategic interaction among firms. Moreover, as Bresnahan (1987) demonstrates in his comparison of product-level versus firm-level price-setting competition, even for the same functional forms for market demands and production costs, assuming a different model of competition can lead to substantially different conclusions about the extent of market power exercised.

The recent world-wide trend toward introducing bid-based wholesale electricity markets has created an increasing number of data-rich economic environments where it is possible to study how firms behave in imperfectly competitive markets using only the assumption of expected profit-maximizing behavior. Participants in these multi-unit auction markets submit

their willingness-to-sell or willingness-to-purchase curves to the market operator and these curves are used to compute market-clearing prices and the quantities bought and sold by each market participant. A willingness-to-sell or willingness-to-buy curve gives the amount of the good a market participant is willing to sell or buy for each possible market-clearing price. If the researcher is willing to assume that a supplier constructs its willingness-to-supply curve to maximize the expected profits that it earns given the offers of its competitors and the bids of demanders, then it is possible to infer a supplier's variable cost function from the bid and offer curves that it and its market participants submit without having to resort functional form assumptions for aggregate demand or an assumed model of competition among firms.

This result follows by the same logic as described above. For the case of a multi-unit auction market, the offers submitted by other suppliers besides the supplier under consideration and the bids of all demanders determine the realized residual demand curve faced by that supplier. For the case that the researcher only has data on market-clearing prices and quantities, the residual demand curve a supplier faces is determined by the functional form assumption for aggregate demand and an assumed model of competition among firms. Because a supplier does not know the offers of other suppliers or all demand bids at the time it submits its willingness-to-supply curve, this supplier must construct its offer curve to maximize the expected profits that it expects to earn given the distribution of residual demand curves that it faces. Wolak (2003a) demonstrates that the assumption that the supplier chooses the form of its offer curve to maximize its expected profits given the distribution of residual demand curves that it faces identifies that supplier's marginal cost function.

Wolak (2003a) applies this logic to a multi-unit auction market for wholesale electricity to estimate generation unit-level variable cost-functions without the first two assumptions described above. The information contained in the offer curves submitted by all market participants and the assumption of expected profit-maximizing offer behavior by the supplier under consideration are sufficient to estimate generation unit-level marginal cost functions for a supplier. Wolak (2007) extends this cost function estimation framework to the case of multivariate cost functions in order to quantify the extent to which marginal costs for a specific generation unit in a given half-hour of the day vary with the level of output during that half-hour and during other half-hours of the day. Wolak (2003b) shows that the information contained in

the offer curves and demand bids can also be used to compute a measure of the ability of a supplier to exercise unilateral market power.

This paper uses the framework in Wolak (2003a,b and 2007) and data on half-hourly offer curves and market-clearing prices and quantities from the New Zealand wholesale electricity market over the period January 1, 2001 to June 30, 2007 to characterize how the four large suppliers in this imperfectly competitive industry exercise market power. To accomplish this we introduce half-hourly measures of the firm-level ability and incentive of an individual supplier to exercise unilateral market power that are derived from a model of expected profit-maximizing offer behavior in a multi-unit auction market. We then show that half-hourly market-clearing prices are highly correlated with the half-hourly values of the firm-level and firm-average measures of both the ability and incentive of the four large suppliers in New Zealand to exercise market power.

We then present evidence consistent with the view that this increasing relationship between the ability or incentive of individual suppliers to exercise market power and higher market-clearing prices is caused by the four large suppliers submitting higher offer prices when they have a greater ability or incentive to exercise unilateral market power. We show that after controlling for changes in input fossil fuel prices and other factors that impact the opportunity cost of producing electricity during that half hour, each of the four suppliers submits a higher offer price into the wholesale market when it has a greater ability or incentive to exercise unilateral market power.

This analysis considers alternative half-hourly measures of both the incentive and ability of each of the four large suppliers to exercise unilateral market power and finds a similar increasing relationship between a greater ability or incentive to exercise unilateral market power and a higher offer price by that supplier into the wholesale market. For all measures of the ability and incentive of a supplier to exercise market power we find that our regression coefficient estimates imply economically significant changes in a supplier's half-hourly offer prices for changes in the half-hourly value of both the ability and incentive of that supplier to exercise market power that occur routinely during our sample period.

To strengthen the case that this increasing relationship between market prices and the ability and incentive of each of the suppliers to exercise unilateral market power is actually caused by the four large suppliers exercising unilateral market power by changing their offer

prices in response to their ability and incentive to exercise market power, we also perform a test of the implications of the null hypothesis that the four large suppliers behave as if they had no ability to exercise market power. We find strong evidence against this null hypothesis and instead find that these hypothesis testing results are consistent with the perspective that these suppliers are exercising all available unilateral market power.

These empirical results lead to the following conclusions about the behavior of the four large suppliers in the New Zealand wholesale electricity market. First, these suppliers are successful at increasing market-clearing prices by raising their offer prices into the wholesale market when they have a greater ability to exercise unilateral market power. Second, even though an individual supplier may have a substantial ability to exercise unilateral market power, it may not in fact exercise this market power because it has no incentive to do so. However, when a supplier with the ability to exercise unilateral market power also has the incentive to do so, we show that it can do this by either raising or lowering its offer price and therefore the market-clearing price, depending on its fixed-price forward market commitments relative to the amount of energy it sells into the short-term market.

The remainder of this paper proceeds as follows. The next section provides a brief overview of the New Zealand wholesale electricity market. Section 3 introduces a simplified model of expected profit-maximizing offer behavior in a bid-based wholesale electricity market that forms the basis for our measures of the ability of a supplier to exercise market power. We then extend this model to the case that the supplier has fixed-price forward contract obligations. This model forms the basis for our measure of the incentive of a supplier to exercise unilateral market power. Section 4 discusses how our measures of the ability and incentive of suppliers to exercise market power are constructed from half-hourly offer and demand data and relates these half-hourly measures to the half-hourly values of market-clearing prices from the New Zealand wholesale electricity market over our sample period. Section 5 present the results of our empirical analysis relating half-hourly offer prices to half-hourly values of that supplier's half-hourly ability or incentive to exercise unilateral market power, after controlling for factors that determine variable cost differences for that suppliers across half-hours of the sample period. Section 6 presents the results of our analysis using alternative measures of the firm-level ability and incentive of suppliers to exercise unilateral market power. Section 7 presents the results of our test of the null hypothesis that some of the large suppliers behave as if they had no ability to

exercise unilateral market power. Section 8 closes with a discussion of the lessons this analysis of how firms exercise unilateral market power has for other electricity markets and other oligopoly markets.

2. The New Zealand Wholesale Electricity Market

In October 1996, a wholesale electricity market was formed by the New Zealand electricity supply industry. This market was a contract between market participants—generation unit owners, retailers, and energy traders—that specified how generation units were dispatched and wholesale prices were determined.

Prior to the start of the wholesale market, the transmission and generation sectors were dominated by the state-owned Electricity Corporation of New Zealand (“ECNZ”), which owned and operated more than 95% of the total New Zealand electricity generating capacity. ECNZ was broken up in three stages. First, in July 1994, the national transmission grid was separated into a stand-alone State-Owned Enterprise (“SOE”) Transpower. In February 1996, before the start of the wholesale electricity market, Contact Energy was formed out of ECNZ generation assets that represented roughly 22% of total electricity production. Contact was a stand-alone SOE in competition with ECNZ until it was privatized in 1999. Finally, about the same time as the privatization of Contact, the remainder of ECNZ was split into three competing SOEs: Genesis, Meridian and Mighty River Power. All three firms, as well as Transpower, remain state-owned.

In response to a perceived lack of competition in both the wholesale and retail markets, the Government announced a series of reforms of the electricity supply industry in April 1998. In addition to the final split of ECNZ, these reforms included the forced separation of distribution and retailing businesses. At the time there were more than 40 distribution firms, each with a very high market share in retailing for customers on their networks. The separation of distribution and retail led to rapid vertical integration between the generation and retail sectors, as Contact Energy and the newly-formed SOE generators bought the retail businesses from the network owners. Two new privately-owned generation and retail firms were created out of the industry reorganization—TransAlta New Zealand and TrustPower—although the former firm disintegrated in 2001.

Since 2001 the industry market structure has been relatively stable. There are five major generation owners: Contact, TrustPower, and the three SOEs, Genesis, Meridian and Mighty

River Power. Each of these generation owners is vertically integrated with a retail business serving a mix of residential, commercial and industrial users. With the exception of TrustPower, all of these firms have more generation capacity than their average retail load obligation, although there are half-hours during our sample period when each of these retailers has retail load obligations that exceed their sales in the short-term market.

As of mid-2007, the installed capacity of the New Zealand market is approximately 9050 MW. Meridian owns 2642 MW in hydroelectric and wind capacity. Contact is the second largest generation owner with 2287 MW of gas-fired, geothermal and hydroelectric generation units. Genesis owns 1886 MW of capacity, including the combination gas and coal-fired Huntly power station and 500 MW of hydroelectric capacity. Mighty River Power owns 1295 MW of a combination of hydroelectric, geothermal and natural gas-fired capacity. Trustpower owns 452 MW of hydroelectric and wind capacity. The remaining 488 MW of generation capacity is owned by a number of small firms, none of them with more than 120 MW of capacity.

The New Zealand electricity system consists of two alternating current subsystems, for the North and South Islands, connected by a 610 km (including 40 km submarine) High Voltage Direct Current (“HVDC”) cable between the largest hydroelectric scheme in the South Island and the Wellington region in the North Island. The maximum total power transfer capability is currently 1040 MW from south to north, the direction of flows for the majority of half-hour periods of the year, and approximately 600 MW from north to south. Hydroelectric energy availability in the South Island is the major determinant of the direction and level of energy flows.

More than 99% of the energy produced in the South Island comes from hydroelectric sources. There is sufficient generation capacity in the South Island to serve its annual electricity requirements, as well as export a substantial amount of energy to the North Island using the HVDC cable. Approximately 24.4% of the North Island supply came from hydroelectric sources in 2007, with the remaining 75.6% split between natural gas-fired (44.6%), coal-fired (11.6%), geothermal (13.0%), wind (3.4%), wood (2.1%), and less than 1% from biogas facilities.

Annual electricity consumption for the entire country in the year ending December 2007 was approximately 38.5 Terawatt hours (TWh), with the commercial sector consuming 23.3% of this total, the industrial sector 43.7%, and the residential sector 33.0%. An important aspect of the New Zealand electricity industry is that much of the population resides in the northern part of

the North Island in the Auckland metropolitan area, whereas many of the major hydroelectric resources are in the southern part of the South Island. As a result, transmission and distribution accounts for a relatively large fraction of the cost of delivered electricity compared to the rest of the world.

3. The Unilateral Market Power Problem in Wholesale Electricity Markets

This section introduces the economic theory underlying the measures used in our empirical work of the ability and incentive of a supplier in a multi-unit auction-based wholesale electricity market to exercise market power. All of these measures depend on the half-hourly willingness-to-supply curves of all producers and the level of half-hourly demand. Before proceeding with this discussion, we first define unilateral market power and why exercising all available unilateral market power is equivalent to a privately-owned firm serving its fiduciary responsibility to its shareholders.

A market participant is said to possess market power if it can take unilateral actions to influence the market price and to profit from the resulting price change. The demand side of most electricity markets is composed of many small buyers and the supply side is typically composed of a small number of large sellers. It is also relatively straightforward for a large supplier to withhold output from the short-term market, whereas it is extremely difficult, if not impossible, for a large demander to do this unless it curtails the consumption of the retail customers that it serves. Consequently, the primary market power concern in wholesale electricity markets is from suppliers taking actions to influence market prices.

It is important to emphasize that a supplier exercising all available unilateral market power subject to obeying the market rules is equivalent to that supplier taking all legal actions to maximize the profits it earns from participating in the wholesale market. Moreover, a firm's management has a fiduciary responsibility to its shareholders to take all legal actions to maximize the profits it earns from participating in the wholesale market. Consequently, a firm is only serving its fiduciary responsibility to its shareholders when it exercises all available unilateral market power subject to obeying the wholesale market rules.

3.1. Measuring the Ability to Exercise Unilateral Market Power in Bid-Based Markets

A supplier to an auction-based wholesale electricity market submits a willingness-to-supply or offer curve which is composed of a series of offer steps for each pricing period. The

length of the step specifies an incremental quantity of energy to be supplied and the height of the step is the price at which the supplier is willing to sell that quantity. The New Zealand market has 48 half-hourly pricing periods each day and suppliers are allowed to submit different price and quantity steps for their offer curves in each half-hour of the day.

Figures 3.1 and 3.2 show the final offer curves submitted by Firm A and Firm B for a peak half-hour period in February 2006. For the lowest-priced offer step, Firm A is willing to supply 920 MW at \$0.03/MWh and if the market price increases to \$60/MWh, it is willing to supply an additional 430 MW, and so on. As the offer price increases, the supplier's cumulative willingness to sell electricity increases along with the offer price, from 920MW at \$0.03/MWh to 1,350MW at \$60/MWh (= 920MW at \$0.03/MWh + 430 MW at \$60/MWh). This increasing relationship between the offer price and the supplier's cumulative willingness to sell yields the upward sloping offer curves for each supplier shown in Figures 3.1 and 3.2. Let $S_k(p)$ denote the offer curve of supplier k . At each price, p , this function gives the total quantity of energy that supplier k is willing to sell.

The offer curves from each supplier can be used to construct the aggregate offer curve for any set of suppliers. This is done by calculating the cumulative quantity that the set of suppliers are willing to supply across the relevant range of prices. Let $S_{123}(p)$ equal the aggregate offer curve for firms 1, 2, and 3. In terms of the individual offer curves, $S_{123}(p) = S_1(p) + S_2(p) + S_3(p)$, which means that $S_{123}(p)$ at price p is equal to total amount of energy that firms 1, 2, and 3 are willing to supply at price p . Figure 3.3 shows the aggregate offer curve for Firm A and Firm B for the firm-level offer curves shown in Figures 3.1 and 3.2. At a price of \$200/MWh, for example, Firm A is willing to supply a total of 1,650 MW and Firm B is willing to supply 835 MW. Therefore, the aggregate offer of both firms at a price of \$200/MWh is 2,485 MW.

Given the offer curves of all generation units in New Zealand, the price each generation unit receives for its output and each buyer pays for its withdrawals is determined by minimizing the as-offered cost of serving actual demand at all locations in the country. The as-offered cost for each generation unit is equal to the offer price times the offer quantity for each quantity increment or partial quantity increment accepted to provide energy summed over all offer price levels for that generation unit. The total as-offered cost of serving load in New Zealand is the as-offered cost for each generation unit summed over all units in the country. The total of all offer quantities accepted by M-co to produce energy is equal to the total demand at all locations in the

transmission network plus total transmission losses. The market price at each location in the transmission network is equal to the increase in the minimized value of the total as-offered cost of serving system demand at all locations in New Zealand as a result of an additional 1 MWh of load at that location.

This methodology for determining prices at each location in the transmission network is called locational marginal pricing and is discussed in detail in Bohn, Caramanis, and Schweppe (1984). These locational marginal prices (LMPs) or nodal prices differ across locations in the transmission network because of transmission losses and transmission congestion. For locations far from generation units, more energy must be injected by distant generation units in order to withdraw an additional 1 MWh from this location because of greater line losses in transferring the electricity from the point of injection to the point of withdrawal. In contrast, for locations close to generation units, the nodal price is lower because the electricity withdrawn at that location does not travel as far. Congestion in the transmission network arises when the amount of electricity that suppliers on one side of a transmission link would like to inject leads to flows on the transmission link that exceed its capacity. In these circumstances, prices on one side of the link must be reduced to lower the flows on the transmission line to its capacity and prices on the other side of the link must be increased to ensure that there is sufficient local generation to serve demand given the actual flows of the transmission link into the area.

Transmission congestion in the New Zealand wholesale market, as measured by the number of half-hours with price differences between locations that are greater than can be explained by line losses, is very infrequent. Line losses also tend to produce persistent price differences across locations in New Zealand because generation-rich nodes (those with low loss factors) and generation-poor nodes (those with high loss factors) within each island tend to remain so regardless of the level of demand throughout New Zealand.

Consequently, during the periods when no transmission constraints are binding, the price at each location in New Zealand is well-approximated by taking the aggregate willingness-to-supply curve across all locations in New Zealand and solving for the price where this curve intersects the total demand in New Zealand. Define $S(p)$ as the aggregate willingness-to-supply curve for a half-hour. It is equal to $S_1(p) + S_2(p) + \dots + S_K(p)$, where K is the total number of suppliers in New Zealand. Let $QD = QD_1 + QD_2 + \dots + QD_M$, where QD_m is the actual real-time demand at node m and M is the total number of nodes in New Zealand.

This no-congestion market-clearing price is the solution in p to the equation $S(p) = QD$. An example of this process is shown in Figure 3.4 for the same half-hour period as in Figures 3.1 to 3.3. In the period, the total market demand is 4,400 MW and based on the aggregated offer curve for all the suppliers, the market price has to be at least \$120/MWh for there to be enough supply offers to meet this demand.

This description of the price-setting process in the New Zealand market allows a graphical description of how suppliers exercise unilateral market power in a bid-based wholesale market, which motivates our measure of the ability of a supplier to exercise unilateral market power. As we discuss below, the basic intuition and insights provided by this single-price, graphical analysis carry over to the case of the nodal price-setting process used in the New Zealand market. To analyze the bidding behavior of an individual supplier using this graphical framework, the above mechanism can be reformulated in terms of the supplier's own offer curve, the offers of other suppliers and the total market demand. Specifically, the price setting equation $S(p) = QD$ can be re-written as:

$$S_1(p) + S_2(p) + \dots + S_k(p) = QD.$$

Suppose that we are interested in measuring the ability of supplier j to exercise unilateral market power. This price-setting equation can be re-written as:

$$S_j(p) = QD - (S_1(p) + S_2(p) + \dots + S_{j-1}(p) + S_{j+1}(p) + \dots + S_k(p)) = QD - SO_j(p),$$

where $SO_j(p)$ is the aggregate willingness-to-supply curve of all firms besides supplier j . Define $DR_j(p) = QD - SO_j(p)$ as the residual demand curve facing supplier j . The residual demand of supplier j at price p is defined as the market demand remaining to be served by supplier j after the willingness to supply curves of all other firms besides supplier j have been subtracted out.

Figure 3.5 provides a graphical version of the above calculation of the residual demand for Firm A in the same half-hour period. The total market demand is 4,400MW and the total quantity offered by all suppliers other than Firm A is 3,350MW at \$300 and 2,560MW at \$50. Therefore, Firm A's residual demand at \$300 is 1,050MW (the market demand of 4,400MW minus 3,350MW of supply by other generators at that price). Its residual demand at \$50 is 1,840MW (the market demand of 4,400MW minus 2,560MW of supply by other generators at that price). Figure 3.6 shows the residual demand curve resulting from performing this calculation for all possible prices for Firm A in this half-hour period.

Figure 3.7 combines Firm A's residual demand curve from Figure 3.6 with Firm A's offer curve from Figure 3.1. Because this is a half-hour period with no transmission congestion, nodal prices differ across locations in New Zealand only because of line losses, which implies small differences between nodal prices across most locations in the New Zealand. If p_m is price at node m and q_m is the amount of energy injected at node m , the quantity-weighted average nodal price is the sum of the product of the nodal price and nodal quantity of energy injected over all M locations in New Zealand divided by the sum of the nodal injections at all locations in New Zealand during that half-hour. Mathematically, this quantity weighted average price $p(avg)$, is equal to

$$p(avg) = \frac{\sum_{m=1}^M p_m q_m}{\sum_{m=1}^M q_m}.$$

Applying this process to all of the nodal prices for this half-hour period yields a quantity-weighted average price of \$120/MWh.

The residual demand curve that a supplier faces summarizes its ability to impact the market price through changes in its offer curve, holding the offer curves of other suppliers constant. This residual demand curve also gives the quantity of energy that the offer curves submitted by the firm's competitors allow it to sell at each possible price. A firm can choose to produce any price and generation quantity pair along its residual demand curve. For example, Figure 3.8 shows the residual demand curve for Firm A calculated above. The realized price was \$120/MWh and the quantity supplied by Firm A was 1,500MW, which gives Firm A generation revenues of \$90,000 in the half-hour. However, if Firm A had reduced the amount of energy it supplied by 15 percent to 1,270 MW, this would have increased the market price to \$250/MWh. This price and quantity combination yields generation revenue of \$158,750, even though Firm A supplies less energy to the wholesale market at this substantially higher price.

As shown in Figure 3.8, Firm A could have increased the market price by 108% with a reduction in its quantity supplied of 15%. We define the ratio of the potential percentage increase in market price to the percentage reduction in quantity supplied as the inverse elasticity of the residual demand curve. In this case the inverse elasticity is $108/15 = 7.2$. Higher values of the inverse elasticity mean that the supplier has greater ability to unilaterally change the market price.

As noted above, a supplier's residual demand curve gives the set of feasible price/quantity pairs that it can choose from to maximize its profits. Firms in imperfectly

competitive markets often speak of “pricing to take what competition gives them” or “pricing at what the market will bear.” These statements can be interpreted as the firm choosing the price/quantity pair along its residual demand curve that maximizes its profits. In this sense, a supplier’s residual demand curve shows the trade-off between a higher system price and lower generation quantity for the supplier because of supply responses of its competitors. The supplier maximizes profits by producing at the output level where the marginal revenue associated with selling an additional MWh equals the marginal cost associated with producing an additional unit. For the residual demand curve and marginal cost curve in Figure 3.9, a profit-maximizing firm will supply a quantity Q_1 , the output level at the point of intersection of the marginal cost and marginal revenue curves. Note that the system price will be P_1 , the intersection of quantity Q_1 with the residual demand curve. Note that this price exceeds the firm’s marginal cost of supplying Q_1 . The firm’s profits are given by the area left of Q_1 , below P_1 and above the marginal cost curve.

Figure 3.10 repeats the process of computing the profit-maximizing level of output for a flatter residual demand curve and the same marginal cost curve as in Figure 3.9. A profit-maximizing supplier will produce the quantity Q_2 at a price of P_2 . Note that the difference between P_2 and the marginal cost of production at Q_2 is smaller than this same magnitude in Figure 3.9, which is a result of the flatter or more elastic residual demand curve in Figure 3.10. The case of a perfectly elastic residual demand curve is shown in Figure 3.11. This residual demand curve is the result of a flat aggregate offer curve of all other suppliers besides supplier j , which implies that there were many other firms willing to supply the entire market at the price P_3 in Figure 3.11. For this residual demand curve, the marginal revenue curve coincides with the residual demand curve, because producing an additional unit of output has no effect on the market price. The firm will produce at the point of intersection of its marginal cost curve with its residual demand curve, which is the output level Q_3 and price P_3 in Figure 3.11.

This example demonstrates the very important point that if a supplier faces a sufficiently elastic residual demand curve, typically because there is a large number of independent suppliers competing to sell energy at close to the same price, then it is unilaterally profit-maximizing for this supplier to produce at the point where the market price is equal to its marginal cost. As we demonstrate below, the firm accomplishes this by submitting an offer curve that is equal to its

marginal cost curve, because the intersection of this offer curve with its residual demand curve produces the desired price/quantity pair.

The examples in Figures 3.9 and 3.10 demonstrate that when a profit-maximizing supplier faces an upward sloping residual demand curve, the firm will find it unilaterally profit-maximizing to produce at an output level that is below the output level at the point of intersection of its marginal cost curve with its residual demand curve. In Figure 3.9, the firm would optimally offer only Q_1 into the market, even though the price P_1 greatly exceeds its marginal cost at that level of production. The firm accomplishes this by submitting an offer curve that lies above its marginal cost curve—that is, an offer curve with an offer price for that level of output above the marginal cost of producing that output. Figures 3.9 and 3.10 demonstrate that the difference between the supplier's profit-maximizing offer price and the supplier's marginal cost will be greater when the inverse elasticity of the residual demand curve is larger. This is an important implication of expected profit-maximizing behavior that will be explored in subsequent sections of this paper.

3.2. A Simplified Model of Expected Profit-Maximizing Offer Behavior

All of the examples presented thus far have assumed the residual demand curve is known when the supplier computes its profit-maximizing output level. Because a supplier's residual demand curve is composed of the offer curves of its competitors and the market rules require all suppliers to submit their offers at the same time, this assumption is not in fact true. However, the economic justification for using the inverse elasticity of a supplier's residual demand curve as a measure of its ability to exercise unilateral market power carries over to the case that suppliers do not observe the actual residual demand curve they face at the time they submit their offers to the wholesale market.

Although a supplier does not know with certainty the market demand and the willingness-to-supply offers of other suppliers when it submits its offers for the pricing period, the supplier does have a very good idea of the set possible realizations of the residual demand curves it might face. The characteristics of each generation unit owned by the supplier's competitors and the market rules can significantly constrain the set of offers curves a supplier can submit. For example, the New Zealand markets rules specify a maximum number of quantity and price steps for each generation unit that owner can submit in their offer curve. The maximum value of the sum of these quantity steps must be less than the capacity of the

generation unit. Producers are able change their willingness-to-supply functions up to 2 hours before the trading period, so there is likely to be very little uncertainty in system demand at the time they submit their final willingness-to-supply offers. All of these factors imply that the supplier has a very good idea of the set of possible realizations of the residual demand curve it might face. For each possible residual demand curve realization the supplier can find the ex post profit-maximizing market price and output quantity pair given its marginal cost curve following the process described above. This is the market price and output quantity pair that the supplier would like to achieve for that residual demand curve realization.

Figure 3.12 illustrates the construction of an expected profit-maximizing willingness to supply curve using this process for the case of two possible residual demand curve realizations. For each residual demand curve realization, intersect the marginal cost curve with the marginal revenue curve associated with that residual demand curve realization. For example, for Residual Demand Curve 1 the marginal revenue curve for this residual demand curve (not shown on the figure) intersects the marginal cost curve at the quantity Q_1 . The output price associated with this output level on Residual Demand Curve 1 is P_1 . Repeating this process for Residual Demand Curve 2 yields the profit-maximizing price and quantity pair (P_2, Q_2) . Note that because both residual demand curves are very steeply sloped, there is a substantial difference between the market price and the marginal cost at each output level. If these two residual demand realizations were the only ones faced the supplier, it would submit an offer curve that passes through both of these points because regardless of the residual demand realization this offer curve would cross at an ex post expected profit-maximizing level of output. The straight line connecting the points (P_1, Q_1) and (P_2, Q_2) is one such expected profit-maximizing offer curve.

To illustrate the impact of more elastic residual demand curves on the offer curves submitted by an expected profit-maximizing supplier, Figure 3.13 repeats the construction of an expected profit-maximizing offer curve for the case of two more elastic residual demand curve realizations. The line connecting the points (P_1, Q_1) and (P_2, Q_2) , which is an expected profit-maximizing offer curve for these two residual demand realizations, is much closer to the supplier's marginal cost curve. Specifically, for each residual demand realization, the price associated with the profit-maximizing level of output for that residual demand curve realization is closer to the marginal cost of producing that level of output than it was in Figure 3.12. This

outcome occurs because each residual demand realization is much more elastic than the residual demand realizations in Figure 3.12.

Figure 3.14 considers the case in which the two residual demand curve realizations are infinitely elastic, meaning that for both realizations the supplier sufficient competition so that the entire market can be satisfied at a fixed price. By the logic described above, the supplier will find it unilaterally profit-maximizing to produce at the intersection of each residual demand curve realization with its marginal cost curve. In this case, the supplier expected profit-maximizing offer curve, the line connecting the profit-maximizing output levels for each residual demand curve realization, is equal to the supplier's marginal cost curve. This result illustrates a very important point that if a supplier faces sufficient competition for all possible residual demand curve realizations then it will find it unilaterally expected profit-maximizing to submit an offer curve equal to its marginal cost curve.

The examples in Figures 3.12 to 3.14 utilize linear residual demand curves. However, the same process can be followed to compute an expected profit-maximizing offer curve for the case of step-function residual demand curves. Figure 3.15 shows how this would be done for the more realistic case of step function residual demand curves with two possible residual demand realizations. For each, residual demand curve realization, the supplier would compute the profit-maximizing level of output and market price for the marginal cost curve given in Figure 3.15. For DR_1 this is the point (P_1, Q_1) and for DR_2 this is the point (P_2, Q_2) . If these two residual demand curve realizations were the only possible residual demands that the supplier could face, then a step function offer curve that passes through these two points would be an expected profit-maximizing offer curve.

Computing the expected profit-maximizing offer curve for a supplier is generally more complex than passing an offer curve through the set of all possible ex post expected profit-maximizing price and output quantity pairs. That is because the market rules can prevent a supplier from achieving the ex post profit-maximizing market price and output quantity pair for all possible residual demand realizations. Specifically, unless all of these ex post profit-maximizing price and quantity pair lie along a willingness-to-supply curve for the supplier that the market rules allow it to submit, it is not possible for the supplier to submit a willingness to supply curve that always crosses the realized residual demand curve at an ex post profit-maximizing price and quantity pair for that residual demand curve realization. Figure 3.16

provides an example of this phenomenon. This figure shows the ex post profit-maximizing price and quantity pairs for three residual demand curves. Note that the profit maximizing point for DR_3 lies below and to the right of the profit maximizing point for DR_1 . This makes it impossible for the supplier to submit a non-decreasing step function offer curve that passes through the three ex post profit-maximizing price and output quantity pairs. In this case, the supplier must know the probability of each residual demand curve realization in order to choose the parameters of its expected profit-maximizing willingness to supply curve.

Figure 3.16 demonstrates that the expected profit-maximizing residual demand curve does not pass through any of these three ex post profit-maximizing price/quantity pairs. Instead, as discussed in Wolak (2003a and 2007), the form of the expected profit-maximizing willingness-to-supply curve depends on both the form of each residual demand curve realization and the probability of that residual demand curve realization. This curve, shown in Figure 3.16, yields market-clearing price and quantity-sold pairs for the firm for each of the three residual demand curve realizations that maximize the expected profits the firm earns subject to this offer curve being in the set of offer curves the market rules allow a supplier to submit. As shown in Wolak (2003a) and Wolak (2007), the supplier chooses the price level and quantity increments that determine its offer curve to maximize its expected profit over the distribution of residual demand curve realizations that it faces.

Nevertheless, the inverse elasticity of the realized residual demand curve at the actual market-clearing price still provides an ex post measure of the ability of a supplier to exercise market power. Specifically, this inverse elasticity quantifies the percentage increase in the market-clearing price that would have occurred if the supplier had reduced the amount of output it sold in the market by a pre-specified percentage. This interpretation of the inverse elasticity of the residual demand curve does not rely on the assumption that the realized output level and market-clearing price maximize the supplier's ex post profits.

As emphasized in Wolak (2003a and 2007), expected profit-maximizing offer behavior does not imply that every point of intersection of the supplier's offer curve with its residual demand curve yields the ex post profit-maximizing price and output quantity pair for the supplier for that residual demand curve realization. Therefore, this result implies that in general, there is no deterministic relationship between the difference between the market-clearing price and the

firm's marginal cost of production at its actual output level divided by the market-clearing price and the value of the inverse elasticity of the residual demand curve the supplier faces.

3.3. Impact of Fixed-Price Forward Market Obligations on Short-Term Market Behavior

The above discussion of expected profit-maximizing offer behavior assumes that the supplier only earns revenues from selling energy in the wholesale market. However, as noted earlier the four large suppliers in the New Zealand market are all vertically integrated. They not only sell energy in the wholesale electricity market, but they also sell electricity to retail customers at retail prices that do not vary with hourly prices in the wholesale market. These fixed-price retail load obligations function very much like fixed-price forward contract obligations, because the vertically-integrated supplier has essentially made a commitment to provide its fixed-price retail load obligation at a pre-determined wholesale price.

For example, suppose that during a half-hour period in February 2006 Firm C had retail load obligations of 880 MW at various fixed prices. That means that Firm C would be obliged to supply 880 MW retail load at those prices regardless of the actual wholesale price. This implies that Firm C has a strong financial incentive to purchase the 880 MW at the lowest price possible. Figure 3.17 shows the offer curve and residual demand curve for Firm C in that period. By increasing the price at which it offered in its generation, Firm C could have moved to the point on its residual demand curve shown by the red dot. This would have increased the market price by 100% (from \$120 to \$240/MWh) and reduced the quantity supplied by Firm C by 35% (from 1,170 MW to 765 MW). As a result, Firm C's generation revenue in this half-hour would have increased from \$70,200 to \$91,800. However, Firm C's net revenue from transactions in the wholesale market would have decreased substantially. At the actual prices and quantities, Firm C sold 1,170 MW and bought 880 MW from the wholesale market, at the market price of \$120/MWh. Therefore, its net position was 290 MW, the difference between 1,170 MW and 880 MW, so Firm C's net revenue would have been \$17,400 (290 MW at \$120/MWh for one half-hour). At the higher price, Firm C would have sold 765MW while still buying 880 MW from the wholesale market, now at the higher market price of \$240/MWh. Firm C's net position would have been -115 MW, and its net revenue -\$13,800 (-115 MW at \$240/MWh for one half-hour). This example demonstrates the importance of the supplier's fixed-price load obligations in considering its incentive to increase the market price.

In general, because a supplier with fixed-price retail load obligations must serve this load at a fixed price no matter what the actual wholesale price is, a wholesale price increase has two opposing effects on the supplier's profits: (1) it increases the supplier's profits from selling energy in the wholesale market; and (2) it decreases the suppliers' profits by raising the cost of serving its retail demand. Consequently, whether and to what degree a price increase is beneficial to a vertically-integrated supplier depends on whether and to what degree the profit increase from withholding supply more than offsets the increase in cost of serving the supplier's retail load covered by a fixed-price forward market obligation. If the profit loss due to the cost increase in (2) exceeds that profit gain in (1), a supplier would lose profits from a market price increase. In that case, the supplier would have no incentive to exercise market power by withholding output from the wholesale market to increase the market-clearing price.

For a supplier, the comparison between its profit gain and loss from a price increase depends on the difference between the supplier's sales in the short-term market and its fixed-price load obligations. For example, suppose that the supplier's sales in the wholesale market are 2,000 MW while its fixed-price load obligation is 1,500 MW. In that case, a \$1 increase in market price would increase the supplier's profits from its generation sales by \$2,000 and increase the cost of its load obligation by \$1,500, implying a net gain of \$500 (or \$1 times the 500 MW difference between the supplier's supply of 2,000 MW and load obligations of 1,500 MW). In that case, the supplier has an incentive to increase the market price through its unilateral actions because it is profitable to do so. However, if the supplier has a significantly larger load obligation of 2,500 MW, then the \$1 increase in market price would imply a net loss of \$500 (or \$1 times the -500 MW difference between its supply and load obligations) as the supplier's profit gain from its generation sales (\$2,000) is less than the increase in its cost to serve the fixed-price retail load obligation (\$2,500).

To understand the incentives to exercise unilateral market power of a supplier with fixed-price retail load obligations or fixed-price long-term contract obligations define the following notation. Let P_R equal the retail price at which the firm is selling Q_R MWh of retail electricity. Let $DR(p)$ equal the firm's residual demand curve for sales in the short-term market and p the market price. For simplicity, assume that c is the constant marginal cost of producing electricity and τ is average cost of distributing wholesale electricity to retail customers. The vertically-integrated sellers in New Zealand also participate in the market for fixed-price long-term

contract obligations. Let P_C equal the quantity-weighted average price of fixed-price forward contract obligations held by the vertically integrated firm and Q_C equal the MWh quantity of fixed-price forward contract obligations. The firm's profits from selling in wholesale market during that pricing period given these forward market commitments is equal to

$$\Pi(p) = (P_R - p)Q_R + DR(p)(p - c) - (p - P_C)Q_C - \tau Q_R.$$

The first term is the profits from retail sales. The second term is the profits from wholesale electricity sales in the short-term market. The third term is the profits or losses from fixed-price forward contract obligations, and the final term is the cost of distributing retail electricity. This expression for the vertically-integrated firm's profits can be re-written as:

$$\Pi(p) = (P_R - \tau - c)Q_R + (P_C - c)Q_C + [DR(p) - (Q_R + Q_C)](p - c).$$

The first and second terms are profits from retailing assuming Q_R cost c \$/MWh to produce and the second term is the profit from sales of fixed price forward contracts assuming Q_C is produced at c \$/MWh. The third term is the only one that depends on the short-term market price. The first and second terms only depend on variables that the supplier cannot influence at the time they are offering into the short-term market.

This form of the firm's profit function shows that the values of Q_R and Q_C , the firm's retail load obligation and fixed-price forward contract obligations, influence its incentive to exercise unilateral market power. Even though the supplier may face a very inelastic residual demand curve, it would have little incentive to reduce the output it sells to raise prices above its marginal cost if the amount it sells in the short-term market, $DR(p)$, is less than the sum of its fixed price forward market obligations, $Q_R + Q_C$. Under these circumstances, the vertically-integrated supplier is a net buyer from the short-term market. It has obligations for purchases of $Q_R + Q_C$ from the short-term market and it only sells $DR(p)$. As a net buyer, the supplier would like the price to be as low as possible. When $DR(p)$ exceeds $Q_R + Q_C$ the vertically-integrated supplier is a net seller in the wholesale market and as such would like to raise the price at which it sells its net output in the short-term market.

The difference between a firm's sales in the short-term market and its fixed-price retail load and forward contract obligations is its residual demand net of its forward market obligations. In terms of the above notation, this net residual demand curve is equal to $DR_F(p) = DR(p) - (Q_R + Q_C)$. Depending on whether a supplier's net residual demand is positive ("net long") or negative ("net short"), the supplier has incentive to either increase or decrease the

market price through its unilateral action. If a supplier is net long (i.e., has a positive net residual demand), it will benefit from a higher market price because it is making net sales into the short-term market. Consequently, the larger a supplier's net residual demand, the greater is the supplier's gain from a market price increase. Conversely, the more a supplier is net short (i.e., a negative net residual demand), the greater the supplier's incentive to decrease the market price because it is a net buyer from the short-term market.

In terms of this net residual demand function, the firm's profit function becomes:

$$\Pi(p) = DR_F(p)(p - c) + F, \text{ where } F = (P_R - \tau - c)Q_R + (P_C - c)Q_C.$$

The first two terms in the profit function written above are collected into the term F because all of the variables comprising of these terms are not affected by the supplier's offers into the short-term wholesale market and are known before the supplier submits its offers. This expression for the vertically integrated supplier's profit function takes the same form as a non-vertically integrated supplier.

To determine the firm's profit-maximizing price and quantity pair we can solve for the value of p that maximizes the above expression. We can also follow a slightly more involved version of the graphical approach shown in Section 3.2. Figure 3.18 shows the net residual demand curve $DR_F(p)$, which is calculated by shifting the original residual demand curve $DR(p)$ to the left by the amount of the supplier's fixed-price forward market obligations, $Q_R + Q_C$. Using $DR_F(p)$ we can compute $MR_F(p)$, the supplier's marginal revenue curve for sales in excess of its fixed-price forward market obligations. Because the firm's production decision must still take account of its forward market position, Figure 3.19 shifts $MR_F(p)$ to the right by the amount of the fixed-price contract obligations. The firm produces at the point where $MR_F(p) + Q_R + Q_C$, the shifted marginal revenue curve, intersects the marginal cost curve MC . This is output level Q_4 in Figure 3.19. The short-term market price is determined by the supplier's original residual demand curve at this level of production, the price P_4 in Figure 3.19.

Figure 3.20 demonstrates the impact of fixed-price forward market obligations on the supplier's expected profit-maximizing price and output quantity pair. For the residual demand curve given in Figure 3.20, a supplier without any forward market obligations would find it optimal to produce at the price and output quantity pair (P_1, Q_1) that was derived in Figure 3.9. A supplier with the level of fixed price forward market obligations facing this same residual demand curve would find it unilaterally profit-maximizing to produce at the price and output

quantity pair (P_4, Q_4) . As shown in Figure 3.20, a firm with fixed-price forward market obligations facing the same residual demand curve finds it unilaterally profit-maximizing to sell more output in the short-term market at a lower price, $Q_4 > Q_1$ and $P_4 < P_1$.

There is even a level of fixed-price forward market obligations that would cause a supplier facing a steep residual demand curve to find it unilaterally profit-maximizing to produce at the point of intersection of its marginal cost curve with its residual demand curve. Specifically, if $Q_R + Q_C$ is chosen to equal $DR(c)$, the value of output at the point of intersection of the residual demand curve with the supplier's marginal cost curve, the supplier will find it unilateral profit-maximizing to produce at $DR(c)$, regardless of the slope or inverse elasticity of the residual demand curve. In other words, a supplier that possesses substantial ability to exercise unilateral market power as measured by the inverse elasticity of its residual demand curve, has no incentive to do so because of the level of its fixed-price forward market obligations.

The relationship in Figure 3.20 carries over to the case of constructing expected profit-maximizing offer curves with fixed-price forward market obligations. Figure 3.21 repeats the computation of the expected profit-maximizing offer curve for the same two residual demand curve realizations for the case of no fixed-price forward market obligations and positive fixed-price forward market obligations. For the case of a positive forward market obligation, the expected profit-maximizing offer curve is much closer to the firm's marginal cost curve than the expected profit-maximizing offer curve derived assuming the firm has no fixed-price forward market obligations. This is a general result of the impact of fixed-price forward market obligations on the expected profit-maximizing offer curve of a supplier. The higher the level of fixed-price forward contract obligations relative to the supplier's actual short-term market sales, the closer is the expected profit-maximizing offer curve to the supplier's marginal cost curve.

Because fixed-price forward market obligations alter the incentive of a supplier to exercise unilateral market power, the net residual demand curve can be used to construct a measure of the incentive, as distinct from the ability, of a supplier to exercise unilateral market power. This measure is the inverse elasticity of the net residual demand curve. In terms of $DR_F(p)$ this inverse elasticity is defined as:

$$1/\varepsilon^F = - [DR_F(p)/p] * [1/DR_F'(p)],$$

which is also equal to the percentage change in the market-clearing price as a result of a one percent change in the net residual demand of the supplier. The inverse elasticity of the net residual demand curve is related to the inverse elasticity of the residual demand curve by the following equation

$$1/\varepsilon^F = - \{ [DR(p) - [Q_R + Q_C]] / DR(p) \} * [1/\varepsilon] \text{ where } 1/\varepsilon = - [DR(p)/p] * [1/DR'(p)].$$

The inverse elasticity of the residual demand curve times the exposure of the supplier to the short-term market is equal to the inverse elasticity of the net residual demand curve. Note that in spite of the fact that the inverse elasticity of the residual demand curve is always positive the inverse elasticity of the net residual demand curve can be negative or zero. Zero occurs if the supplier's short-term market sales equals its fixed-price forward market obligations, $DR(p) = [Q_R + Q_C]$. A negative inverse elasticity occurs if the supplier's short-term market sales are less than its fixed-price forward market obligations, $DR(p) < [Q_R + Q_C]$.

The same caveats apply to the use of the inverse elasticity of the net residual demand curve when it is applied to step function residual demand curves such as those that exist in the New Zealand electricity market. Specifically, the researcher must choose percentage changes in the market-clearing quantity and then compute the implied change in the market price from the residual demand curve. To compute values of the two inverse elasticities that are internally consistent, the most straightforward way is to compute the inverse elasticity of the residual demand curve and use the above relationship that relates this magnitude to the inverse elasticity of the net residual demand curve. For the same reasons as described above for the case of the inverse elasticity of the residual demand curve, expected profit-maximizing offer behavior with fixed-price forward market obligations does not imply a deterministic relationship between the inverse elasticity of the net residual demand curve and the difference of the market price and the marginal cost of the supplier's highest cost generation unit operating in that period divided by the market price. For similar reasons, the inverse elasticity of the net residual demand curve is still a measure of the incentive of a supplier to exercise market power.

3.4. Pivotal Supplier and Net Pivotal Supplier as Measures of Unilateral Market Power

The residual demand curve and net residual demand curve can be used to derive additional measures of the ability and incentive of a supplier to exercise unilateral market power. Different from the inverse elasticity, these measures typically depend on the behavior of the residual demand curve and net residual demand curve at prices significantly higher than the

market-clearing price. As a consequence, these measures capture a more extreme ability and incentive to exercise unilateral market power.

Figure 3.22 shows the construction of a residual demand curve for the case in which the aggregate willingness-to-supply curve of all other suppliers reaches its capacity before system demand is met. As shown in the figure, this yields a residual demand facing the supplier that is positive for all possible prices. Because the real-time demand for electricity is perfectly inelastic and the production of electricity is subject to capacity constraints, it is possible for the residual demand curve facing a supplier to become perfectly inelastic at some positive output level. A supplier that faces a residual demand curve that is positive for all possible positive prices is said to be a pivotal supplier because some of its supply is necessary to serve the market demand regardless of the offer price.

The output level at which the supplier's residual demand curve become perfectly inelastic is called the pivotal quantity and it is shown in Figure 3.22 as the quantity associated with the vertical portion of the residual demand curve. Mathematically, a supplier is pivotal if $DR(p_{\max}) > 0$ where p_{\max} is the highest possible price that could occur in the market. The quantity $DR(p_{\max})$ is called the pivotal quantity. If a supplier is pivotal, this means that regardless of the offer price it submits, at least the pivotal quantity must be accepted from that supplier. A pivotal supplier has the ability to set the market price as high as it would like if it is willing sell only the pivotal quantity. Although a pivotal supplier clearly has a substantial ability to exercise unilateral market power, it may not have an incentive to do so because of its fixed-price forward market obligations. In particular, if the supplier's fixed-price forward market obligations exceed its pivotal quantity, $DR(p_{\max})$, then the supplier would have no incentive to exploit the fact that it is pivotal for the reason that it is a net buyer of energy at output levels equal to or below its pivotal quantity.

The net residual demand curve can be used determine whether a pivotal supplier would have an incentive to exploit the fact that it is pivotal. Specifically, if a supplier is net pivotal, then clearly it has such an incentive. A supplier is said to be net pivotal if $DR_F(p_{\max}) > 0$. The quantity $DR_F(p_{\max})$ is called the net pivotal quantity. By the definition of the net residual demand function, if a supplier is net pivotal and it has positive fixed-price forward market obligations, then the supplier is also pivotal. This means that regardless of the offer price it submits, at least $DR(p_{\max})$, the pivotal quantity (not the net pivotal quantity) of energy must be

accepted from the supplier. Different from a pivotal supplier, a net pivotal supplier has a very strong incentive to exercise unilateral market power the larger is the net pivotal quantity because it earns the short-term price on its net sales at the market-clearing price, $DR_F(p)$.

To summarize, a supplier can be pivotal and therefore have a significant ability to raise short-term prices. However, this supplier has little incentive to exploit its pivotal status if its fixed-price forward market obligations exceed its pivotal quantity, i.e., it is not net pivotal. Conversely, if a supplier is net pivotal, then it is also pivotal and has both a substantial incentive and ability to exercise unilateral market power. This incentive to exercise unilateral market power is greater the larger is the supplier's net pivotal quantity.

It is important to emphasize that a supplier cannot determine whether it is pivotal until the level of demand is realized and all supply offers of its competitors are known. Because the market rules require all suppliers to submit their offer at the same time and the market demand is not known when these offers are submitted, no supplier knows with certainty if it will be pivotal when it submits its offers. However, there are number of factors that can help suppliers predict when it might be pivotal. For example, an unexpectedly high level of demand or a large generation or transmission outage can create system conditions when one or more suppliers is pivotal.

Figures 3.23 to 3.25 depicts an example of the tradeoff that a supplier faces in deciding whether to submit offers into the short-term market to exploit the fact that it is pivotal. This figure considers the case of two residual demand curve realizations. For the low residual demand curve realization, $DR_L(p)$, the supplier is not pivotal. For the high residual demand curve realization, $DR_H(p)$, the supplier is pivotal. Let $0 < \theta < 1$ denote the probability of the high residual demand realization and $1 - \theta$ the probability of a low residual demand realization. Figure 3.23 draws $S_1(p)$, the expected profit-maximizing offer curve, for these two residual demand realizations assuming that the firm does not exploit the fact that it is pivotal for $DR_H(p_{\max})$.

However, if the probability of the high residual demand curve realization is sufficiently high, then it may be expected profit-maximizing for the supplier to exploit the fact that it is pivotal by submitting the willingness-to-supply curve, $S_2(p)$, shown in Figure 3.24, that crosses its residual demand curve at the point $DR_H(p_{\max})$. By doing so, the supplier forgoes the ability to sell any output if the low residual demand curve realization occurs. However, this may be

expected profit-maximizing if the probability of being pivotal times the profits the supplier earns from selling $DR_H(p_{\max})$ at p_{\max} exceeds the expected profits from submitting the willingness-to-supply curve in Figure 3.24 and selling $DR_H(p_H)$ in the high demand state and $DR_L(p_L)$. Let $C(q)$ denote the variable cost of producing output level q . An expected profit-maximizing supplier will decide to exploit the fact it is pivotal and submit an offer curve that sets p_{\max} in the high residual demand realization if the following inequality holds:

$$\theta(DR_H(p_{\max})p_{\max} - C(DR_H(p_{\max}))) > \theta(DR_H(p_H)p_H - C(DR_H(p_H))) + (1 - \theta)(DR_L(p_L)p_L - C(DR_L(p_L))),$$

meaning that the expected profits of selling in the high demand states at p_{\max} and selling zero in the low demand state exceeds the expected profits from selling at p_H in the high demand state and p_L in the low demand state.

Figure 3.25 shows an example of when the supplier is likely to find it expected profit-maximizing to submit $S_2(p)$ instead of $S_1(p)$. The long thin vertical expected profits from submitting $S_2(p)$, labeled $E(\pi(S_2(p)))$, is larger than the expected profits of submitting $S_1(p)$, labeled $E(\pi(S_1(p)))$. The above inequality illustrates several points about the likelihood a supplier will exploit its pivotal status. The higher the values of p_{\max} , the size of the supplier's pivotal quantity, and the probability the supplier is pivotal, the greater is the likelihood that the supplier will submit an offer curve that exploits the fact that it is pivotal.

Factoring in the impact of fixed-price forward market obligations complicates the analysis slightly although the basic insight about the determinants of when a supplier will exploit the fact that it is net pivotal remains. The supplier compares the expected profits from selling at p_{\max} during high residual demand curve realizations when it is net pivotal to the expected profit from submitting the expected profit-maximizing offer curve that does not exploit the fact that it is net pivotal. If the supplier assesses that the former expected profits are higher, then it will submit an offer curve that exploits the fact that it is pivotal. This logic suggests that when a supplier believes that the probability of being net pivotal is high, it will significantly increase its offer prices. The empirical validity of this prediction will be explored in the next section.

4.0. Empirical Evidence on How Suppliers Exercise Market Power

This section uses supplier offers, water reservoir levels, and market outcomes to demonstrate a number of empirical regularities in the behavior of the four large suppliers and market outcomes in the New Zealand market. First, summary statistics are presented on the behavior of half-hourly measures of both the unilateral ability and incentive to exercise unilateral

market power for each of the four large suppliers. These half-hourly measures of the ability and incentive to exercise unilateral market power are shown to be highly positively correlated with the value of the quantity-weighted average half-hourly market-clearing price.

To demonstrate that this observed positive correlation between the average half-hourly firm-level unilateral ability and incentive to exercise market power and half-hourly market prices is the direct result of market participant behavior, the second line of empirical evidence demonstrates that expected profit-maximizing offer behavior implies that a supplier's half-hourly offer price—the price at which it is willing to sell a pre-specified amount of energy to the short-term wholesale market—should be positively correlated with both its ability and incentive to exercise unilateral market power during that half-hour. Econometric analysis is then used to quantify the empirical relationship between the half-hourly offer price of each supplier and the half-hourly value of an index of that supplier's unilateral ability to exercise unilateral market power, after controlling for other exogenous factors impacting half-hourly market outcomes such as water levels and fossil fuel prices. Further econometric analysis examines the empirical relationship between the half-hourly offer price of each supplier and the half-hourly value of an index of that supplier's unilateral incentive to exercise unilateral market power. We find that when each of the four suppliers has a greater ability or greater incentive to exercise unilateral market power, they submit substantially higher half-hourly offer prices for a pre-specified quantity of energy.

4.1. Market Outcomes and the Unilateral Ability and Incentive to Exercise Market Power

Section 3 derived measures of the unilateral ability and incentive of a supplier to exercise market power that can be computed on a system-wide basis or separately for the North and South Islands using the half-hourly level of demand and the willingness-to-supply curves of all market participants. In this section, we derive modifications of these measures that the theory of expected profit-maximizing offer behavior derived in Section 3 implies should be related to the half-hourly market-clearing price.

As shown in Section 3.2, the form of the residual demand curve that a supplier faces determines its ability to exercise unilateral market power. The inverse of the elasticity of the residual demand curve evaluated at the market-clearing price is one measure of the ability of a supplier to exercise unilateral market power. This inverse elasticity measures the percent change

in the market-clearing price that would result from the supplier producing one percent less output than it actually produced during that half-hour period.

Under the simplified model of expected profit-maximizing offer behavior described in Figures 3.12 and 3.13, this inverse elasticity measure can be directly related to the market-clearing price and the marginal cost of the highest cost unit owned by that supplier operating during that half-hour period. The logic underlying the construction of the expected profit-maximizing offer curve in Figure 3.12 implies that the point (P_1, Q_1) is the ex post profit-maximizing price/quantity pair for the firm for the residual demand realization $DR_1(p)$ and the point (P_2, Q_2) is the ex post profit-maximizing price/quantity pair for the firm for the residual demand realization $DR_2(p)$. The first-order conditions for ex post profit-maximization for these two residual demand realizations are:

$$(P_1 - C_1)/P_1 = -1/\epsilon_1 \text{ and } (P_2 - C_2)/P_2 = -1/\epsilon_2 \quad (4.1)$$

where C_i ($i=1,2$) is the marginal cost for supplier i at output level Q_i ($i=1,2$) and $-1/\epsilon_i$ ($i=1,2$) is the inverse of the elasticity of the residual demand curve for that residual demand realization.

Recall that the inverse elasticity is defined in terms of the residual demand curve as:

$$-1/\epsilon_i = [DR_i(P_i)/P_i] * [1/DR_i'(P_i)] \quad (4.2)$$

where $DR_i'(P_i)$ is the slope of residual demand curve i evaluated at price P_i , and $DR_i(P_i)$ is the value of residual demand curve evaluated at price P_i . Using this definition of the inverse elasticity, the two equations in (1) can be rearranged to equal:

$$P_i = C_i - [DR_i(P_i)/DR_i'(P_i)], \quad i=1,2. \quad (4.3)$$

Equation (4.3) implies that the market-clearing price is equal to the marginal cost of the highest cost unit owned by that supplier operating during that half-hour plus the level of the residual demand curve divided by the absolute value of the slope of the residual demand curve.

Define η_i ($i=1,2$), the inverse semi-elasticity of the residual demand curve i , as:

$$\eta_i = - (1/100)[DR_i(P_i)/DR_i'(P_i)]. \quad (4.4)$$

This magnitude gives the \$/MWh increase in the market-clearing price associated with a one percent reduction in the amount of output sold by the supplier. In terms of this notation, equation (4.3) becomes

$$P_i = C_i + 100\eta_i, \quad i=1,2. \quad (4.5)$$

Thus, the simplified model of expected profit-maximizing offer behavior implies that higher market-clearing prices should be associated with higher values of the inverse semi-elasticity.

As discussed in Section 3, because offer curves in the New Zealand wholesale market are step functions, residual demand curve realizations do not strictly satisfy the assumptions implied by the simplified model of expected profit-maximizing offer behavior presented there, so that equation (4.5) will not hold with equality. However, the general model of expected profit-maximizing offer behavior described in Section 3 implies that when a supplier has a greater ability to exercise unilateral market power as measured by the size of η_i , the \$/MWh price increase that results from reducing the amount it sells in the wholesale market by one percent, that supplier's offer price is likely to be higher.

Computing the slope of the residual demand curve at the market-clearing price for a step-function residual demand curve requires choosing the output change used to compute the finite-difference approximation to the slope. These output changes should be large enough to ensure that enough price steps on the residual demand curve are crossed so that a non-zero slope is obtained, but not too large that the implied output change is judged as implausible for the supplier to implement. We also want to choose a procedure for selecting the output changes to ensure that the value of slope obtained is not sensitive to the size of the output changes used to compute it.

Figure 4.1 describes the details of the process we use to compute the slope of the residual demand curve for Firm B for a peak half-hour period in February 2006. Suppose that $Q^* = 901$ MW is the output sold by Firm B at the market-clearing price for this half-hour period of $P^* = \$145/\text{MWh}$. We want to approximate the slope of the residual demand curve in the vicinity of (P^*, Q^*) . Consider a 10% price change window on either side of P^* , and look for the closest steps on the residual demand curve to (P^*, Q^*) that lie outside this 10% price window. The closest point below P^* that has price less than 0.9 times P^* is $(\$129, 969)$. Call this point (P_1, Q_1) . Above P^* the closest point with price greater than 1.1 times P^* is $(\$164, 871)$. Call this point (P_2, Q_2) . The slope of the residual demand curve $DR(P^*)$ at (P^*, Q^*) according to this procedure is given by the formula:

$$DR'(P^*) = (Q_1 - Q_2)/(P_1 - P_2) = (969 - 871)/(129 - 164) = -2.8 \quad (4.6)$$

The resulting inverse semi-elasticity at (P^*, Q^*) for this residual demand curve gives the \$/MWh price increase from a 1% reduction in output and is equal to:

$$\eta = -(1/100)DR(P^*)/DR'(P^*) = - (1/100)Q^*/ DR'(P^*) = -(1/100) 901/(-2.81) = 3.21. \quad (4.7)$$

This semi-elasticity quantifies the ability of Firm B to raise prices during this half-hour period by reducing its output by 1%. This magnitude implies that if Firm B reduces its output by 1% relative to $Q^* = 901$ MW, the increase in the market price would be \$3.21/MWh. Figure 4.2 shows the same calculation for Firm B in the half-hour period exactly one year later. For this period, $-(1/100)Q^*/DR'(p) = 0.29$. That is, a 1% reduction in output would produce an increase in the market price of \$0.29/MWh, a significantly lower price increase from the same 1% output reduction. Note that this inverse semi-elasticity is significantly lower, despite the fact that Firm B is producing over 180 MW more during this half-hourly period than in the same period during 2006. These two figures demonstrate the usefulness of the inverse semi-elasticity as a measure of the ability of a supplier to exercise unilateral market power because high and low values of this measure can occur for both high and low output levels of the supplier.

To demonstrate the robustness of our inverse semi-elasticity estimates to the price change window used for the calculation, Table 4.1 compares the results from calculating the inverse semi-elasticity for the four large suppliers in each half-hour from January 1, 2001 to June 30, 2007, using four different values for the price change window: 1%, 5%, 10% and 15%. For each supplier, the overall mean value of the semi-elasticity is shown for each price window. For example, using a price window of +/-15% the mean inverse semi-elasticity for Firm C is 1.27, compared to a mean inverse semi-elasticity of 1.21 using a price window of +/-1%. The table also shows the Pearson correlation coefficients between the half-hourly inverse semi-elasticities calculated for different price windows. These show that there is high correlation (in all cases greater than 0.80) between the values calculated for different price windows. This provides strong empirical evidence that our inverse semi-elasticities are not sensitive to the choice of the price window used to compute them. For the remainder of this chapter all results are shown based on the inverse semi-elasticities calculated with a 10 percent price window.

To compare time series behavior of the inverse semi-elasticities across firms, Figure 4.3 plots the 30-day moving average of the half-hourly values of the inverse semi-elasticities for the four largest firms from January 1, 2001 to June 30, 2007. The half-hourly inverse semi-elasticities follow a very similar pattern across the four firms and certain suppliers have persistently larger values than other suppliers. The maximum value of the smoothed inverse semi-elasticities shown in the figure is 10, with the values for Firm A peaking at close to 20 during early 2003 and the peak values for Firm C for this time period also exceeding 10. Over

the entire sample period, Firm A's smoothed inverse semi-elasticities tend to be the highest, followed by Firm C, then by Firm B, and finally by Firm D.

To provide a clear picture of the magnitude of persistent differences across the four suppliers in this index of the ability to exercise unilateral market power, Figure 4.4 presents the sample mean of the half-hourly values of η_{ihd} , the semi-elasticity for supplier i during half-hour h of day d . Each point on the graph in Figure 4.4 for supplier i is equal to $\eta_{ih}(\text{mean}) = \frac{1}{D} \sum_{d=1}^D \eta_{ihd}$, where D is the total number of days in the sample period of January 1, 2001 to June 30, 2007. Firm A has the highest value of $\eta_{ih}(\text{mean})$ for all half-hours and Firm D the lowest for all half-hours. Firm B is slightly higher than Firm D for all half-hours and the values for Firm C are roughly midway between the values for Firm B and Firm A.

It is important to emphasize that these inverse semi-elasticities only depend on the quantity sold in the market and the form of the residual demand curve faced by supplier under consideration at that quantity. For example, the inverse semi-elasticity for Firm A for a given half-hour depends only on the half-hourly market-clearing quantity, the half-hourly offer curves of all other suppliers besides Firm A, and the level of system demand during that half-hour, but it measures the \$/MWh increase in the half-hourly market price that would result from Firm A supplying 1% less output during that half-hour.

To demonstrate the very close relationship between half-hourly market-clearing prices and the half-hourly ability of the four large suppliers to exercise unilateral market power (as measured by the inverse semi-elasticity of their residual demand curves), Figure 4.5 plots the 30-day moving average of the half-hourly values of the quantity-weighted average of the nodal prices and a 30-day moving average of the half-hourly values of the unweighted average of the four values η_{ihd} for Firms A to D, which is equal to $\eta_{hd}(\text{firm}) = \frac{1}{4} \sum_{i=1}^4 \eta_{ihd}$. Define p_{hdm} as the price at node m during half-hour h of day d and q_{hdm} as the total amount of energy injected at node m during half-hour h and day d . Figure 4.5 shows that the time series pattern of $p_{hd}(\text{avg}) = \frac{\sum_{m=1}^M p_{hdm} q_{hdm}}{\sum_{m=1}^M q_{hdm}}$, the quantity-weighted average of the nodal prices for half-hour h of day d , closely tracks $\eta_{hd}(\text{firm})$. During periods when the average index of the ability of these suppliers to exercise unilateral market power is high, the quantity-weighted average of the nodal prices they are paid is also very high. Specifically, during mid-2001, early 2003, and early 2006 the average index of the ability of suppliers to exercise unilateral market power is high and the

quantity-weighted average nodal price is high. Conversely, during periods when the average index of the ability of these suppliers to exercise unilateral market power is low, the quantity-weighted average of the nodal prices is significantly lower. This occurs during 2002, 2004, and 2005.

Figure 4.6 plots that the sample half-hourly means of $\eta_{hd}(\text{firm})$, $\frac{1}{D} \sum_{d=1}^D \eta_{hd}(\text{firm})$, and the sample half-hourly means of the quantity-weighted average nodal prices, $\frac{1}{D} \sum_{d=1}^D p_{hd}(\text{avg})$, for our sample period. The average pattern throughout the day of the average half-hourly market-wide ability of the four suppliers to exercise unilateral market power very closely tracks the average half-hourly pattern of the quantity-weighted average price throughout the day. Figure 4.6 clearly demonstrates that over our sample period from January 1, 2001 to June 30, 2007, a greater average half-hourly ability of each supplier to exercise unilateral market power is coincident with a higher average half-hourly market-clearing price.

Even if a supplier possesses a substantial ability to exercise unilateral market power, it may not submit willingness-to-supply curves that reflect this ability if it has no incentive to exercise unilateral market power. As shown in Section 3, a supplier with fixed-price forward market obligations approximately equal to its sales in the short-term wholesale market has little incentive to exercise unilateral market power, even if it has a substantial ability to do so. This logic suggests that half-hourly measures of the unilateral incentive of each supplier to exercise unilateral market power should be correlated with both market-clearing prices and the level of offer prices that each supplier submits.

Inverse semi-elasticities for the net-of-forward market obligations residual demand curves can be computed from these inverse semi-elasticities to obtain measures of the incentive (as opposed to ability) of individual suppliers to exercise unilateral market power. Under the simplified model of expected profit-maximizing offer behavior described in Section 3, the inverse semi-elasticities of the net-of-forward obligations residual demand curve can be directly related to the market-clearing price and the marginal cost of the highest cost unit owned by that supplier operating during that half-hour period.

The logic underlying the construction of the expected profit-maximizing offer curve with forward market obligations drawn in Figure 3.21 implies that the point of intersection between the offer curve and each residual demand realization is an ex post profit-maximizing price/quantity pair for the firm for each residual demand realization given the forward market

obligations of the supplier, Q_C . For the two residual demand curve realizations in Figure 3.21, the first-order conditions for ex post profit-maximization for these two residual demand realizations are:

$$(P_1 - C_1)/P_1 = -1/\epsilon_1^C \text{ and } (P_2 - C_2)/P_2 = -1/\epsilon_2^C \quad (4.8)$$

where C_i ($i=1,2$) is the marginal cost for supplier i at the output level Q_i ($i=1,2$) and $-1/\epsilon_i^C$ ($i=1,2$) is the inverse elasticity of the net-of-forward market obligations residual demand curve for that residual demand realization.

Recall from Section 3 that the inverse elasticity of the net-of-forward market obligations residual demand curve at price P_i and forward market obligation Q_C is equal to:

$$-1/\epsilon_i^C = [(DR_i(P_i) - Q_C)/P_i] * [1/DR_i'(P_i)] = -1/\epsilon_i [(DR_i(P_i) - Q_C)/DR_i(P_i)]. \quad (4.9)$$

The first equality defines this inverse elasticity and the second demonstrates that it is equal to the inverse elasticity of the residual demand curve multiplied by the firm's exposure to short-term market prices. This exposure is measured by the difference between the supplier's short-term market sales, $DR_i(P_i)$, and its forward market obligations, Q_C , divided by its short-term market sales.

Using this definition of the inverse elasticity net-of-forward market obligations, the two equations in (4.8) can be rearranged to equal:

$$P_i = C_i - [(DR_i(P_i) - Q_C)/DR_i'(P_i)], \quad i=1,2. \quad (4.10)$$

Equation (4.10) implies that if an expected profit-maximizing supplier has fixed-price forward market obligations, the market-clearing price is equal to the marginal cost of the highest cost generation unit operating during that half-hour owned by the supplier plus the value of the net-of-forward market obligations residual demand curve, $DR_i^C(p) = (DR_i(P_i) - Q_C)$, divided by the slope of this residual demand curve.

Define η_i^C ($i=1,2$), the net inverse semi-elasticity of the net-of-forward market obligations residual demand curve i , as:

$$\eta_i^C = - (1/100)[(DR_i^C(P_i)/DR_i^{C'}(P_i))] = \eta_i[(DR_i(P_i) - Q_C)/DR_i(P_i)]. \quad (4.11)$$

The first equality defines η_i^C in terms of the net of fixed-price forward market obligations residual demand curve. The second equality demonstrates that it is equal to the inverse semi-elasticity of the residual demand multiplied by the supplier's exposure to short-term prices. This value of η_i^C gives the \$/MWh increase in the market-clearing price associated with a one percent reduction in the net position of the supplier, the difference between its short-term market sales

and its fixed-price forward market obligations. In terms of this notation, equation (4.10) becomes

$$P_i = C_i + 100\eta_i^C, \quad i=1,2. \quad (4.12)$$

This equation demonstrates that the simplified model of expected profit-maximizing offer behavior with fixed-price forward market obligations implies that higher offer prices and higher market-clearing prices are associated with higher values of the inverse semi-elasticity of the net-of-fixed price forward market obligations residual demand curve after controlling for the variable cost of the highest cost generation unit in that supplier's portfolio of generation units operating during that half-hour period, C_i in equation (4.12).

To compute the half-hourly value of the inverse semi-elasticity of the net-of-forward market obligations residual demand curve for each of the four largest suppliers, we use the second equality in equation (4.11) which computes this index of the incentive of a supplier to exercise unilateral market power by multiplying the inverse semi-elasticity of the residual demand curve by that supplier's exposure to short-term wholesale prices at the market-clearing price P^* , $(DR(P^*) - Q_C)$, divided by the supplier's short-term market sales, $DR(P^*)$. This approach to computing η_i^C ensures that the same estimate of the slope of the step-function residual demand curve is used to compute both η_i and η_i^C .

As discussed in Section 3, the assumptions required for the validity of the simplified model of expected profit-maximizing offer behavior with fixed-price forward market obligations do not hold because suppliers submit non-decreasing step functions rather than increasing continuous functions as their willingness-to-supply curves. It is important to emphasize that even if the assumptions necessary for the strict validity of the simplified model of expected profit-maximizing offer behavior do not hold, η_i^C is still a valid measure of the half-hourly incentive of a supplier to exercise unilateral market power. It equals the \$/MWh increase in the market-clearing price that results from the supplier a 1% net position than it actually had during that half-hour period. As shown in the first-equality of equation (4.11), this measure depends on the half-hourly offers of all other suppliers and the supplier's short-term market sales minus its fixed-price forward market obligation.

Figure 4.7 graphs the 30-day moving average of the net inverse semi-elasticities over the sample period of January 1, 2001 to June 30, 2007 computed as described above. For the value of Q_C in equation (11), we use the half-hourly value of the retail load obligation of that supplier.

Because there is a small, but sometimes important, fixed-price forward contract market in New Zealand and a small amount of retail load pays a retail price that varies with the half-hourly wholesale price, there is the potential for a small amount of measurement error between the true value of Q_C and the supplier's retail load obligation.

Figure 4.7 demonstrates the mitigating influence of fixed price forward contracts on the ability of suppliers to exercise unilateral market power. All of the inverse semi-elasticities of the residual demand curve are reduced significantly in absolute value as a result of multiplying them by the half-hourly value of the net exposure of the supplier to short-term prices, $[(DR_i(P_i) - Q_C)/DR_i(P_i)]$. This net exposure can be negative if the supplier sells less in the short-term market than its fixed-price forward market obligations, Q_C . This explains why some of the smoothed values of η_i^C are negative for certain suppliers during portions of the sample period.

As shown in Figure 4.3, all four suppliers had more than double the ability to exercise unilateral market power in early 2003 relative to mid-2001, as measured by smoothed half-hourly semi-elasticities during the two time periods. Only Firm C translated this larger ability into a large incentive to raise short-term prices as measured by the value of η_i^C . Consequently, one explanation for the slightly longer period of higher prices that prevailed during mid-2001 is that a larger number of suppliers had a significant incentive to exercise unilateral market power during mid-2001 versus early 2003.

Figure 4.8 plots the 30-day moving average of the half-hourly values of the quantity-weighted average of the nodal prices and a 30-day moving average of the half-hourly values of $\eta_{hd}^C(firm)$. Figure 4.8 shows that the time series pattern of $p_{hd}(avg)$, the quantity-weighted average of the nodal prices for half-hour h of day d , closely tracks the time series pattern $\eta_{hd}^C(firm)$. During the half-hour periods when this average index of the incentive of these suppliers to exercise unilateral market power is larger, the quantity-weighted average of the nodal prices is high. Specifically, during mid-2001, early 2003, and early 2006 the average index of the incentive of suppliers to exercise unilateral market power is high and the quantity-weighted average nodal price is high. Conversely, during periods when the average index of the incentives of these suppliers to exercise unilateral market power is close to zero, the smoothed quantity-weighted average of the nodal prices is significantly lower. This occurs during 2002, 2004, and 2005.

This section has shown that both the ability and incentive of all four suppliers to exercise unilateral market power are positively correlated with market-clearing prices. The ability to exercise unilateral market power is clearly a necessary condition for a supplier to exercise unilateral market power because a supplier must face an upward-sloping residual demand curve to be able to raise market prices by withholding its output. However, even a supplier with a substantial ability to exercise unilateral market power may not exploit this ability unless it has an incentive to do so. As noted above, the difference between a supplier's short-term market sales and its fixed-price forward market obligations determines the supplier's incentive to exercise unilateral market power.

5. Offer Behavior and Ability and Incentive to Exercise Market Power

The previous section has demonstrated that the ability and incentive to exercise unilateral market power is very highly correlated with the level of market prices. This section explores the extent to which this relationship is due to suppliers exercising unilateral market power by raising their offer prices during periods when they have an increased ability and incentive to exercise market power. As discussed in Section 3, the theory of expected profit-maximizing offer behavior implies that suppliers exercising all available unilateral market power will submit higher offer prices when they have a greater ability and incentive to exercise unilateral market power. This section provides empirical confirmation for this implication of expected profit-maximizing behavior.

We find that after controlling for differences over days of the sample and half-hours of the day or half-hours of the day during each month of our sample period in an individual supplier's opportunity cost of producing electricity from their generation units, higher values of three different indexes of a supplier's unilateral ability to exercise market power are associated with a higher offer price for the quantity of energy dispatched during that half-hour period by that supplier. A similar statement holds for three analogous indexes of the supplier's unilateral incentive to exercise market power. After controlling for opportunity cost differences over time, higher values of each index of the incentive to exercise unilateral market power are associated with a higher offer price for the quantity of energy dispatched during that half-hour period by that supplier. The absolute values of the regression coefficient estimates associated with the incentive of a supplier to exercise unilateral market power are uniformly higher for all market participants than the corresponding coefficient estimates for the regressions using the unilateral

ability measure. This outcome is consistent with the discussion in Section 3 that the incentive to exercise unilateral market power is a key determinant of a supplier's offer price if it has significant fixed-price forward market obligations, as is the case for all of four large suppliers under consideration.

In order to describe our empirical analysis a definition of a supplier's half-hourly offer price is required. Figure 4.9 presents the actual offer curve for Firm A for a half-hour period in February 2006. The dispatched quantity of energy for Firm A during that half-hour is 1,508 MW. The offer price along Firm A's willingness-to-supply curve for that half-hour period is found by extending a vertical line up from the horizontal axis at 1,508 MW until it intersects Firm A's willingness-to-supply curve. In this case, the offer price for the dispatched quantity for Firm A is equal to \$145/MWh, which is the offer step directly above the quantity level 1,508 MW. In general, the offer price for output level Q^* for supplier k during half-hour period h is computed as the solution to the following equation in P : $Q^* = S_{hk}(P)$, where $S_{hk}(P)$ is supplier k 's willingness-to-supply curve during half-hour period h .

As equations (4.5) and (4.12) in Section 4.1 demonstrate, the simplified model of expected profit-maximizing offer behavior by a supplier facing a distribution of downward sloping residual demand curves implies that, after controlling for the opportunity cost of the highest cost generation unit operating during that half-hour period (the term C_i in these two equations), a supplier's offer price at the quantity of energy that it sells in the short-term market should be an increasing function of the value of the inverse semi-elasticity, if the supplier has no fixed-price forward market obligations, and increasing in the net inverse semi-elasticity if the supplier has fixed-price forward market obligations. Although the conditions necessary for the strict validity of the simplified model of expected profit-maximizing offer behavior outlined in Section 3 do not hold for the New Zealand market, we still expect these two implications of the model to hold. Specifically, when a supplier has a greater unilateral ability or incentive to exercise unilateral market power, after controlling for its opportunity cost of selling energy from its highest cost generation unit operating during that hour, the offer price it sets for the amount of energy that it sells in the short-term market should be higher.

Let $P_{jhd m}(\text{actual})$ equal the offer price at the actual level of output sold by supplier j during half-hour h of day d during month of sample m , $\eta_{jhd m}$ the inverse semi-elasticity of supplier j 's residual demand curve during half-hour h of day d during month of sample m and

η_{jhd}^C the inverse net semi-elasticity of supplier j 's net-of-forward-market-obligation residual demand curve during half-hour h of day d during month of sample m . We take two approaches to controlling for differences across half-hours during our sample period in the variable cost of the highest cost generation unit owned by that supplier operating during that half-hour period. The first approach assumes that this variable cost can be different for each supplier for every day during our sample period and each half-hour during the day. The following regressions are estimated for each supplier j :

$$P_{jhd}(\text{offer}) = \alpha_{dmj} + \tau_{hj} + \beta_j \eta_{jhd} + \varepsilon_{jhd} \text{ and } P_{jhd}(\text{offer}) = \gamma_{dmj} + \mu_{hj} + \delta_j \eta_{jhd}^C + v_{jhd}, \quad (5.1)$$

where the α_{dmj} and γ_{dmj} are day-of-month d for month-of-sample m fixed effects and the τ_{hj} and μ_{hj} are half-hour-of-the-day fixed effects. The ε_{jhd} and v_{jhd} are mean zero and constant variance regression errors. Input fossil fuel prices and water levels change at most on a daily basis. Because there is a different fixed effect for each day and month combination during our sample period, these fixed effects completely account for the impact of daily changes in fossil fuel prices and water levels during our sample period on the variable cost of the highest cost generation unit owned by supplier j that is operating during each half-hour period in the day. Consequently, these day-of-sample fixed-effects completely control for any differences across days of the sample in input fossil fuel prices and water levels. The half-hourly fixed-effects account for differences across half-hours of the day in this variable cost. This strategy for controlling for variable cost changes across half-hours of the sample implies more than 2,400 possible variable cost values over the sample period for each supplier. Multiplying this figure by four implies more than 9,600 possible variable costs of the highest cost generation unit operating during a half-hour that could set the market-clearing price during our sample.

Our second strategy for controlling for the opportunity cost of producing electricity from the highest variable cost unit operating during half-hour period-of-the-day h during month of the sample m for supplier j uses different half-hour-of-the-day fixed-effects for each month of the sample period. The two equations estimated are:

$$P_{jhd}(\text{offer}) = \alpha_{hmj} + \beta_j \eta_{jhd} + \varepsilon_{jhd} \text{ and } P_{jhd}(\text{offer}) = \gamma_{hmj} + \delta_j \eta_{jhd}^C + v_{jhd}, \quad (5.2)$$

where α_{hmj} and γ_{hmj} are half-hour-of-the-day for each month-of-the-sample fixed effects to control for the differences in the opportunity cost of producing electricity from the highest variable cost unit operating during half-hour period-of-the-day h during month-of-the-sample m for supplier j . The ε_{jhd} and v_{jhd} are once again mean zero and constant variance regression

errors. Because there are 48 half-hour periods in the day and 78 months during our sample period from January 1, 2001 to June 30, 2007, there are $48 \times 78 = 3,744$ values of the α_{hmj} and the same number of values of the γ_{hmj} for each supplier j . These fixed-effects imply that the variable cost of producing electricity from the highest cost generation unit operating during half-hour 12 in month 3 of the sample period can be different from this same variable cost during all other months of the sample period. Moreover, the variable cost of producing electricity from the highest cost generation unit operating during half-hour 12 in month 3 can differ from the variable cost of producing electricity in any other half-hour of any other month of the sample period, including month 3.

These fixed-effects allow for a substantial amount of variability in the time path of the variable cost of the highest cost unit operating in the North and South Island of New Zealand during each half-hour of our sample period. There are 3,744 fixed effects for each supplier to account for differences in the variable cost of the highest cost unit in their portfolio operating during each half-hour of the sample period. Multiplying this figure by 4 implies 14,976 different possible variable costs of the highest cost unit operating owned by the four large suppliers that could set prices during our sample period.

The fixed-effects in model (5.1) and model (5.2) should be more than sufficient to account for differences in the variable cost of the highest cost generation unit operating during each half-hour of the sample period in the portfolio of generation units owned by each of the four large suppliers. The opportunity cost of producing electricity from hydroelectric generation units should not differ significantly across half-hours of the day or days of the month in a hydroelectric dominated system. The opportunity cost of water depends on current water storage levels and the distribution of future water inflows and outflows. New information about these variables arrives daily, but the best estimates of future inflows and outflows changes slowly as do water storage levels. Our day-of-sample fixed effects are more than sufficient to account for changes in the opportunity cost of water over our sample period.

The variable cost of producing electricity from individual fossil fuel generation units is unlikely to change significantly during individual months of our sample period, which implies that fixed-effects that allow these half-hourly variable costs to change each month of the sample period should provide for far more fluctuations in the variable cost of the highest cost unit producing electricity during each half-hour of our sample period than is likely to be necessary to

capture the amount of variability that actually exists in these variable costs. Regressions of model (5.1) including the value of the relevant daily fossil fuel price and daily water levels to account for daily changes in the variable cost of operating fossil fuel generation units and daily changes in the opportunity cost of water did not quantitatively change any of our results. This outcome is not surprising given the high level of agreement between our estimates of β_j and δ_j using day-of-sample and half-hour-of-the-day fixed-effects and different half-hour-of-the-day fixed effects for each month of the sample period.

Table 5.1 presents the estimated values of β_j and δ_j and the estimated standard errors for each of the four suppliers using the day-of-sample and half-hour-of-the-day fixed-effects. Table 5.2 presents estimates of the same parameter values for the different half-hour-of-the-day fixed effects for each month of the sample period. The values of β_j and δ_j are positive, precisely estimated and economically meaningful for all regressions. Focusing on the day-of-sample and half-hour-of-the-day fixed-effects model, holding all other factors constant, if the residual demand curve faced by Firm C has an inverse semi-elasticity that is one unit higher, the offer price associated with the amount of output that it sells in the short-term market is predicted to be \$1.41/MWh higher, because of the greater ability Firm C has to exercise market power implied by the inverse semi-elasticity of its residual demand curve.

Table 5.3 computes the half-hourly sample mean and standard deviation of $\eta_{j\text{hdm}}$ for each h. For each supplier, a row of the table is the sample mean and sample standard deviation across all days and months of our sample period of the value $\eta_{i\text{hdm}}$ for that half-hour of the day. This table can be used to demonstrate the economic significance of our estimates of β_j . For example, for Firm C, the standard deviation of $\eta_{j\text{hdm}}$ for h=37 is equal to 6.811. This implies that holding opportunity cost of water and the price of the input fossil fuel constant, a one standard deviation change in the value of $\eta_{j\text{hdm}}$ for half-hour 37 implies a \$9.60/MWh higher offer price and a two standard deviation change a \$19.20/MWh higher offer price according to the parameter estimates in Table 5.1. For Firm A, the mean and variance of the inverse semi-elasticities over the sample period are even higher. The value of β_j for Firm A implies that a one standard deviation change in the value of the inverse semi-elasticity of its residual demand curve during half-hour 23, holding all other factors constant, implies an offer price increase of \$4.50/MWh. Changes of this magnitude in the value of its inverse semi-elasticity for half-hour 23 for Firm A during our sample period are not unusual.

For Firm D the value of β is significantly higher than it is for all of the other suppliers, on the order of \$3.81/MWh. However, as shown in Table 5.4 the mean value of the inverse semi-elasticity is the lowest of all of the suppliers and the variance is also the smallest. Nevertheless, the magnitude of β for Firm D implies that even for one standard deviation changes in the value of its inverse semi-elasticity, economically significant changes in Firm D's offer price are predicted to occur because of its increased ability to exercise unilateral market power.

The values of δ , the coefficient associated with $\eta_{j\text{hdm}}^C$, the inverse semi-elasticity of the net of forward market obligations residual demand curve, are substantially larger in absolute value than the corresponding value of β , the coefficient associated with $\eta_{j\text{hdm}}$, for all suppliers. The value of δ for Firm C implies that if the value of the inverse semi-elasticity of the net forward market obligations residual demand curve for Firm C increases by one unit, then Firm C's offer price for the amount it sells in the short-term market is predicted to increase by \$4.31 because of the substantially greater incentive Firm C has to exercise unilateral market power. Table 5.4 lists the half-hourly sample means and standard deviations of $\eta_{j\text{hdm}}^C$ for each supplier. This table demonstrates that a one unit change in the value of $\eta_{j\text{hdm}}^C$ is a fairly frequent occurrence. For a number of half-hours of the day, a 3 unit change in $\eta_{j\text{hdm}}^C$ is less than a two standard deviation change. For example, during half-hour 37, a two standard deviation change in the value of $\eta_{j\text{hdm}}^C$ implies a more than \$20/MWh increase in Firm C's offer price.

It is important to emphasize that different from the case of inverse semi-elasticity of the residual demand curve, which can only be positive, the inverse semi-elasticity of the net of forward market obligations residual demand curve can be negative if the supplier's fixed-price forward market obligations exceed the amount of energy that it sells in the short-term market. As shown in Figure 4.7, this was frequently the case for Firm A as well as for Firm B and Firm D during the sample period. The results in Table 5.1 for Firm A imply that, keeping all other factors constant, if a negative value of $\eta_{j\text{hdm}}^C$ for Firm A becomes larger in absolute value by one unit, Firm A's offer price is predicted to be \$5.08/MWh lower because of its greater incentive to exercise unilateral market power by driving the price down. As shown in Table 5.4, a one unit change in $\eta_{j\text{hdm}}^C$ is less than a one standard deviation change for many half-hours of the day. The results in Table 5.1 also imply that, keeping the opportunity cost of water and the price of the input fossil fuel constant, if the value of the inverse semi-elasticity of the net-of-forward-market-obligations residual demand curve facing Firm A increases by one unit, the offer price for the

amount of energy it sold in the short-term market is \$5.08/MWh higher because of the greater incentive Firm A has to exercise unilateral market power.

Thus, once fixed price forward contract obligations are introduced into a wholesale market, suppliers with the ability to exercise unilateral market power can do so either by increasing or decreasing prices. A supplier with a substantial ability to exercise unilateral market power that is net short relative to its forward market obligations, meaning that it has more fixed-price forward market obligations than the amount of energy it sold in the short-term market, has an incentive to exercise market power by driving down the wholesale price, which reduces the cost of closing out its net short position through purchases from the short-term market. The results shown in Table 5.1 confirm this for logic for all suppliers. Alternatively, when a supplier is long relative to its forward market position, meaning that its sales in the short-term market exceed its fixed-price forward market obligations, a higher value of the $\eta_{j\text{hdm}}^C$ implies that it will raise its offer price because it has an incentive to use its ability to exercise market power to raise the market-clearing price.

The estimate for δ_j for Firm D is by far the largest of the five values reported in Table 5.1. However, as shown in Table 5.4 the standard deviations of the inverse elasticity of the net of fixed-price forward market obligations for Firm D are very small in absolute value relative to the values for the other three suppliers. Nevertheless, even multiplying the estimate of δ_j for Firm D by a one standard deviation change in the value of its inverse elasticity yields predicted offer price changes of more than \$10/MWh for many half-hours of the day. Because the $\eta_{j\text{hdm}}^C$ for Firm D takes on both positive and negative values during the sample period, there are times when Firm D submits a substantially lower offer price, all other factors held constant, because it has an incentive to use its ability to influence market prices to lower the market-clearing price because its short-term market sales are less than its forward market obligations. Alternatively, when it is long relative to its forward market position, a higher value of the $\eta_{j\text{hdm}}^C$ for Firm D implies that it will raise its offer price because it has an incentive to use its ability to exercise market power to raise the market-clearing price.

It is important to emphasize that the goal of our modeling effort is to determine whether higher offer prices are systematically associated with higher values of $\eta_{j\text{hdm}}$ and $\eta_{j\text{hdm}}^C$ and whether the magnitude of this relationship is economically significant. The results of our analysis presented in Tables 5.1 and 5.2 provide strong confirmation of a positive and

economically significant relationship between a supplier's half-hourly offer price and the half-hourly values of $\eta_{j\text{hdm}}$ and $\eta_{j\text{hdm}}^C$. The magnitude of this relationship is substantially larger for the measure of the incentive to exercise unilateral market power relative to the measure of the ability to exercise unilateral market power. This result is consistent with the logic in Section 3 that a supplier with the ability to exercise unilateral market power must also have the incentive to do so in order to find it expected profit-maximizing to submit offer prices that exploit it.

It is important to emphasize that the regressions (5.1) and (5.2) are predictive regressions in the sense discussed in Reiss and Wolak (2007). As noted above, the economic theory of expected profit-maximizing offer behavior described in Wolak (2003a, and 2007) does not imply these regressions yield the precise causal relationship between half-hourly offer prices and the half-hourly indexes of the ability and incentive of market participants to exercise unilateral market power. This fact does not invalidate the interpretation of these regressions as providing predictive statistical evidence consistent with the view that after controlling for the level of input fossil fuel prices and the opportunity cost of water, when any of the four suppliers has a greater ability or incentive to exercise unilateral power market power as measured by these indexes, each supplier submits a significantly higher half-hourly offer price and this higher offer price results in substantially higher market-clearing price.

6. Analysis with Pivotal Measures of the Ability and Incentive to Exercise Market Power

We now present an analysis of the relationship between a supplier's offer price and indexes of the ability and incentive to exercise market power based whether the supplier is pivotal and net pivotal as defined in Section 3. Although there is no simple relationship between a supplier's offer price and its status as a pivotal supplier or net pivotal supplier that can be derived from the assumption of expected profit-maximizing offer behavior, periods when a supplier expects it is pivotal or net pivotal are likely to cause it to raise its offer price, particularly for the pivotal quantity of energy. In fact, a number of the market power mitigation mechanisms in United States wholesale markets are based on this supposition. The short-term market operator takes the offers and bids of all market participants and determines whether a supplier is pivotal or a set of suppliers are jointly pivotal. If this is the case then the offers of this supplier or this set of suppliers are mitigated to some reference offer level that is based on that supplier's variable cost of production. Our analysis tests whether suppliers recognize that an increased likelihood of being pivotal or net pivotal causes them to raise their offer prices.

Recall that supplier j is pivotal during half-hour h of day d of month m if its residual demand is positive for all finite prices. Define the indicator variable $\text{Piv}_{j\text{hdm}}$ to equal 1 if supplier j is pivotal during half-hour h of day d of month m and zero otherwise. A related measure of the ability of supplier j to exercise unilateral market power is the pivotal quantity of energy for supplier j , which is the maximum of zero and the residual demand of supplier j evaluated at the highest observed offer during that half-hour period, p_{max} . If $\text{DR}_{j\text{hdm}}(p_{\text{max}})$ is the value of the residual demand curve at p_{max} for supplier j during half-hour h of day d of month m , then the value of the pivotal quantity $\text{PQuant}_{j\text{hdm}}$ equals $\max(0, \text{DR}_{j\text{hdm}}(p_{\text{max}}))$. Note that when supplier j is not pivotal the value of the pivotal quantity is zero and when the supplier is pivotal the value of $\text{PQuant}_{j\text{hdm}}$ equals $\text{DR}_{j\text{hdm}}(p_{\text{max}})$.

The analogous measure of the incentive of a supplier to exercise unilateral market power is the indicator variable for whether a supplier is net pivotal meaning that the pivotal quantity for that supplier exceeds its fixed-price forward market obligation. If $\text{QC}_{j\text{hdm}}$ is supplier j 's fixed-price forward market obligation in half-hour period h of day d and month m , then if $\text{DR}_{j\text{hdm}}(p_{\text{max}})$ is greater than $\text{QC}_{j\text{hdm}}$, the supplier is deemed to be net pivotal. Define the indicator variable $\text{NPiv}_{j\text{hdm}}$ to equal 1 if supplier j is net pivotal during half-hour h of day d of month m and zero otherwise. The second measure of the incentive to exercise unilateral market power is net pivotal quantity, which is defined as maximum of zero and the difference between the pivotal quantity and the supplier's fixed-price forward market obligation. Define $\text{NPQuant}_{j\text{hdm}}$, the net pivotal quantity for supplier j during half-hour h of day d of month m as $\max(0, \text{DR}_{j\text{hdm}}(p_{\text{max}}) - \text{QC}_{j\text{hdm}})$. If supplier j is not net pivotal then the value of $\text{NPQuant}_{j\text{hdm}}$ is equal to zero and if the supplier is net pivotal then $\text{NPQuant}_{j\text{hdm}} = \text{DR}_{j\text{hdm}}(p_{\text{max}}) - \text{QC}_{j\text{hdm}}$.

Table 6.1 presents summary statistics on pivotal indicator and net pivotal indicator variables for each year of our sample period from January 1, 2001 to June 30, 2007. Firm A has by far the highest pivotal and net pivotal frequency. For all but 2001, it is pivotal in more than 50 percent of the half-hour periods of the year. Next is Firm C with annual pivotal frequencies that range from 10 to 20 percent. Firm B's annual pivotal frequency ranges from slightly more than 3 percent to slightly more than 10 percent. Firm D has the lowest annual pivotal frequency of the four suppliers. It is important to note that one supplier being pivotal or net pivotal during a half-hour period does not preclude other suppliers from being pivotal or net pivotal during this

same half-hour period. Typically, when one large supplier is pivotal or net pivotal, other suppliers are as well.

Table 6.1 shows that for all suppliers but Firm A, being net pivotal is an extremely rare event. For all but 2001 for Firm A, the net pivotal percentage never exceeds one percent. For most of the years of the sample, the remaining suppliers are never net pivotal during any half-hour of the year. Firm B and Firm C are net pivotal only in 2001 and Firm C's net pivotal frequency is less than one-tenth that of Firm B. Therefore, we would not recommend putting much weight on the net pivotal regression results for Firm C because it is net pivotal for such a small number of half-hours during the sample period.

Table 6.2 presents linear regressions of the offer price at the supplier's dispatched quantity of energy on these two indicators of the ability of the supplier to exercise unilateral market power and the two indicators of the incentive of the supplier to exercise unilateral market power. All of these regressions include day-of-sample and half-hour-of-the-day fixed-effects similar to the regressions presented in Table 5.1. Table 6.3 presents linear regressions of the half-hourly offer price on half-hourly values of these same four variables with half-hour-of-the-day fixed-effects for each month of the sample similar to the regressions presented in Table 5.2. For all suppliers and all measures (except for the net pivotal dummy and net pivotal quantity for Firm C in Table 6.2), we find that a higher ability and incentive to exercise unilateral market power as measured by respectively, the pivotal indicator variable and pivotal quantity and net pivotal indicator and net pivotal quantity, predict higher offer prices for the supplier's dispatched quantity of energy.

Although the point estimates for Firm C in Table 6.2 for the net pivotal dummy and net pivotal quantity coefficients are negative, they are not statistically different from zero, which is to be expected given the near-zero frequency that Firm C is net pivotal during our sample period. Although the net pivotal dummy and net pivotal quantity parameter estimates for Firm C in Table 6.3 are both positive, they are not statistically different from zero, which provides further evidence that the near-zero frequency that Firm C is net pivotal during our sample period makes it impossible to estimate these coefficients with any degree of precision.

The estimates in Table 6.2 imply that keeping water levels and input fossil fuel prices constant, if Firm A is pivotal then the offer price for its dispatched quantity is expected to be \$12.45 higher. For Firm C this corresponding figure is \$11.40. According to Table 6.2, being

pivotal is predicted to increase a supplier's offer price by at least \$10.22, the coefficient estimate for Firm B. For all suppliers the coefficient on the pivotal quantity is also positive and precisely estimated. For example, a 50 MW pivotal quantity for Firm A implies a roughly \$1.00 higher offer price. For Firm C this same pivotal quantity of energy implies a \$1.65 higher offer price.

For the net pivotal indicator, the predicted offer price increases are much larger for two of the four firms and the predicted increase in the offer price for a net pivotal quantity change is an order of magnitude larger for two of these firms. For Firm B, being net pivotal predicts a \$67.58 increase in its offer price and for Firm D being net pivotal implies a \$220.40 increase in its offer price. For Firm A, a 20 MW net pivotal quantity predicts a \$1.56 higher offer price and for Firm B a 20 MW net pivotal quantity predicts a \$10.22 higher offer price. These net pivotal results should be interpreted with caution for all suppliers but Firm A because of the very small number of net pivotal events during the sample period for the remaining suppliers. Despite the very infrequent occurrence of being net pivotal for Firm B and Firm D, different from Firm C, these regressions yield precise estimates that imply these suppliers will adjust their price offers upward by economically meaningful magnitudes when they are net pivotal and increasingly so the larger is their net pivotal quantity.

These regression results and the results presented in Table 5.1 and 5.2 provide strong evidence that the higher market prices that occur when the four large suppliers have a greater unilateral ability and incentive to exercise market power, as shown in Figures 4.5 and 4.9, is due to the fact that these suppliers submit higher offer prices in order to raise market prices. In addition, when these suppliers have a substantial ability to exercise market power and have an incentive to exercise market power by lowering their offer price, they also do so. Taken together, the empirical evidence presented in this section suggests a causal link between the unilateral ability and incentive of suppliers to exercise market power and the offer prices they submit for the quantity of energy they sell in the short-term market. These higher or lower offer prices produce higher or lower market-clearing prices that are consistent with the unilateral ability and incentive of suppliers to exercise market power.

7. Do Thermal Suppliers Behave as if They Have No Ability to Exercise Market Power?

The final piece of evidence in favor of the view that the four large suppliers exercise all available unilateral market power is a test of the null hypothesis that suppliers behave as if they had no ability or incentive to exercise unilateral market power. As discussed in Section 3, a

supplier that has no ability or incentive to exercise unilateral market power can be expected to submit an offer curve equal to its aggregate marginal cost curve of supplying electricity. The complication with implementing this test for hydroelectric suppliers is that estimating their no-market-power opportunity cost of supplying energy is a massively complex computational problem. However, for fossil fuel suppliers we know that the opportunity cost of producing electricity from their generation units depends on the price of the input fossil fuel, the heat rate of the generation unit and the variable operating and maintenance cost of the generation unit. Consequently, as demonstrated in Section 3, a fossil fuel supplier with no ability to exercise unilateral market power will submit an offer price for each fossil generation unit equal to the unit's variable cost.

Our test of the null hypothesis that no supplier has the ability or incentive to exercise unilateral market power is based on the simple insight that offer prices of fossil fuel generation unit owners with no ability to exercise unilateral market power should not be predicted by any other factors besides those that impact the variable cost of the generation unit. In particular, if fossil fuel suppliers do not have any ability to exercise unilateral market power, the offer price for the amount of energy they sell into the short-term market should not be impacted by the system hydro storage level. In contrast, if higher offer prices are associated with lower water levels, then this is consistent with a supplier that has the ability to exercise unilateral market power taking advantage of this fact to raise their offer prices and market-clearing prices in response to the incentives that it faces.

To investigate this null hypothesis we regress the offer price for the quantity of energy sold from each fossil fuel generation unit during the half-hour periods of sample when the unit was available to supply energy on a number of factors that control for the variable cost of producing electricity from this generation unit at different levels of output and daily level hydro storage in Terawatt-hours (TWh). Let $P_{khd m}(\text{offer})$ equal the offer price of the energy sold in the short-term market from fossil fuel generation unit k during half-hour h of day d and month m . Let Hydro_{dm} equal the amount of hydroelectric energy in storage on day d of month m . Let $\text{QINC}_{ikd h m}$ equal a set of $I(k)$ dummy variables each of which equals 1 if the dispatch quantity from fossil fuel generation unit k during half-hour h of day d in month m lies in the 10 MW quantity increment i . For each generation unit we take the maximum and minimum output observed during the sample period and divide this range into 10 MW increments. For example,

if 250 MW is the lowest output level and 360 MW is highest output level, then $I(k)$ equals 11, meaning that there are 11 possible 10 MW bins that the supplier could produce in during the sample period. These quantity bins are chosen to account for the fact that the heat rate of fossil fuel units can be different for different output levels. Define YR_{zdhm} as an indicator variable that equals one if half-hour h of day d and month of sample m is in year z , where $z=2001,2002,\dots,2007$. Define MTH_{wdhm} as an indicator variable that equals 1 if half-hour h of day d and month-of-sample m is in month-of-the-year $w=1,2,3,\dots,12$. We estimate the following regression for each fossil fuel unit:

$$P_{kdhm}(offer) = \sum_{i=1}^{I(k)} \alpha_{ik} QINC_{ikdhm} + \sum_{z=2001}^{2007} \gamma_{zk} YR_{zkdhm} \quad (7.1)$$

$$+ \sum_{i=1}^{I(k)} \sum_{z=2002}^{2007} \theta_{izk} YR_{zkdhm} * QINC_{ikdhm} + \sum_{i=1}^{12} \delta_{ik} MTH_{wkdhm} + \beta_k Hydro_{dm} + \varepsilon_{kdhm}.$$

This linear regression controls for differences in the variable cost of fossil fuel units across the 10 MW quantity increments of output levels for the unit (the first summation), across each year of the sample (the second and third summations), and within the months of the year (the fourth summation) in order to assess whether the level of hydroelectric storage provides incremental explanatory power, beyond these variables that control for differences in the generation unit's variable cost of production, in predicting the offer price.

Table 7.1 presents the results of estimating (7.1) for the major fossil fuel units (or, in one case, group of units) operating in the New Zealand market during our sample period. In all cases, the estimated value of β_k , the coefficient associated with the value of system hydro storage for unit k , is found to be negative and precisely estimated. The null hypothesis that β_k is equal to zero is overwhelmingly rejected for all eight units, which provides strong evidence against the null hypothesis that the owners of these fossil fuel units behave as if they had no ability to exercise unilateral market power. The implied change in offer behavior from these generation units as a result of changes in the water level are also economically meaningful. For example, if the value of system hydro storage decreases by 1 TWh, then the offer price for the Plant 6 is predicted to increase by \$24.31 and by \$24.12 for the Plant 8. The predicted increases in the offer prices for a 1 TWh reduction in the value of system hydro storage for Plant 5 and Plant 7 are roughly half these values. Plant 1 and Plant 3 have predicted offer price increases for a 1 TWh reduction in system hydro storage of \$17.40 and \$19.61, respectively. Note that the difference between the minimum and maximum system hydro storage levels during our sample

period is 3.1 TWh, so these estimates predict very large changes in the offer prices of fossil fuel units for the observed changes in hydrological conditions.

Although these parameter estimates are inconsistent with the hypothesis that these fossil fuel generation unit owners have no ability to exercise unilateral market power, the signs and magnitudes of the estimated values of the β_k are consistent with the hypothesis that the owners of these generation units have a significant ability to exercise unilateral market power and that this ability to exercise unilateral market power increases with the level of system hydro storage. These results are also consistent with the results presented in the previous section which showed that the offer price for the quantity of energy sold in the short-term market by each of the four suppliers is increasing in that supplier's ability and incentive to exercise unilateral market power.

8. Conclusions about How Firms Exercise Market Power

The three lines of empirical inquiry presented in this paper are broadly consistent with the implications of expected profit-maximizing offer behavior by the four large suppliers in response to the extent of competition they face from other suppliers on a half-hourly basis. This conclusion does not depend on any assumptions about the functional form of aggregate demand in the market or any model of strategic interaction among firms. Because of the data-rich multi-unit auction environment that we study, ex post half-hourly measures of the ability of a supplier to exercise market power using the offers submitted by all suppliers and the level of system demand can be computed without either of these assumptions. We find that each of the four large suppliers submits a higher half-hourly offer price when it has a higher half-hourly unilateral ability to exercise market power. The half-hourly offer price increases predicted by the parameter estimates from our econometric model for typical changes in the half-hourly ability of each supplier to exercise market power are economically significant in the sense that the implied offer price increases can be in the range of \$10/MWh to \$20/MWh during peak periods of the day.

We find even larger (in absolute value) predicted changes in a supplier's half-hourly offer prices in response to changes in its half-hourly incentive to exercise market power for typical changes in the values of these indexes. Our index of the half-hourly incentive of a supplier to exercise market power can be positive or negative, depending on the supplier's exposure to short-term market-clearing price during that half-hour period. If a supplier is net long—its short-term market sales exceed its fixed-price forward market obligations for that half-hour—then its

index of the incentive to exercise market power is positive. If a supplier is short—its sales are less than its fixed-price forward market obligations for that half-hour—then its index of the incentive to exercise market power is negative. Our regression results predict that sizeable increases in the supplier's offer price occur during half-hour periods when this index of the supplier's incentive to exercise market power is large and positive and sizeable decreases in the supplier's offer prices occur during the half-hour periods when this half-hourly index of the supplier's the incentive to exercise market power is large in absolute value and negative. These results emphasize that the extent a supplier actually exploits a lower degree of competition from other firms depends on the incentive it has to do so, as measured by the degree to which the revenues the supplier receives depends on the short-term market-clearing price. In addition, how the supplier exploits its ability to influence the short-term market price depends on the sign of its exposure to short-market prices. This result implies that a portion of the high degree of volatility in half-hourly short-term wholesale electricity prices is the result of changes in the sign of the half-hourly incentive of suppliers to exercise unilateral market power.

These relationships between the half-hourly ability and incentive of suppliers to exercise market power and the offer price than they submit also hold for half-hourly indexes of the ability and incentive to exercise market power based on the pivotal and net supplier concept. Sizeable offer price increases are predicted for each of the suppliers during the half-hour periods when they are pivotal. Finally, we provide strong evidence against the null hypothesis that the half-hourly offer curves submitted by owners of fossil fuel generation units are the result of those suppliers behaving as if they have no ability to exercise market power.

Taken together, the empirical results in this paper demonstrate that although prices in a multi-unit auction wholesale electricity market depend on supply and demand conditions, actual supply conditions depend on the offer curves submitted by market participants to the wholesale market. These offer curves are direct result of the unilateral expected profit-maximizing actions of suppliers given factors that they are unable to control such as the level of demand at all locations in the New Zealand, amount of water inflows to hydroelectric generation units and the price of fossil fuels and other inputs consumed to produce electricity. Therefore, the ability and incentive of large suppliers to exercise unilateral market power are important determinants of the supply conditions that determine short-term wholesale prices, even after the impact of exogenous factors such as water availability and fossil fuel prices have been taken into account.

References

- Berry, Steven, T. Levinsohn, James, L. and Pakes, Ariel (1995), "Automobile Prices in Market Equilibrium," *Econometrica*, 63(4), 841-890.
- Bohn, Roger, E., Caramanis, Michael C., and Schweppe, Fred, C. (1984) "Optimal Pricing in Electrical Networks over Space and Time," *The RAND Journal of Economics*, Volume 15(3), 360-376.
- Bresnahan, Timothy (1981), "Departures from Marginal-Cost Pricing in the American Automobile Industry," *Journal of Econometrics*, 17, 201-227.
- Bresnahan, Timothy (1987), "Competition and Collusion in the American Automobile Market: The 1955 Price War," *Journal of Industrial Economics*, 45(4, June) (special issue), 457-482.
- Bresnahan, Timothy (1989). "Empirical methods for industries with market power". In: Schmalensee, R., Willig, R. (Eds.), *Handbook of Industrial Organization*, vol. 2. North-Holland, Amsterdam.
- Bulow, Jeremy, and Pfleiderer, Paul (1983), "A Note on the Effect of Cost Changes on Prices," *Journal of Political Economy*, 1983, 91(1), 182-185.
- Goldberg, Penny K., (1995), "Product Differentiation & Oligopoly in International Markets: The Case of the U.S. Automobile Industry," *Econometrica*, 63(4), 891-952.
- Nevo, Aviv (2001), "Measuring Market Power in the Ready-to-Eat Cereal Industry," *Econometrica*, 69(2), 307-342.
- Porter, Robert H., (1983), "A Study of Cartel Stability: The Joint Executive Committee, 1880-1886," *Bell Journal of Economics*, 14(2, Autumn), 301-314.
- Reiss, Peter, C. and Wolak, Frank A. (2007), "Structural Econometric Modeling: Rationales and Examples from Industrial Organization," *Handbook of Econometrics*, Volume 6A, edited by James J. Heckman and Edward E. Leamer, 4277-4415.
- Rosse, James, N. (1970), "Estimating Cost Function Parameters Without Using Cost Data: Illustrated Methodology," *Econometrica*, 38(2), 256-275.
- Wolak, Frank A. (2000), "An Empirical Analysis of the Impact of Hedge Contracts on Bidding Behavior in a Competitive Electricity Market," *International Economic Journal*, 14(2), 1-40.
- Wolak, Frank A. (2003a), "Identification and Estimation of Cost Functions Using Observed Bid Data: An Application to Electricity Markets," in M. Dewatripont, L.P. Hansen, and S.J.

Turnovsky, eds., *Advances in Economics and Econometrics: Theory and Applications, Eight World Congress, Volume II*. New York: Cambridge University Press, pp. 133-169.

Wolak, Frank A. (2003b) “Measuring Unilateral Market Power in Wholesale Electricity Markets: The California Market 1998 to 2000,” *American Economic Review*, May 2003, 425-430.

Wolak, Frank A. (2007), “Quantifying the Supply-Side Benefits from Forward Contracting in Wholesale Electricity Markets,” *Journal of Applied Econometrics*, 22, 1179-1209.

Figure 3.1: Offer curve for Firm A for a peak half-hour period in February 2006

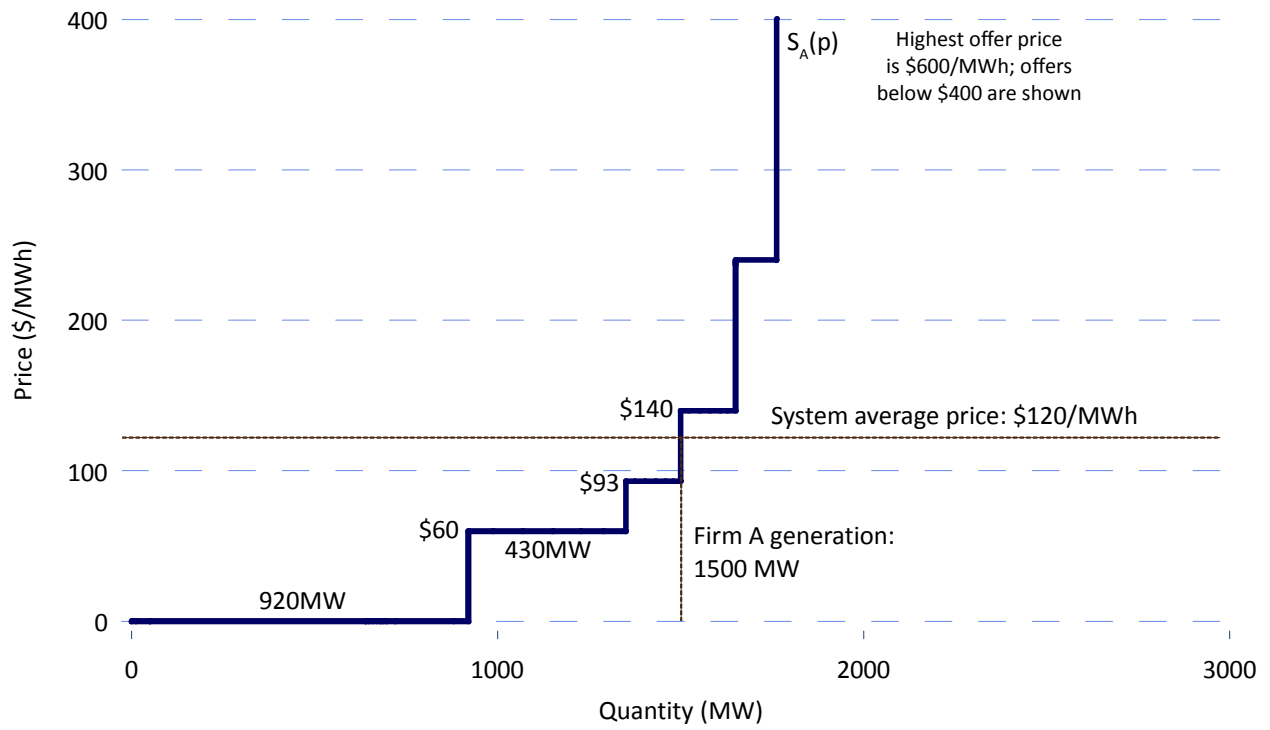


Figure 3.2: Offer curve for Firm B for the same peak half-hour period in February 2006




Figure 3.3: Combined offer curve for Firms A and B for the half-hour period




Figure 3.4: Aggregate offer curve for all generators for the half-hour period

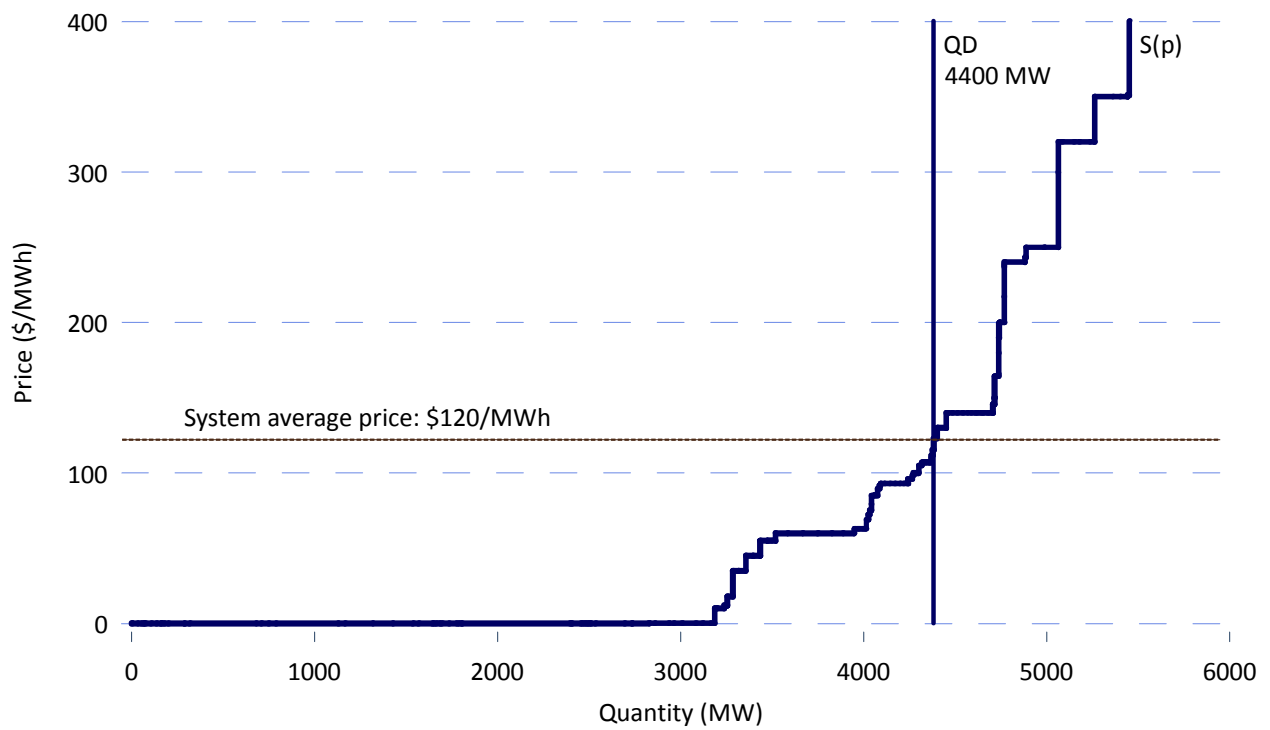


Figure 3.5: Calculation of residual demand for Firm A in the half-hour period

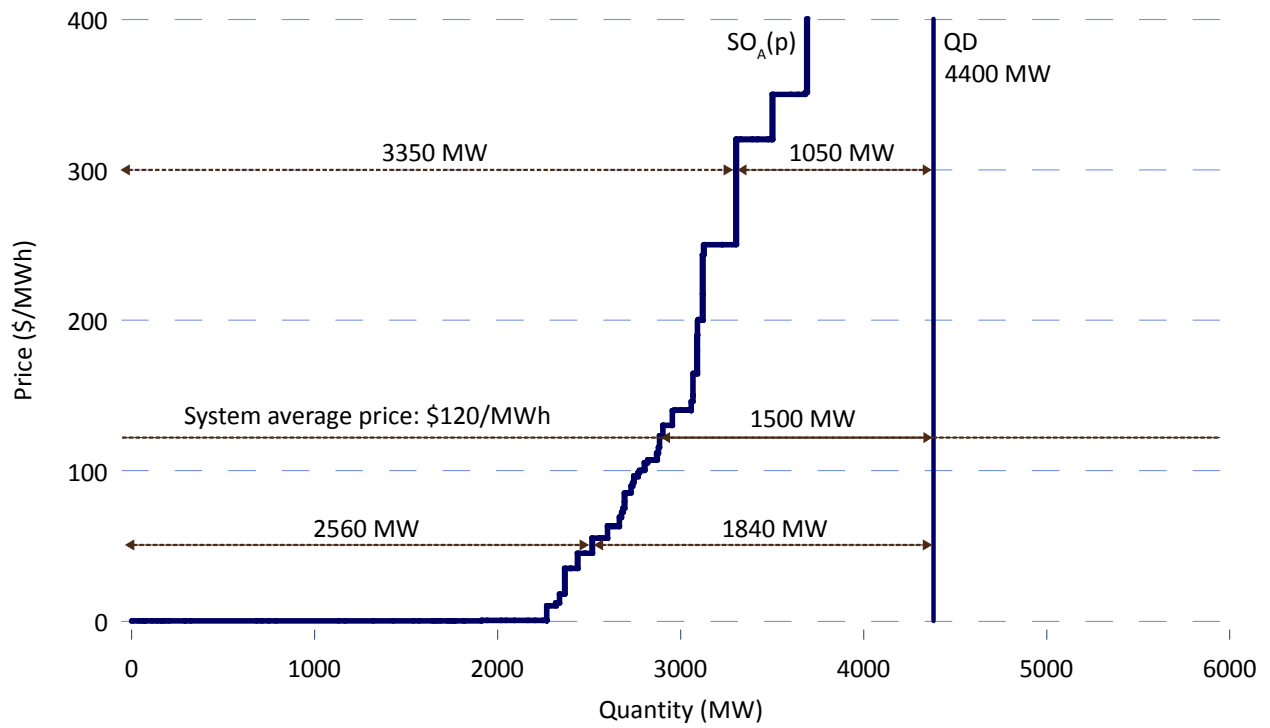


Figure 3.6: Residual demand for Firm A for the half-hour period

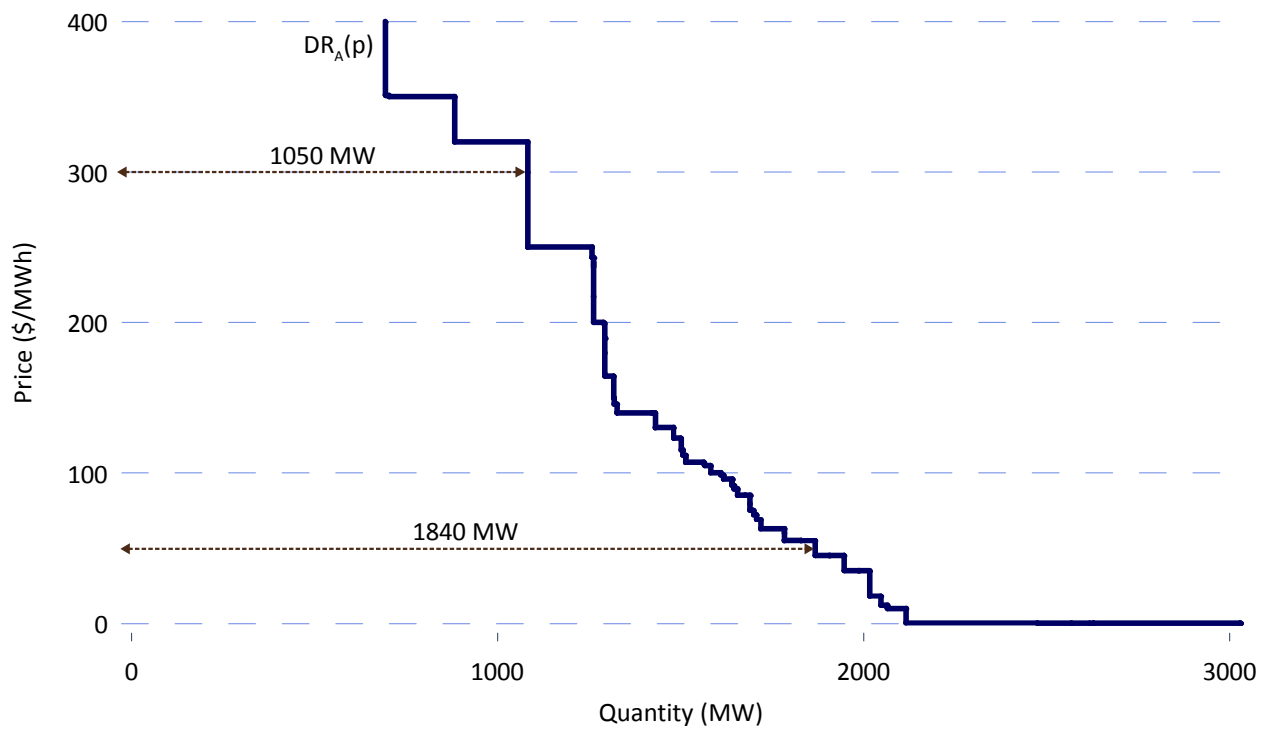


Figure 3.7: Residual demand and offer curve for Firm A for the half-hour period

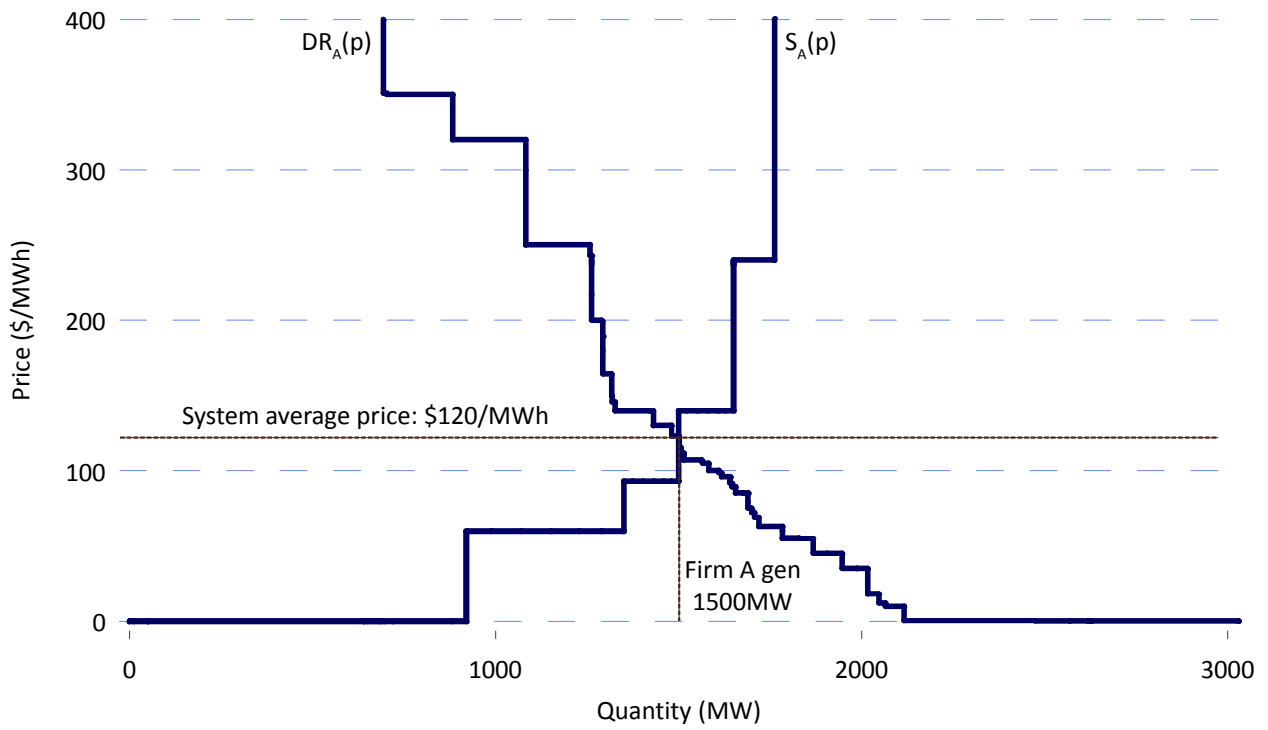


Figure 3.8: Example showing residual demand and the calculation of inverse elasticity

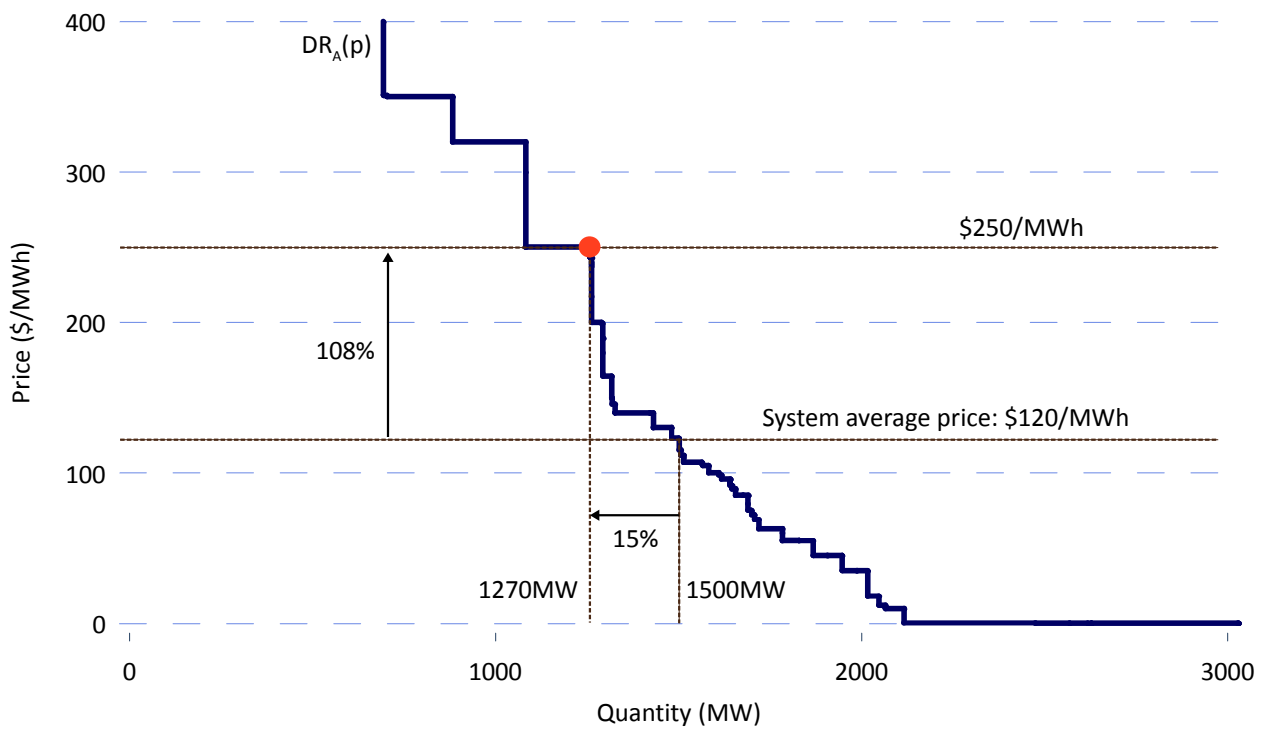


Figure 3.9: Profit-maximizing choice of price and quantity




Figure 3.10: Profit-maximizing price and quantity with elastic residual demand




Figure 3.11: Profit-maximizing price and quantity with perfectly elastic residual demand




Figure 3.12: Derivation of offer curve (steep residual demands)




Figure 3.13: Derivation of offer curve (flatter residual demands)




Figure 3.14: Derivation of offer curve (perfectly elastic residual demands)




Figure 3.15: Impact of Step Functions on Optimal Offer Curve




Figure 3.16: Expected Profit-Maximizing Offer Curve




Figure 3.17: Effect of fixed-price obligations on Firm C for a half-hour period in February 2006

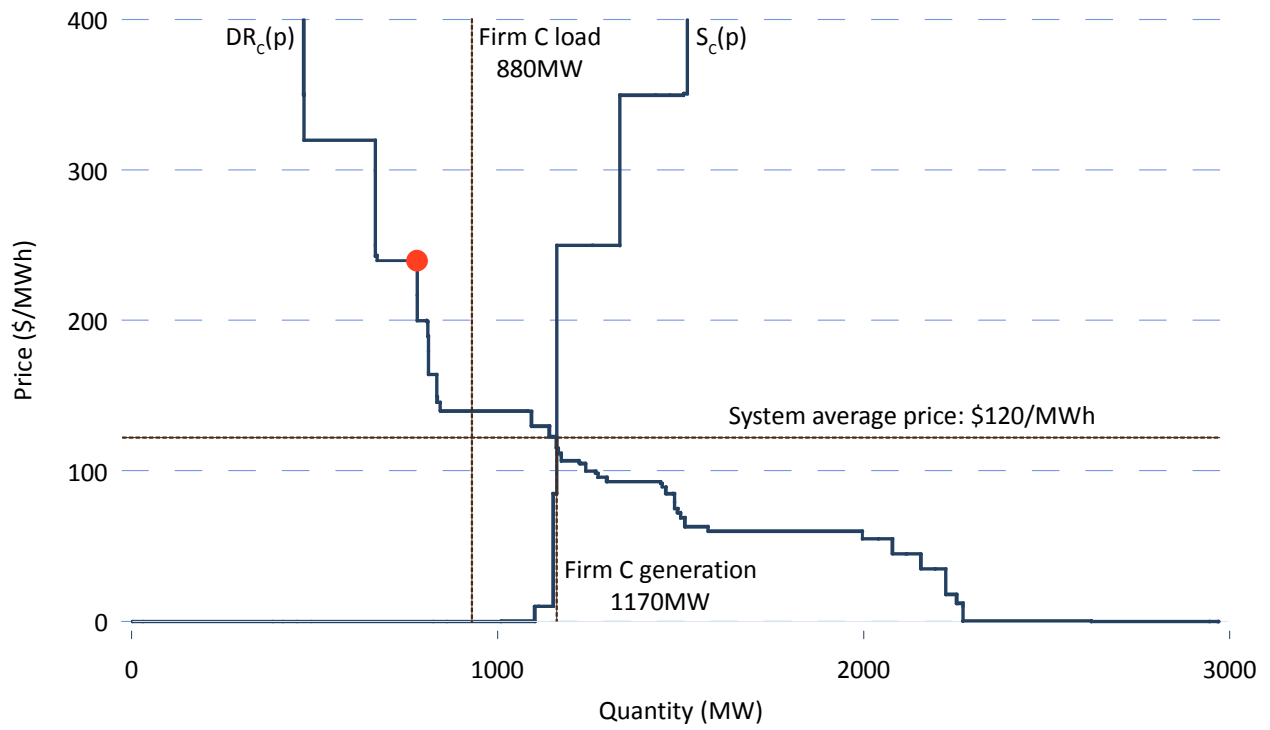


Figure 3.18: Profit-maximization with fixed-price contracts, part I




Figure 3.19: Profit-maximization with fixed-price contracts, part II




Figure 3.20: Price and quantity with and without fixed-price contracts

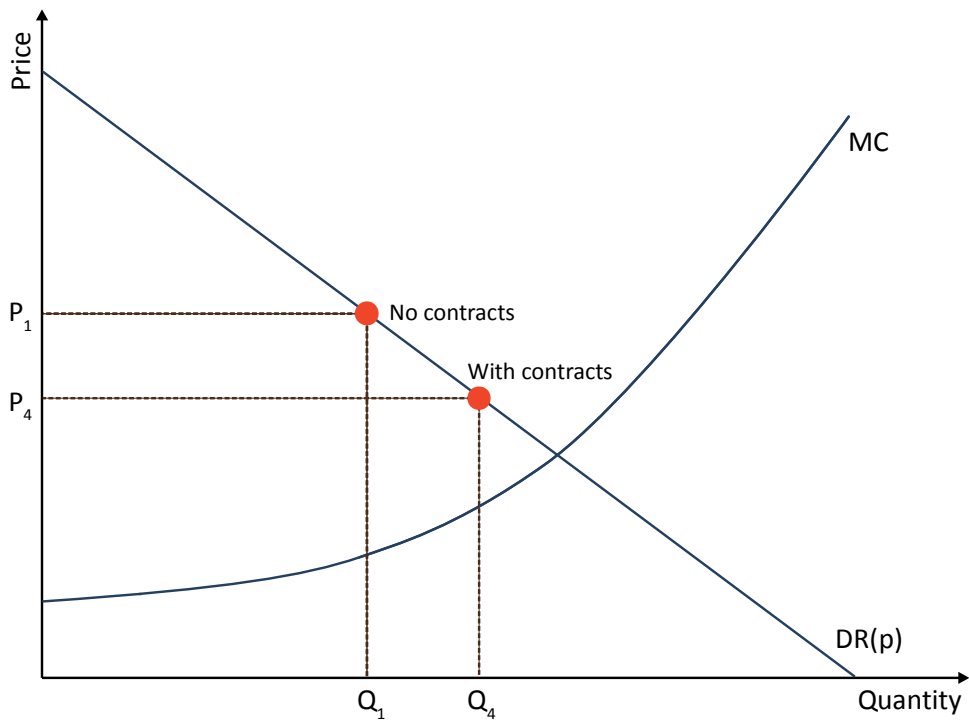


Figure 3.21: Derivation of offer curves with and without fixed-price contracts




Figure 3.22: Definition of Pivotal Supplier and Pivotal Quantity




Figure 3.23: Offer curve determination with pivotal residual demand – Step 1




Figure 3.24: Offer curve determination with pivotal residual demand – Step 2




Figure 3.25: Offer curve determination with pivotal residual demand – Step 3




Figure 4.1: Elasticity calculation for Firm B, peak half-hour period in February 2006




Figure 4.2: Elasticity calculation for Firm B, peak half-hour period exactly 1 year later

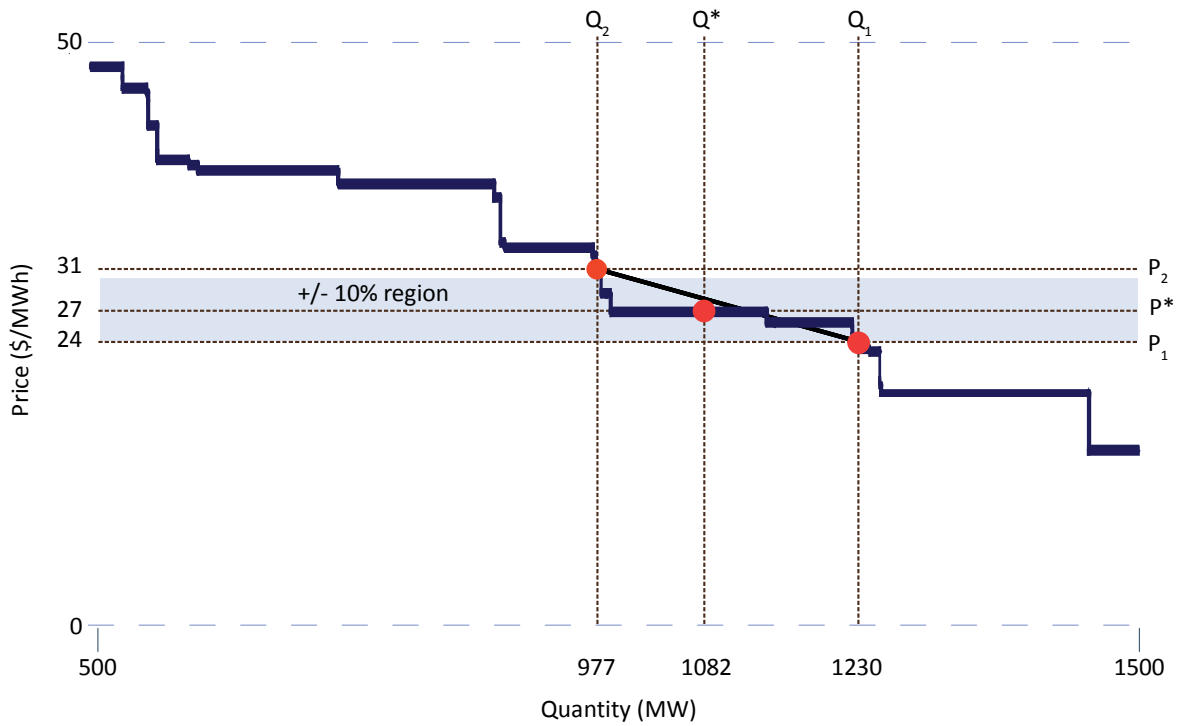


Figure 4.3: Half-hourly inverse semi-elasticities by firm, 30-day rolling average




Figure 4.4: Half-hourly mean inverse semi-elasticities by firm, 2001–07




Figure 4.5: Mean inverse semi-elasticities and system price, 30-day rolling average




Figure 4.6: Half-hourly mean inverse semi-elasticities and system price, 2001–07

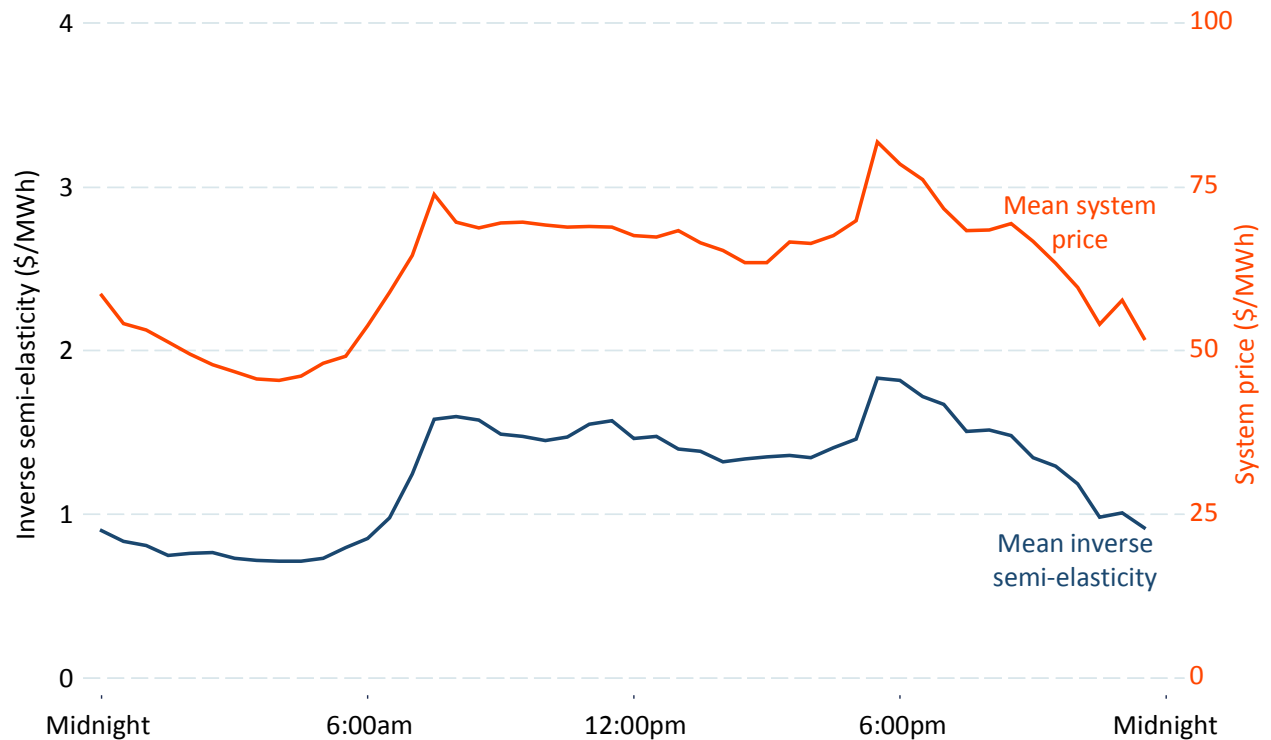


Figure 4.7: Half-hourly net inverse semi-elasticities by firm, 30-day rolling average




Figure 4.8: Mean net inverse semi-elasticities and system price, 30-day rolling average




Figure 4.9: Example showing calculation of offer prices




Table 4.1: Correlation between semi-elasticity results for different price windows

Price window	Firm A				Firm B			
	15%	10%	5%	1%	15%	10%	5%	1%
15%	1.00				1.00			
10%	0.96	1.00			0.96	1.00		
5%	0.88	0.92	1.00		0.92	0.96	1.00	
1%	0.81	0.85	0.93	1.00	0.82	0.85	0.90	1.00
Mean semi-elasticity	2.05	2.07	2.07	1.99	0.88	0.88	0.86	0.83

Price window	Firm C				Firm D			
	15%	10%	5%	1%	15%	10%	5%	1%
15%	1.00				1.00			
10%	0.95	1.00			0.97	1.00		
5%	0.90	0.95	1.00		0.91	0.94	1.00	
1%	0.84	0.90	0.94	1.00	0.85	0.88	0.94	1.00
Mean semi-elasticity	1.27	1.28	1.25	1.21	0.73	0.74	0.75	0.74

Table 5.1: Dependent variable = offer price at dispatch quantity for supplier j

	Firm A	Firm B	Firm C	Firm D
β_j (s.e.)	0.46 (.017)	0.56 (.040)	1.41 (.031)	3.81 (.062)
δ_j (s.e.)	5.08 (.108)	4.02 (.146)	4.31 (.101)	21.63 (.335)

Note: Day-of-sample and half-hour fixed effects are included in all regressions.

Table 5.2: Dependent variable = offer price at dispatch quantity for supplier j

	Firm A	Firm B	Firm C	Firm D
β_j (s.e.)	0.67 (.020)	0.73 (.040)	1.16 (.029)	4.54 (.064)
δ_j (s.e.)	7.27 (.129)	3.39 (.154)	3.38 (.092)	22.86 (.354)

Note: Month-of-sample interacted with half-hour fixed effects are included in all regressions.

Table 5.3: Half-hourly summary statistics for η_i by firm

Half-hour	Firm A		Firm B		Firm C		Firm D	
	Mean	Std. dev.	Mean	Std. dev.	Mean	Std. dev.	Mean	Std. dev.
1	1.541	2.636	0.644	2.079	0.993	2.420	0.416	0.913
2	1.429	2.530	0.587	1.276	0.975	2.807	0.339	0.640
3	1.447	2.410	0.584	1.416	0.888	2.035	0.310	0.608
4	1.351	2.166	0.533	0.921	0.822	1.483	0.287	0.585
5	1.412	2.474	0.526	1.054	0.813	1.511	0.289	0.659
6	1.409	2.594	0.505	0.962	0.860	1.991	0.281	0.613
7	1.333	2.322	0.509	1.620	0.800	2.455	0.279	0.643
8	1.333	2.318	0.472	0.799	0.789	1.468	0.278	0.621
9	1.333	2.491	0.486	0.965	0.748	1.218	0.275	0.588
10	1.348	2.354	0.456	0.687	0.761	1.422	0.269	0.546
11	1.342	2.285	0.501	1.107	0.780	1.634	0.288	0.642
12	1.393	2.723	0.587	1.993	0.882	2.444	0.311	0.698
13	1.445	2.627	0.654	1.669	0.939	2.496	0.364	0.699
14	1.584	3.191	0.688	1.210	1.137	3.897	0.485	1.028
15	1.917	4.402	0.882	2.197	1.409	5.288	0.760	2.667
16	2.418	6.425	1.229	6.171	1.589	4.039	1.045	3.220
17	2.570	7.223	1.194	5.489	1.537	3.994	1.057	4.789
18	2.463	6.399	1.134	3.108	1.622	4.799	1.066	4.715
19	2.372	5.881	1.120	4.086	1.501	4.333	0.968	2.636
20	2.300	5.690	1.161	4.631	1.526	4.469	0.926	2.268
21	2.364	6.734	1.068	3.243	1.442	3.706	0.958	2.774
22	2.479	6.608	0.997	2.816	1.487	4.349	0.942	2.420
23	2.677	9.769	1.004	2.491	1.549	4.374	0.988	2.709
24	2.668	9.224	1.008	2.480	1.647	5.656	0.970	2.493
25	2.366	6.058	0.999	3.021	1.562	5.810	0.924	2.770
26	2.458	6.747	1.043	4.225	1.486	4.647	0.920	2.844
27	2.348	5.341	0.962	3.524	1.408	3.411	0.872	2.338
28	2.319	6.026	0.967	3.707	1.402	3.847	0.851	2.302
29	2.198	5.043	0.890	2.019	1.322	3.034	0.852	2.471
30	2.247	6.291	0.965	3.881	1.305	3.477	0.834	2.532
31	2.293	6.303	0.933	3.261	1.366	4.191	0.817	2.341
32	2.254	5.510	0.951	3.088	1.394	3.937	0.839	2.179
33	2.263	4.978	0.877	1.713	1.402	3.845	0.850	2.437
34	2.318	5.427	0.974	3.242	1.437	4.428	0.896	2.420
35	2.375	4.528	1.057	2.809	1.445	3.843	0.954	2.619
36	2.853	6.571	1.364	4.375	1.823	4.874	1.257	4.240
37	2.712	5.981	1.301	4.435	1.989	6.811	1.241	3.660
38	2.672	5.361	1.191	2.690	1.784	5.458	1.186	3.423
39	2.599	6.263	1.203	3.762	1.687	5.058	1.168	3.888
40	2.454	6.112	1.079	2.861	1.435	3.618	1.023	3.723
41	2.448	6.388	1.082	3.108	1.452	4.059	1.042	4.053
42	2.402	5.690	1.060	2.568	1.482	4.036	0.954	2.524
43	2.242	4.855	0.954	2.347	1.298	3.253	0.884	2.194
44	2.218	5.787	0.894	3.299	1.298	4.765	0.768	2.229
45	2.093	4.885	0.831	2.304	1.148	2.636	0.656	1.630
46	1.747	2.897	0.694	1.501	0.977	1.843	0.511	0.988
47	1.758	3.442	0.715	2.311	0.990	1.962	0.562	1.765
48	1.705	4.103	0.613	1.365	0.917	2.047	0.419	1.073

Table 5.4: Half-hourly summary statistics for η_i^c by firm

Half-hour	Firm A		Firm B		Firm C		Firm D	
	Mean	Std. dev.	Mean	Std. dev.	Mean	Std. dev.	Mean	Std. dev.
1	0.029	0.356	0.143	0.632	0.346	1.020	-0.029	0.290
2	-0.010	0.415	0.142	0.518	0.354	1.086	-0.033	0.242
3	-0.033	0.348	0.145	0.586	0.333	0.822	-0.028	0.222
4	-0.055	0.345	0.132	0.459	0.315	0.676	-0.022	0.189
5	-0.061	0.410	0.130	0.577	0.314	0.667	-0.022	0.194
6	-0.062	0.431	0.125	0.508	0.346	1.004	-0.028	0.189
7	-0.077	0.401	0.126	0.981	0.311	1.020	-0.023	0.206
8	-0.084	0.406	0.099	0.266	0.308	0.638	-0.028	0.209
9	-0.087	0.414	0.099	0.308	0.290	0.519	-0.026	0.163
10	-0.083	0.401	0.093	0.274	0.295	0.604	-0.024	0.173
11	-0.066	0.359	0.101	0.569	0.298	0.706	-0.019	0.177
12	-0.044	0.364	0.111	0.761	0.332	0.961	-0.030	0.191
13	0.010	0.379	0.131	0.856	0.333	0.962	-0.010	0.188
14	0.045	0.445	0.112	0.468	0.358	1.281	0.006	0.213
15	0.100	0.594	0.127	0.623	0.401	1.516	0.050	0.416
16	0.236	0.834	0.180	1.151	0.410	1.341	0.117	0.652
17	0.293	1.036	0.132	0.980	0.366	1.196	0.115	0.556
18	0.299	1.033	0.137	0.822	0.386	1.527	0.122	0.584
19	0.276	0.910	0.130	0.737	0.347	1.168	0.110	0.446
20	0.259	0.778	0.119	0.672	0.357	1.273	0.100	0.413
21	0.272	0.821	0.121	0.677	0.347	1.085	0.094	0.440
22	0.274	1.003	0.129	0.716	0.370	1.537	0.092	0.366
23	0.311	1.087	0.102	0.686	0.385	1.255	0.097	0.427
24	0.324	1.187	0.092	0.796	0.418	1.803	0.085	0.412
25	0.295	1.038	0.113	0.914	0.381	1.543	0.078	0.568
26	0.306	1.147	0.110	0.940	0.362	1.143	0.070	0.502
27	0.303	1.121	0.104	0.871	0.343	0.956	0.071	0.378
28	0.272	0.938	0.106	0.895	0.359	1.199	0.061	0.423
29	0.245	0.769	0.107	0.793	0.354	0.984	0.052	0.379
30	0.243	0.992	0.099	0.927	0.356	1.168	0.039	0.390
31	0.242	0.957	0.105	0.838	0.377	1.427	0.040	0.335
32	0.269	0.976	0.097	0.856	0.361	1.119	0.051	0.312
33	0.265	0.912	0.094	0.705	0.372	1.193	0.057	0.294
34	0.311	1.002	0.077	0.640	0.359	1.229	0.071	0.364
35	0.324	0.734	0.109	1.019	0.349	1.113	0.095	0.479
36	0.454	1.466	0.170	1.381	0.421	1.415	0.155	0.842
37	0.417	1.067	0.142	1.304	0.479	2.335	0.165	0.760
38	0.392	0.947	0.145	1.089	0.413	1.495	0.152	0.793
39	0.350	1.056	0.140	1.226	0.395	1.329	0.145	0.777
40	0.296	1.031	0.135	1.198	0.351	1.074	0.113	0.580
41	0.300	0.990	0.116	0.989	0.355	1.198	0.112	0.653
42	0.293	0.960	0.117	0.904	0.366	1.208	0.082	0.482
43	0.278	0.805	0.109	0.795	0.335	1.003	0.068	0.361
44	0.228	0.764	0.086	0.574	0.363	1.612	0.048	0.380
45	0.185	0.753	0.135	1.378	0.350	0.954	0.008	0.296
46	0.091	0.549	0.120	0.836	0.318	0.715	-0.017	0.265
47	0.115	0.496	0.125	1.001	0.291	0.716	0.008	0.448
48	0.068	0.574	0.127	0.686	0.304	0.875	-0.040	0.568

Table 6.1: Summary statistics for pivotal variables

	Firm A	Firm B	Firm C	Firm D
Gross pivotal				
2001	49.1%	4.7%	13.8%	2.7%
2002	61.2%	10.4%	12.9%	5.2%
2003	52.9%	4.6%	17.6%	1.4%
2004	60.9%	10.2%	20.2%	5.4%
2005	53.5%	3.4%	17.0%	1.2%
2006	52.4%	6.1%	16.6%	1.7%
2007	51.2%	4.1%	13.2%	0.6%
Net pivotal				
2001	2.25%	0.43%	0.02%	-
2002	0.67%	-	-	0.02%
2003	0.13%	-	-	0.02%
2004	0.48%	-	-	-
2005	0.07%	-	-	-
2006	0.06%	-	-	-
2007	-	-	-	-

Table 6.2: Dependent variable = offer price at dispatch quantity for supplier j

Regression on:	Firm A	Firm B	Firm C	Firm D
(a) Pivotal dummy (0/1) (s.e.)	12.45 (.267)	10.22 (.491)	11.40 (.331)	16.07 (.903)
(b) Pivotal quantity (MW) (s.e.)	0.020 (.0004)	0.034 (.0018)	0.033 (.0010)	0.086 (.0049)
(c) Net pivotal dummy (0/1) (s.e.)	10.62 (1.09)	67.58 (4.02)	-16.62 (15.92)	220.4 (16.03)
(d) Net pivotal quantity (MW) (s.e.)	0.078 (.007)	0.511 (.032)	-0.093 (.312)	1.642 (.133)

Note: Day-of-sample and half-hour fixed effects are included in all regressions.

Table 6.3: Dependent variable = offer price at dispatch quantity for supplier j

Regression on:	Firm A	Firm B	Firm C	Firm D
(a) Pivotal dummy (0/1) (s.e.)	17.37 (.310)	13.68 (.544)	14.33 (.322)	24.34 (.993)
(b) Pivotal quantity (MW) (s.e.)	0.030 (.0004)	0.048 (.0020)	0.047 (.0011)	0.135 (.0055)
(c) Net pivotal dummy (0/1) (s.e.)	19.37 (1.36)	70.35 (4.32)	5.76 (15.68)	259.7 (17.13)
(d) Net pivotal quantity (MW) (s.e.)	0.122 (.008)	0.585 (.035)	0.350 (.307)	1.937 (.143)

Note: Month-of-sample interacted with half-hour fixed effects are included in all regressions.

Table 7.1: Dependent variable = offer price at dispatch quantity for fossil fuel plant/unit k

	Plant 1	Plant 2	Plant 3	Plant 4
β_k	-17.40	-2.34	-19.61	-21.13
(s.e.)	(.457)	(.135)	(.340)	(.448)
	Plant 5	Plant 6	Plant 7	Plant 8
β_k	-8.05	-24.31	-11.01	-24.12
(s.e.)	(.674)	(.377)	(.459)	(.335)

Note: Regressions include year-of-sample fixed effects interacted with generation quantity in 10MW bins, as well as month-of-year fixed effects. The dependent variable in each regression is the offer price from either a single generation unit, or a group of units.