Status Quo Effects In Bargaining: An Empirical Analysis of OPEC

Kyle Hyndman*
New York University

25 November 2005

Abstract

We conduct an event analysis on OPEC quota announcements to determine their impact on the stock returns in the oil industry. We find that announcements to reduce the quota are followed by positive excess returns over pre-announcement levels, announcements of no action are met with negative excess returns and announcements to increase the quota have no significant impact on stock market returns. This suggests that there is an asymmetric ability on the part of OPEC to secure agreements. In particular, when demand has increased, agreements are easily forthcoming, while when times are bad the probability of a disagreement is substantially higher. We present further empirical as well as anecdotal evidence to support our interpretation. A bargaining model with one-sided private information which generates such predictions is also discussed. We also show that our model explains observed patterns in cheating by OPEC countries.

JEL Classification: D70, D82, L13, L71
Keywords: Status Quo, OPEC, Collusion, Disagreement, Bargaining

*I would like to thank my advisor Debraj Ray for suggesting the topic and for his advice and encouragement throughout the process. Conversations with Alberto Bisin, Alessandro Lizzeri, David Pearce, Ronny Razin, Giovanni Serio, Ennio Stacchetti and Henry van Egteren were also extremely beneficial. Correspondence may be directed to the author at Department of Economics, New York University, 7^{th} Floor, 269 Mercer Street, New York, NY 10003. Tel: (212) 998-3745. email: kbh212@nyu.edu. Financial support from NSF grant #0241070 (Debraj Ray) is also gratefully acknowledged.

1 Introduction

In this paper we are concerned with the following question: are agreements easier to achieve when times are good or when times are bad? This is a question which has been well-studied in many fields of economics. There is a large literature in industrial organisation asking whether collusion is more difficult to sustain during booms or recessions. Empirically, the results are mixed. The general view, espoused by Scherer and Ross (1990, Ch. 8), is that collusion is generally much more difficult to sustain during recessions. More recently, however, Reynolds and Wilson (2005) argue that markups are countercyclical but that prices are more variable during recessions than booms. From a theoretical perspective, the results are also mixed — the answer typically depends crucially on the assumptions about the demand process.[1]

This is also an important question in labour economics; here the concern is whether strikes are more frequent/persistent during recessions rather than during booms. In this literature, the empirical results are quite robust. Strikes are typically more frequent during booms than recessions but, conditional on a strike occurring, it is longer lasting in a recession.[2] While the empirical results are clear, the theoretical explanations seem mixed. Most models take a bargaining approach in which delay is used as a tool (either through signalling or screening) to convey private information. Often, to make the predictions of the model consistent with the empirical findings, one needs to make the assumption that commitment is more difficult to sustain during booms.

Our interest in studying the Organisation of Petroleum Exporting Countries (OPEC) puts us at the crossroads of these two literatures. OPEC is a cartel which restricts output to obtain excess profits and must clearly be concerned with incentives that individual countries may have to overproduce and the collective stick it wields to prevent cheating on quotas. However, given the many folk theorems present in the literature, there is not a unique way to split the gains from cartelisation of the oil market. Indeed, the problem of splitting these gains appears to involve a great deal of closed-door negotiations, which often bears a striking resemblance to a union and firm negotiating over a contract, with individual OPEC members posturing in order to extract the largest possible split. However, often this posturing leads to extended periods of inaction, leading to lower profits for each member. Without even mentioning politics, this is what makes examining OPEC both very interesting and challenging.

There have been many studies of OPEC behaviour, beginning with Griffin (1985) and continuing up to Ramcharran (2002) and beyond. The typical approach in these studies is to estimate a supply curve for each OPEC country and try to describe their behaviour according to one of four competing hypotheses — a competitive model, a cartel model, a target revenue model or a property rights model. These empirical approaches are useful, however, we believe that they are not particularly well-suited to studying the influence of the business cycle on OPEC behaviour.

[1]For example, Bagwell and Staiger (1997) present a very simple model in which collusion starts being more difficult to sustain during recessions and then switches to the opposite prediction as the correlation between demand states moves from one extreme to the other. A similar result appears in Athey, Bagwell and Sanchirico (2004).

[2]Kennan and WIlson (1993) provide a thorough survey of both the empirical and theoretical literature on strike incidence/duration.

Instead, to determine if a relationship exists, we turn to the methodology of event studies. Indeed, we find that there is a relationship between OPEC behaviour and the state of the economy, and this relationship takes us back to our opening discussion. In particular, we find evidence that disagreements are substantially more likely during bad times than during good times.

Our event study allows us to quantify how markets react to production quota announcements by OPEC, which, allows us to infer something about OPEC's decision-making capabilities. However, one will not find OPEC listed on any stock market. Therefore, in order to conduct such an analysis, our approach must necessarily be indirect. Clearly, OPEC adjusts its quota with the intent of influencing the price of oil to maximise its profits [or some other objective function]. Therefore, we may consider analysing how OPEC's various announcements affects the price of oil in the period around the announcement. Beyond this, we can exploit the close relationship between OPEC and various western oil companies. Through its influence on oil prices, OPEC indirectly influences the profitability of western oil companies. In general, for many oil companies, we would expect any action by OPEC which increases the price of oil to increase their profitability; similarly actions leading to lower oil prices should lower their profitability. Moreover, the incentives of many western oil companies and OPEC members may be in substantial alignment due to a relationship rather similar to sharecropping. Thus any action taken by OPEC which is profitable for itself, should also be profitable for the oil company in question.[3]

As is standard in event studies, we measure the impact of announcements by examining the pattern of abnormal returns in the days surrounding each event. More specifically, we divide OPEC announcements into three classes: (1) increases in the quota, (2) decreases in the quota and (3) no change in the quota (or status quo). We then examine how the price of oil and stock market indices of oil industry groups respond to each class of announcement. Upon presenting the results, we argue that they offer us interesting insights into the behaviour of OPEC and their decision-making apparatus — as the title of the paper suggests, we find evidence that the status quo is given extra significance in certain states of the world.[4]

After discussing the methodology of event studies in Section 2, Section 3 presents our two key empirical results. The first result is negative. In particular, we find no evidence of abnormal returns, either positive or negative, when OPEC announces an increase in the quota. Our second result is that significant abnormal returns do accrue following both status quo announcements and announcements of reductions in the quota. However, matters are not so simple. In the case of status quo announcements, we find significantly negative abnormal returns, while for quota reductions we find significantly positive abnormal returns. These results immediately suggest that the market is surprised by status quo announcements and quota reductions, but not by announcements of

[3] For example, Chevron-Texaco holds a 40 percent interest in oil fields offshore the Niger Delta. They note, "In addition, the Nigerian government levies taxes and royalties on Chevron's revenues and oil production, such that total Nigerian revenue from the partnership exceeds 85 percent. This is typical of overseas oil operations in which the government is the primary partner."

[4] To be sure, our empirical results rely on the assumption that markets are largely efficient and rapidly adjust to public information, such as demand shocks or other phenomena that affect the price of oil. Given the size and liquidity of oil markets as well as its importance for the overall economy, we are comfortable making this assumption.

increases in the quota. Both of these results hold generally for the two measures of oil prices and the three stock market indices that we consider. However, why should status quo announcements lead to negative returns, while reductions lead to positive excess returns?

First, note that abnormal returns occur only if the market is taken by surprise. A finding of no abnormal returns for a particular event indicates two things; either the event was well-predicted by the market or the event does not impact the returns of a company [clearly in the current context, the latter is not likely true]. Now consider status quo announcements. The fact that returns are significant indicates that the market is taken by surprise. Moreover, the fact that the returns are negative, implies that demand conditions must have worsened. Had a status quo announcement followed a positive demand shock, surely returns would be positive, since the effect would be to push up the price of oil. Next consider announcements to reduce the quota. Our result is that this leads to positive abnormal returns. Again this indicates that the market is surprised by the announcement. The fact that the returns are positive tells us that demand conditions are likely to have been weak. Finally, consider announcements to increase the quota. Here we do not find any evidence of abnormal returns, either positive or negative. This suggests that the stock market accurately predicts such announcements and is not surprised when they finally come about.

Thus our results from OPEC reiterate a familiar refrain: agreements appear more difficult when times are bad and easier when times are good. In looking for an explanation, one is first drawn to the seminal work of Rotemberg and Saloner (1986) or its many extensions, such as Bagwell and Staiger (1997). However, with only demand uncertainty, in Bagwell and Staiger's most collusive equilibrium, there is a precise mapping from demand to the level of collusion. Given a negative demand shock, the market would correctly predict a lower level of collusion and abnormal returns would not occur. Therefore, in order to generate surprise, there must be another source of randomness in the model. Recent work by Athey, Bagwell and Sanchirico (2004), introduces randomness to firms' marginal costs of production and can generate both surprise and a greater difficulty in sustaining collusion during bad times. In general, in the low demand state, firms try to implement the monopoly price, but firms experiencing a sufficiently low cost realisation, are allowed to undercut the price without fear of punishment. In contrast, in the high state, firms always choose the monopoly price level and there is no desire to undercut since future punishments are quite severe with persistent Markovian demand fluctuations. Thus market share instability is a key feature of the ABS model; this is where the model fails.[5]

From Figure 1 one can see that, with the exception of Saudi Arabia during the first Persian Gulf War, the quota shares of OPEC members have been remarkably stable over time. Therefore, any model which rationalises our empirical results must *not* rely on changing shares to generate predictions.[6] In some sense, this limits the applicability of repeated game models of collusion,

[5] The same can also be said of Staiger and Wolak's (1992) model in which firms choose capacity before demand is realised. For low demand shocks, firms engage in mixed strategy pricing, leading to large swings in market shares.

[6] Indeed, we present anecdotal evidence in Section 4.3 pointing to OPEC's apparent inflexibility in negotiating over shares. One conjecture is that OPEC members do not want to alter shares in response to one country's idiosyncratic shock for fear that the new shares would then become the "reference point" for any future negotiations.

since these models *necessarily* predict changing market shares. Our preferred modeling approach is, therefore, that of bargaining. It has often been noted in the popular press that it takes difficult negotiations to simply change the aggregate production quota, to say nothing of the degree of negotiation required to actually change the shares of aggregate production allocated to OPEC members. Moreover, in a study of various cartels, Levenstein and Suslow (2002, p. 16) argue that "[b]argaining problems were much more likely to undermine collusion than was secret cheating."[7]

Figure 1: OPEC-Member Quota Shares

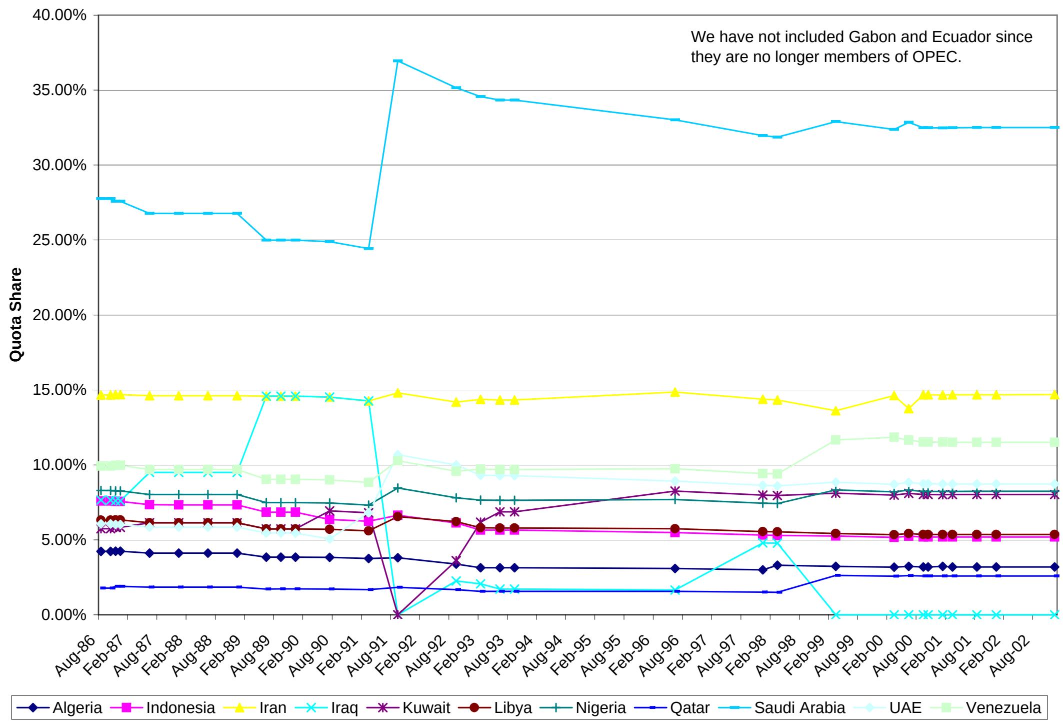

After arguing in Section 4 that our empirical results survive a number of robustness checks, we turn to a theoretical discussion in Section 5. As indicated above, we adopt a bargaining approach. In a static version of the model (*i.e.,* demand is constant), there is a proposer and a responder. We assume that the quota shares are fixed so that the proposer may only offer to adjust the *aggregate* quota. (See footnote 6 and Section 4.3 for further discussion on this point.) If the offer is accepted, then the new quota is implemented, while if not, the status quo obtains. However, we assume that there is one-sided private information. In particular, the responder is subject to a random shock to its marginal cost of production. Importantly, this shock has a potentially different effect depending on the location of the status quo quota. In particular, when the quota is initially very low — so that the proposer wants to increase the quota, the incentives of the proposer and responder types are substantially aligned — all responder types would prefer some increase in the quota. Therefore,

[7]Consistent with our empirical results, they go on to say that "[B]argaining issues may arise as a result of a decline in demand" (p. 18).

agreements to increase the quota are relatively easy when the quota is low. However, as the status quo increases, up to the proposer's optimal quota, the interests of the proposer and more and more responder types start to diverge, making agreements more difficult.

Similarly, when the status quo is initially very high (so that the proposer wants to lower the quota), agreements are relatively easy. However, to the extent that there are enough *low cost* types, the model predicts that it is relatively more difficult to decrease the quota, since there may be a positive mass of responder types who would prefer a higher still quota. Therefore, on the space of initial quotas, the probability of reaching an agreement resembles a horizontally flipped check mark: \/. Moreover, a discontinuity generally appears in the model at the proposer's optimal quota. In particular, if the distribution of responder types is skewed towards low cost types, there is a downward jump discontinuity at the proposer's ideal quota, thus making it potentially extremely difficult to lower the quota.

In Section 5.2, we discuss some issues that arise when we allow Markovian demand shocks and a far-sighted proposer. In this case, we show that many of the results derived in the static model pass through to the dynamic setting. In particular, we show that it is *most* difficult to lower an initially high quota when demand is high but that it is also particularly easy to increase the quota when it is initially low. We end our theoretical discussion in Section 5.3 with a brief look at two alternative explanations for our findings and argue that they are wanting in one or more respects, not least is the ability to make accurate predictions beyond the main empirical findings reported in Section 3. In contrast, we show that a simple extension of our model is able to explain the observed relationship between cheating and quotas for OPEC members.

We conclude the paper in Section 6 with some speculative comments about the current oil market and OPEC specifically.

2 Empirical Methodology

Event studies have been a popular empirical tool in many fields for a great many years. A summary of the standard method can be found in Campbell *et al* [1997]; however, for completeness, we give a brief description. Initially, consider an event which occurred at time T^* and is believed to have an impact on the returns of stock i. We proceed in two steps. First, we must get an estimate of the normal returns of stock i. To do this we estimate:

$$R_{it} = \alpha + \beta R_{mt} + \epsilon_{it} \tag{1}$$

over the period $t = T^* - K - L_1, \ldots, T^* - L_1 - 1$, where L_1 is the number of days before the event in our window, K is the *normal returns* estimation period and R_{mt} denotes the market return at time t. In studies with stock prices, K is typically taken to be 250 or 150 (Campbell *et al* [1997]). Second, we calculate abnormal returns during the event window as:

$$\hat{\epsilon}^*_{it} = R_{it} - \hat{\alpha} - \hat{\beta} R_{mt} \tag{2}$$

where here $t = T^* - L_1, \ldots, T^* + L_2$, and L_2 is the number of days after the event.

Now consider a collection of events which occurred at times $t_1, \ldots, t_N$ and define ϵ^*_{ij} to be the vector of estimated abnormal returns for event j and stock i. Then define

$$\bar{\epsilon}^*_i = \frac{1}{N}\sum_{j=1}^{N} \epsilon^*_j$$

to be the sample average of abnormal returns for a sample of N events. Finally define the cumulative abnormal return (CAR) for stock i to be:

$$\overline{CAR}_i(t) = \sum_{k=T^*-L_1}^{t} \bar{\epsilon}^*_i(k)$$

where $t = T^* - L_1, \ldots, T^* + L_2$ and $\bar{\epsilon}^*_i(k)$ is the k^{th} element of the vector $\bar{\epsilon}^*_i$. Then, using standard techniques we can test for significance of abnormal returns and cumulative abnormal returns over the event window.

This method has been shown to be quite robust to certain deviations from the normality assumption on the errors but it still suffers from a number of drawbacks which we would like to remedy in the present analysis. As in most event studies, we have the problem of choosing the appropriate event window. We do not have a general data-driven method for calculating the window; instead, we conform to most event studies and experiment with windows of various length ranging from 1 to 20 days on either side of the event. The standard method has also been criticised because it does not account for changes in variance due to increased uncertainty during the event period (see Lockwood and Kadiyala [1988]). However, as reported in section 4, tests for changes in variance indicate that this is not a significant problem in the present setting.

Remark 1. An alternative approach to the event study is taken by Guidolin and La Ferrara (2004) in which events are classified into one of a number of types and a dummy variable is created for each event type which takes value 1 in a small window around all events of a given type and zero otherwise. The model is then estimated over the entire sample and the influence of an event is given by the sign and magnitude of the coefficient on the dummy variable. We choose not to follow this approach as we believe that using the entire sample *glosses* over market expectations at the time of each event in a way that the traditional event study does not. We simply note that the results from this approach are not qualitatively different from the approach taken in our paper.

3 Empirical Results

Our events are announcements by OPEC regarding their production quotas. Announcements are divided into three types — quota reductions, quota expansions and status quo announcements. We have data from 1986, when the current quota system was agreed upon, through to September 2002. OPEC typically has between two and four regularly scheduled meetings but at times will announce *extraordinary* meetings if it believes market conditions are particularly uncertain.[8] Our sample contains 50 announcements by OPEC. Twelve of these announcements are for reductions, 16 for increases and 22 maintain the status quo quota.[9]

As we have argued previously, in adjusting its quota, OPEC attempts to influence the price of oil to increase their profits. Through this direct effect on oil prices, OPEC's actions, therefore, indirectly influence the profitability of oil companies [recall that the relationship between a western oil company and an OPEC member often resembles sharecropping]. Therefore, the results of our event study should be largely similar if we conduct it on oil prices or on indices of oil companies. Indeed, our conjecture is that the effects will be greater the *closer* is the variable of interest to actual oil price movements. Therefore, we expect strongest results for oil prices, weaker results for exploration and drilling companies (as they are one step removed from oil price movements) and even weaker results for integrated companies (this follows since a substantial portion of their revenues are derived from refining operations, which use oil as an input). Given that our index of all oil companies is a weighted average of exploration & drilling, integrated and service companies, we expect these results to lie somewhere between those of exploration & drilling and integrated oil companies.

Our stock market data includes market indices for integrated companies, exploration & drilling companies as well as oil companies as a whole. For these indices, we take the S&P 500 as our market basket.[10] We also have data on the two-month forward price and the spot price of oil (West Texas Intermediate Crude Oil). Daily data from 1 January 1986 to 30 September 2002 were obtained from Datastream. The index for oil companies as a whole is a composite of sub-indices for integrated, exploration & drilling and service companies. By market capitalisation, integrated companies comprise 65%, exploration & drilling comprise 21% and service companies the remaining 14%. Note that the index for integrated companies includes companies such as Exxon-Mobil and Chevron-Texaco, which are truly integrated, but it also includes companies such as Valero and Sunoco which derive most of their revenues from refining operations. In this section, we focus exclusively on the results and interpretation of our event study using the full sample of events at our disposal. In the next section, we take pains to show that our results are robust. To do this, we provide further statistical analysis, discuss possible endogeneity concerns regarding the timing of meetings and provide some anecdotal evidence in favour of our interpretation.

3.1 Cumulative Abnormal Returns

In this subsection we report results using the traditional method of Campbell *et al* [1997]. We take as our event window 20 days before and after the date of the event and 150 days prior to the

[8] These extraordinary meetings are typically announced at the conclusion of a regularly scheduled meeting and usually occur between one and three months later. We discuss this further in section 4.2 and appendix B.

[9] The events were obtained from various issues of the New York Times and official OPEC press releases.

[10] In particular, we have the series OILEPAM, OILINAM and OILGSAM from Datastream, which are indices of the exploration & production, integrated and the energy sector as a whole for the Americas.

beginning of the event window as our normal returns estimation period.[11]

Figure 2, which plots the cumulative abnormal returns functions for exploration & drilling companies, is typical of the results that we obtain in our event study. Identical plots for the other indices used in our event study (all oil companies and integrated oil companies) as well as for the spot and two-month forward prices of oil are given in Figure 7 in Appendix A.[12] In particular, we find negative cumulative abnormal returns for status quo announcements which start to accrue in the days leading up to the announcement and become significantly negative *after* the announcement. Indeed, for exploration & drilling companies and for the oil industry as a whole, the cumulative abnormal returns are significantly less than zero at the 5% level (one-sided test) towards the end of the event window. It is also important to note that these abnormal returns are quite large in percentage terms. Twenty days following the announcement, the cumulative abnormal returns are between -2 and -3.5%. Thus it appears that OPEC, by its inaction, has the ability to significantly alter stock returns in the days following such announcements.

Figure 2: CAR Function: Exploration Companies

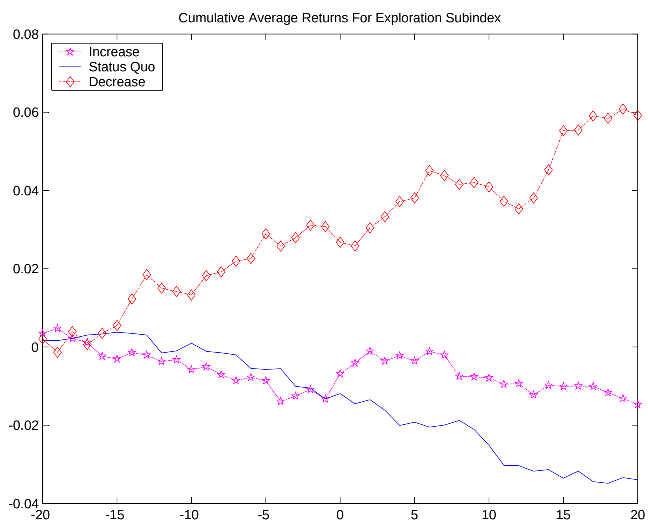

Next consider reductions in the quota. The results are weaker, though still telling. For both the two-month forward price of oil and exploration companies, the cumulative abnormal returns become significantly positive at the end of the event window. Indeed, here the effect is quite strong; in the former case, twenty days following the event, the cumulative abnormal returns are nearly 10% while in the latter case they are approximately 6%. While the CAR function is significantly different from zero only in these two cases, if one looks more closely at the plot of the CAR functions, there is an interesting pattern which is consistent across all groups. In particular, in the days leading up

[11]We experimented with various event window lengths and normal returns periods. The magnitude of some estimates vary a little; however, the pattern and significance, in general, do not.

[12]The CAR functions for oil prices are derived in a slightly different manner. In particular, we consider $R_{it} = \alpha + \epsilon_{it}$ as our normal returns model (see Campbell *et al* [1997] for more details). In addition our normal returns estimation period is taken to be 150 days and the event window is 15 days before and after the event date.

to an announced reduction in the quota, positive abnormal returns accrue, peaking approximately 3 days before the announcement and then make a significant fall, only to recover again at the end of the window. Finally, after some days and once OPEC's resolve becomes clear the markets reaction turns favourable — indeed, if one restricts analysis to days 12 through 20, significantly positive abnormal returns accrue over this time.[13] This finding is significant for all of the sub-indices at the 5% level.

What could be the cause of this? One might suspect that before the commencement of an OPEC meeting, the market is hopeful of an agreement and so positive returns accumulate. However, as the meetings begin, perhaps due to uncertainty or a profit-taking motive, prices decline significantly. Finally, after the announcement and in the days following, OPEC's resolve becomes more apparent and the positive effects of the quota reduction take hold.

One does wonder if there is a perverse explanation for this finding, unrelated to market fundamentals or expectations. However, the two events which had the biggest impact on oil company stock prices and oil prices (11 September 2001 and the Iraqi invasion of Kuwait) do not coincide properly with OPEC meetings (or their outcome) in order to generate such a result. Thus, one is tempted to believe that market psychology about OPEC's resolve is, in fact, generating the result. Furthermore, as with the finding for status quo announcements, the cumulative returns at the end of the window are between 2.3 and 5.9%. So again, the impact of such announcements by OPEC is rather large.

Finally we consider announcements of quota increases. Here the results fall into two different categories. With the exception of integrated oil companies, announcements of increases in the quota were met with no significant abnormal returns - either positive or negative, while for integrated companies, the impact of a quota increase seems generally positive, though not significant.[14]

Remark 2. As can be seen from Tables 3 through 5, the CAR functions for quota reductions were not all significantly different from zero at the end of the event window. However, if one takes the difference between the CAR function for reductions and that for status quo announcements, the result is significantly positive at the end of the event window. Thus, the market reaction is very different for an announcement of a quota reduction than it is for a status quo announcement.

3.2 Interpretation and Discussion

As alluded to in the introduction, the results indicate that OPEC has an asymmetric ability to secure agreements depending on whether the economy is in a good state or a bad state of the world. First, the fact that neither significantly positive nor significantly negative abnormal returns are associated with an increase in the quota implies that such events are anticipated by the stock market. Therefore, when such an announcement is made, say due to a positive demand shock, it does not come as a surprise. Second, the fact that abnormal returns do occur for the other

[13]Day 12 was chosen as it appeared to be a trough for all three sub-indices. Furthermore, in two cases there was a significantly negative return on day 11.

[14]However, the abnormal returns are significantly greater than those which accrue following inaction by OPEC.

announcements implies greater uncertainty in OPEC's ability to agree to a reduction following a negative shock. To see this, suppose that demand is lower and OPEC should reduce its quota. Given the realisation of demand, the market forms expectations about whether an agreement will occur or not; however, these beliefs are much more diffuse than in the case of a demand increase. Thus when the announcement is made, the market reacts to either announcement and abnormal returns of the kind reported above are quite likely to occur.

To us, this is an eminently reasonable explanation of what happens here and in many other bargaining situations. However, it is just one possible interpretation and so far rests only upon a logical argument. Thus, before we can be satisfied that it is the correct interpretation, we must provide further evidence. Before proceeding, we first dispel two alternative explanations. First, it may be that the market correctly predicts the distribution over outcomes but is uninformed about demand and supply conditions. Therefore, abnormal returns occur but not for the reasons that we argue. However, we find this argument fundamentally implausible — timely information on all aspects of the oil market is widely available by the Energy Information Agency and the International Energy Agency. Therefore, lack of information on market conditions cannot be a suitable explanation for the results we report. Moreover, this explanation would also rely on a particular correlation between demand shocks and OPEC meetings, which, given that most meetings are regularly scheduled cannot be true. Second one may argue that an agreement to adjust the quota, either up or down, signals that OPEC as a cartel is strong but that status quo announcements signal weakness in the cartel. To see this argument also fails, suppose that it were true. Then we would not see an asymmetric pattern of abnormal returns; instead, we should observe positive abnormal returns for both increases and decreases in the quota. In the remainder of this section we attempt to show that OPEC generally increases the quota during good times, while during bad times, they sometimes lower the quota and sometimes leave it unchanged.

We first take a backward-looking approach. If we adopt the view of OPEC as a residual supplier of oil, then we may define "good times" as a situation in which residual demand has risen from the previous year, where residual demand is defined to be OECD consumption minus non-OPEC production. Beyond this, many have argued that inventories play an important role in determining the price of oil. In terms of prospects for OPEC, "good times" would be represented by a decrease in inventories from the previous year. Therefore, we may consider the following simple estimation:

$$\begin{aligned} y_i &= \alpha + \beta \Delta \text{stocks}_i + \gamma \Delta \text{resid}_i + \epsilon_i \\ y_i^* &= \begin{cases} 1, & \Delta \text{quota} > 0 \\ 0, & \text{o.w.} \end{cases} \end{aligned}$$

where Δstock_i is the year-on-year change in crude oil stocks at the time of event i and Δresid_i is the year-on-year change in residual demand for oil. If our interpretation is correct, we expect $\beta < 0$ and $\gamma > 0$. Indeed, as Table 1 indicates, this is precisely what we find.

Our second approach is more forward-looking. In making its decisions, OPEC must form expectations about the expected path of residual demand for crude oil, changes in inventories and

Table 1: **Probit Regression**

	Coef.	Std. Err.	z
Δstocks	-0.0056145	0.0020409	-2.75
Δresid	0.0002775	0.0001328	2.09
Constant	-0.3548164	0.2208406	-1.61
	$n = 49$	$LR(2) = 9.64$	

the price of oil. Our strategy, therefore, is to forecast these variables for the quarter immediately following an OPEC meeting. With these forecasts, we can get some idea of the *oil market balance* and determine, in some sense, the *required* course of action by OPEC. We can then compare the required course of action with the actual outcome to determine any observable patterns of OPEC behaviour. The data used for this analysis were provided by the Energy Information Agency.

If one believes in an efficient market, then it should be the case that the price of oil acts as a summary statistic for the oil market balance. That is, if prices are expected to rise it indicates tightness of the market, say because of strong demand growth. Thus prices which are expected to rise may indicate a need for OPEC to increase its quota, while prices which are expected to fall, may indicate a need to reduce the quota. Consider first status quo announcements. In 17 out of the 21 status quo announcements, we found that the two-month forward price of oil was expected to fall. Moreover, of the four cases in which the price was expected to rise, barring any action by OPEC, one coincided with the first Persian Gulf War and the remaining three were after the Asian Financial Crisis.[15]

Now consider quota increases. If the market expected inaction from OPEC (and if demand is increasing, which we show to be the case) then one would predict rising oil prices. Therefore, forecasts indicating declining oil prices supports our hypothesis that increases in the quota are largely predicted. Indeed, we find this to be so. Depending on the exact method of forecasting, between 10 and 17 (of 17) of the events indicate falling prices. Finally, consider quota reductions. If oil prices were forecast to rise, one could argue that the market expects a reduction in the quota; on the other hand, it is consistent with our hypothesis that reductions are not as accurately predicted if prices are expected to fall. Indeed, we find in 9 (of 12) cases that prices are predicted to fall. What is more, in three instances (Dec. 1986, Mar. 1991 and Jan. 2001), the price of oil, though forecast to fall, actually saw a large increase in the following month(s).

Remark 3. Note that we cannot conclusively say that just because oil prices were forecast to decrease for 17 of 21 status quo announcements, it means that the outcome of the meeting was either known or unknown. However, this is not our goal. Our goal is to show that status quo announcements occur, on average, during periods of low demand. On this front, we believe our result does provide strong evidence. Indeed, suppose that demand were predicted to rise — the only way in which one could justify forecasts indicating falling oil prices is if the market also predicted that OPEC would

[15] As the reader shall presently see, this result is in line with the anecdotal evidence pointing to a shift in OPEC behaviour around the time of the Asian Crisis.

increase the quota. However, this would imply that the market is frequently wrong. In this case, one would expect a positive impact on oil prices in the period immediately after the announcement. However, one can see from Figures 7 and 8 that this does not occur. Therefore, the fact that most of our forecasts indicate falling expected prices, leads us to conclude that status quo announcements occur during periods of weak demand.

While we may think of oil prices as a summary statistic for the oil market balance, we may also conduct a similar forecasting exercise on some important determinants of oil prices; the two that we focus on here are residual demand for oil and OECD crude oil inventories. Good times are identified by an expected increase in residual demand and an expected fall in inventories, respectively, with bad times indicated by the opposite forecasts. Consider first residual demand forecasts. For the 21 status quo announcements, between 8 and 12 (depending on the method of forecasting) of the announcements involved falling residual demand. We also repeat this analysis for both quota reductions and quota increases. For quota reductions, approximately 9 (of 12) indicate declining demand, while for quota increases between 10 and 12 (of 17) indicate increasing demand. Thus there is weak evidence in favour of our interpretation. Now consider inventories. We forecast the change in inventories over the previous years. For between 16 and 18 (of 21) instances, we find that inventories are forecast to be *above* the previous year's level, indicating little upward pressure on oil prices.[16] Moreover, if we simply compute month-on-month forecasts, we find that for over half of the status quo announcements, inventories were expected to increase, again indicating little pressure on oil prices (see Kaufmann *et al* (2004)). Taken as a whole, our forecasts on oil prices as well as net consumption growth and inventories indicate that the underlying market conditions called, more often, for a reduction in the quota when the actual course taken by OPEC was to remain with the status quo.

4 Robustness Checks

This section has a number of purposes. First, we subject our empirical results to further robustness checks. In particular, we report a number of non-parametric tests on the nature of abnormal returns and measures of variance. Second, we discuss some potential issues with our empirical methodology and the provided interpretation. Finally, we provide some anecdotal evidence in favour of our interpretation.

4.1 Non-Parametric Tests

Even though the (cumulative) abnormal returns may not be significant in any time period, we are still able to say something about their nature. Specifically, a Wilcoxon signed rank test tests for whether or not more than half of the abnormal returns are of a particular sign. For each of the

[16] This is not a perfect measure because if we think that demand also grows year-over-year, then the same inventory will indicate *fewer* days of oil in reserve the higher is demand. However, we believe our measure has some merit; of the 12 times OPEC reduced the quota, 11 of 12 times the change in inventories was forecast to rise, while for 8 of the 17 times that OPEC increased the quota the change in inventories was forecast to fall.

three sub-indices we test the null hypothesis that 50% of the returns are positive for each class of event. Considering status quo announcements, with the exception of the integrated companies sub-index, we are able to reject this hypothesis in favour of the alternative that more than 50% of the abnormal returns are negative. Moreover, for the two-month forward price of oil, there is weak evidence that more than half of the abnormal returns were positive for quota reductions.

Table 2: **Wilcoxon Signed Rank Tests**

Index	Q↑	Q↔	Q↓
Two-Month Forward Price	1.001	-2.5301	1.7526
Oil Companies	0.1328	-2.3357	0.2365
Integrated Companies	0.9492	-1.8174	0.4438
Exploration Companies	-1.1954	-1.9988	0.6706

In criticising the standard event study techniques Lockwood and Kadiyala [1988] and Cyree and Degennaro [2002] hint that changes in variance during the event window may be important. For each event in the data we can conduct a simple F-test, where the null hypothesis is that the variance term σ^2 is the same during the normal returns period and during the event window. For the oil sub-index we estimated that for 7 of the 50 events the variance was different during the two periods. The corresponding numbers for the integrated companies and exploration companies sub-indices are 5 and 3, respectively. This suggests that changes in variance are not as important as one might think. Furthermore, a non-parametric test similar in spirit to the Wilcoxon signed rank test never rejects the null hypothesis that variance is constant for each of the three different classes of event.

4.2 Potential Issues With Empirical Methodology and Interpretation

In the previous section, we reported results based on the entire sample of events (OPEC announcements) for the three oil company indices and two measures of oil prices. However, a number of issues may arise which we address here. First, part of our identification comes from the assumption that the timing of the events themselves are exogenous. Under this assumption, through the abnormal returns, we are capturing the information value of the actual event. While OPEC typically has between two and four regularly scheduled meetings a year, it will sometimes hold extraordinary meetings.[17] Typically OPEC gives substantial advanced notice of an extraordinary meeting, but this still raises some endogeneity concerns and a potential reverse causality. Regarding extraordinary meetings, one may hold one of two views. First, it may be that extraordinary meetings are held during times of extraordinary demand or price fluctuations, which may imply that OPEC is more likely to come to an agreement on necessary quota adjustments.[18] Under this scenario, announcements of increases or decreases in the quota would be *less* surprising, while status quo announcements, signaling disagreement, would be *more* surprising. Therefore, if we look at only

[17] In our sample, there were nine extraordinary meetings; each possible announcement occurred three times.

[18] For example, in the model we present below, all else equal, given two negative demand shocks, the model predicts that there is a greater chance of agreement for the larger shock.

normal meetings, we should see that the CAR functions for status quo announcements are more negative, while the CAR functions for reductions and increases are closer to zero. On the other hand, it may be that extraordinary meetings are held at times of extraordinary discord amongst OPEC members, implying that disagreement is even more likely. In this case, by restricting attention to normal meetings, the opposite predictions would hold true. In Appendix B, we briefly report the results of the event study excluding extraordinary meetings. The results are not much different from those of section 3; in some cases, the CAR functions are greater (in absolute value) than those of section 3. This is particularly apparent if one compares the CAR functions for the two-month forward price of oil using all meetings with those based only on ordinary meetings. The results also lend some support to the former conjecture — namely, extraordinary meetings occur signal that OPEC is *more* likely to agree.

Second, it is our interpretation that when demand is low/decreasing, OPEC finds it difficult to agree to lower production, while when demand is high/increasing, there are no internal impediments within OPEC to reaching an agreement. However, there remains the possibility that there is some omitted variable that is driving our empirical results but has nothing to do with OPEC negotiations. Such a variable, if it exists, would invalidate our interpretation. However, note that to generate our empirical results, such an omitted variable must satisfy a number of properties. First, because there are no abnormal returns for quota increases, the variable only has an effect during negative demand shocks. Second, it may or may not directly influence the returns of the companies in our sample, but it must affect OPEC decisions — in particular, for some realisations of the variable, it causes inaction by OPEC, while for others it causes OPEC to reduce the quota. Third, the variable must be unobserved by the market but the resolution of uncertainty surrounding this variable is resolved at the time of the OPEC meeting.

In our opinion, it is difficult to think of a variable which, unrelated to OPEC's bargaining problem, satisfies all of these criteria. Three possibilities emerge; however, we feel that none of them meet all of the requirements necessary to invalidate our interpretation. The first possibility, which satisfies the three criteria is that of outside political pressure. Following a negative demand shock, consuming nations may lobby OPEC not to lower production. While lobbying would be partially observable to a market, its effectiveness likely would not be, until OPEC actually makes its announcement. It also seems reasonable that lobbying would be more common during bad times as here the interests of OPEC and consuming nations are opposed, whereas interests are more in line when demand is high — both would like increased production; consuming nations want higher production to lower current prices, while OPEC wants higher production so as not to encourage exploration by non-OPEC countries or research into enhanced fuel efficiency. However, where the explanation of outside lobbying fails is that it is not distinct from OPEC's bargaining problem. In particular, it points to a specific dimension over which OPEC must bargain — how accommodative to be towards consuming nations in any quota decisions.

Finally, it is reasonable that non-OPEC supply shocks could influence OPEC's optimal decision. However, this explanation fails on two grounds. First, such shocks are observable by the market;

second, there is no reason why such shocks should only influence OPEC's decision when demand is already low. Another possibility is that of cost shocks to OPEC producers. Such shocks may well be unobservable but we still do not find this to be a plausible explanation. First, for many of the largest OPEC producers, it does not seem reasonable that there would be such wide variations in the cost of production so as to make it unprofitable to lower production following a negative demand shock. Moreover, while theoretically (see *e.g.,* Athey *et al* (2004)) cost may interact with collusion only when demand is low, empirically, it leads to predictions different from observed OPEC behaviour. Indeed, it seems to us a difficult task to find an omitted variable satisfying the aforementioned three criteria that is capable of generating our empirical findings. Therefore, we believe that our interpretation that our empirical results inform us of OPEC's differing ability to secure agreements over the business cycle remains plausible.

4.3 Anecdotal Evidence

There is plenty of anecdotal evidence which suggests that OPEC is a good test of our intuition that agreements are more difficult when times are bad rather than when times are good. For example, in its June 1988 production announcement maintaining the status quo, the New York Times reports[19] that both Iraq and the United Arab Emirates exempted themselves from the quota. This happened despite the fact that oil prices were low and overproduction was widely noted. Again, in 1993, Kuwait held up an agreement that would have lowered production in an attempt to strengthen weak oil prices with its demands for even more production.[20] The example of Kuwait also indicates that often OPEC's decision making is *all-or-nothing*; it would seem that OPEC could have *both* lowered production minimally *and* accommodated Kuwait's desire for greater post-war production and this would have been the jointly profitable action; however, OPEC chose instead to maintain the (inefficient) status quo, rather than let one country gain disproportionately.[21]

A more detailed analysis of these reports can also provide further insight into the market psychology following OPEC's decision not to change the quota. Here we find some evidence of a *structural break* occurring around the time of the Asian Financial Crisis. Specifically, from December 1987 to June 1997 there were 16 status quo announcements made by OPEC. In only two of these events could one unambiguously make the case that the alternative choice on the table was an increase in production.[22] In some cases the status quo appeared to be the decision that was called for, while in the majority of others, the published news reports cite disagreement in being able to reduce the quota or that OPEC's actions signal an even higher likelihood of cheating.[23]

[19] See, in particular, "Divided OPEC Puts Off Output and Price Moves," New York Times, 15 June 1988, D2.

[20] See, in particular, "Price of Crude Oil Tumbles As Kuwait Rebuffs OPEC," New York Times, 11 June 1993, D15.

[21] Indeed, this episode would seem to go against the predictions of Athey *et al* (2004); Kuwait received a very low shock and so, according to their model, should have been able to over-produce without fear of reprisals. This did not happen.

[22] The two events where one could not make this argument occurred in November 1994 and November 1996. For more details, see "OPEC Limits Oil Output Through '95," New York Times, 22 November 1994, D1 and "OPEC Agrees To Freeze Oil Quotas For 6 Months," New York Times, 29 November 1996, D11.

[23] For other examples, see also 26 November 1993, D14, 28 March 1994, D2, 23 November 1995, D7 and 8 June

In contrast, after the Asian Financial Crisis, old rivalries between Iran and Saudi Arabia seemed to have lessened, leading to a shift in OPEC philosophy. In the eight status quo announcements between July 1997 and September 2002, in only one case can it be read as a failure of OPEC to reduce the quota to induce a price increase, while in three cases the action appears to have been taken to drive prices up even further. We will have more to say on this subject in the conclusion where we speculate about OPEC and the oil market in the near future.

5 Theoretical Discussion: One-Sided Private Information

In this section we present a simple model which rationalises the main empirical results derived in the previous section. We first discuss a simple static version of the model in which very sharp predictions are possible. We then extend the model to allow for demand fluctuations and far-sighted behaviour. However, before we proceed, recall the three main points which came out of our empirical discussion:

1. Agreements are very likely following an increase in demand (or when demand is increasing).
2. Agreements are less likely following a decrease in demand (or when demand is decreasing).
3. The shares of each member are very constant over time, suggesting that bargaining is mostly over setting the total quantity and less focused on reallocating shares. The constancy of shares is largely maintained for both increases and decreases in the quota.

5.1 The Static Game

Consider the following set up. There are two players in a cartel, facing a linear inverse demand curve given by $p(A, Q) = A - Q$. Given an initial quota of Q_0 and shares $(s, 1-s)$, the proposer proposes a new quota, Q. If the offer is accepted, then the new quota is implemented according to the division $(s, 1-s)$, while if the offer is rejected, the status quo quota is implemented, again according to $(s, 1-s)$. That is, we assume that the relative shares are fixed and the only variable that the cartel members can negotiate over is the level of aggregate production. We also abstract entirely from any issues of cheating and simply assume that the cartel members can sign binding agreements.

We initially assume that the proposer has a constant marginal cost of production, which we normalise to zero, while the responder has a cost of production given by $c(q) = cq + (b-1)q^2$. For simplicity, we will normalise $c = 0$. We introduce private information into the model in the following way: at the time of making a proposal, the cost function of the responder is unknown by the proposer. Instead, she knows that b is a draw from some distribution, F.

We assume that the cost functions of responder types vary in the quadratic term primarily for starkness. What is important for the result is a *fanning out* property, which can be seen in Figure 3. That is, when the initial quota is zero, the incentives of the proposer and all responder types are

1996, A31 — all of the New York Times.

aligned — all would accept some (small) increase in the quota. However, as the quota increases, the profit functions of the responder types begin to *fan out*, implying that their interests begin to diverge. By introducing the uncertainty in the quadratic term, the curvature of the profit function changes in b, making our results much starker than if we had assumed that the shocks were linear in q.

This is the basic model. Our task is then to examine its equilibrium and determine conditions under which agreement is more difficult during "bad" times than during "good" times, as our empirical study indicates. In this static setting, bad times will be indicated by an initial quota which is relatively high in comparison to A, while good times will be indicated by a quota which is relatively low in comparison to A. For notational simplicity, we make one simplifying assumption: we write our expressions *as if* the responder faced the inverse demand curve $A - bQ$ and b is drawn from a distribution, F, with non-negative support.[24] Finally, since the shares, $(s, 1-s)$ will not play any role in the analysis, we will simply omit them from the relevant expressions.

Remark 4. It may seem a bit odd that there is only one-sided private information in this model. One may, instead, think of a model in which *both* the proposer and responder independently draw b from the same distribution. This may be the more appropriate modeling choice if one believes that all players are completely symmetric. However, in the present setting, all players are not symmetric — OPEC members vary by size, proven reserves and relations with western (consuming) nations. It may be that some countries are simply more subject to shocks, say due to political instability or ethnic tensions, or that some countries have differing interests *vis-à-vis* oil consuming nations. Finally, in this setting, it is very easy to delineate between good times and bad — we can use the proposer as a guide post. If the quota is above its optimal level, this may indicate bad times (say because demand has just fallen), while if the quota is below its optimal level, this may indicate good times (say because demand has just risen).

With these preliminaries in mind, let us begin to solve the model. Given an offer of Q and a status quo quota of Q_0, the responder will accept if and only if $Q(A - bQ) \geq Q_0(A - bQ_0)$, which can be rewritten as:

$$b \left\{ \begin{array}{c} \geq \\ \leq \end{array} \right\} \frac{A}{Q + Q_0} \left\{ \begin{array}{c} \text{if } Q_0 > Q \\ \text{if } Q_0 < Q \end{array} \right\} \tag{3}$$

Therefore, from the perspective of the proposer, the probability that the offer is accepted is given by:

$$P(A, Q, Q_0) = \begin{cases} 1, & \text{if } Q = Q_0 \\ 1 - F(\frac{A}{Q+Q_0}), & \text{if } Q < Q_0 \\ F(\frac{A}{Q+Q_0}), & \text{if } Q > Q_0 \end{cases} \tag{4}$$

[24] Another interpretation is the following. Suppose that the proposer and the responder produce a differentiated product and that b represents the slope coefficient of the responder. That b is uncertain can be interpreted thusly: the responder enters into side agreements with buyers, selling their product at a price potentially different from the market price. However, the proposer does not know the relative bargaining power in these negotiations of the responder or how receptive she will be to the buyers' entreaties for a lower price. Under this view, Saudi Arabia would be the responder and other OPEC members do not know how accommodating it will be to U.S. demands for "reasonable" oil prices.

Figure 3: Profit Functions For Proposer and Various Responder Types

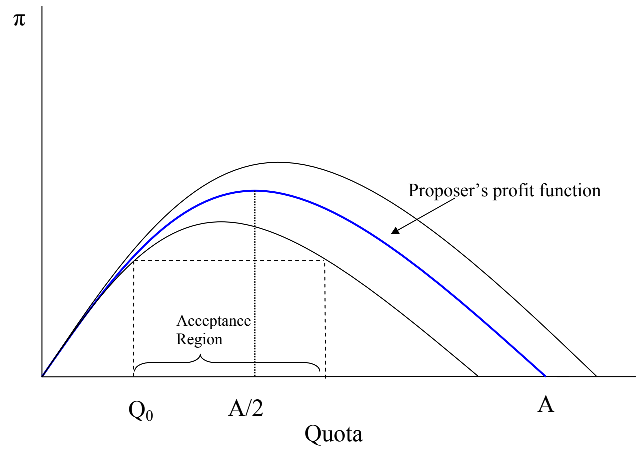

This allows us to write the expected profit function of the proposer as:

$$\mathbb{E}[\pi_p(A, Q, Q_0] = Q_0(A - Q_0) + P(A, Q, Q_0)\left[Q(A - Q) - Q_0(A - Q_0)\right] \tag{5}$$

Notice that, despite the fact that $P(A, Q, Q_0)$ is discontinuous whenever $Q = Q_0$, the expected profit function is continuous in Q, and so a maximum exists. That being said, the expected profit function need not be a globally concave function of Q, for at $Q = Q_0$, there will typically be a kink. This will not generally pose many problems for our purposes. Intuitively, the proposer's ideal point is $Q = Q_0 = \frac{A}{2}$; therefore, at the optimal solution, it must be the case that whenever $Q_0 < \frac{A}{2}$, $Q^* \in [Q_0, \frac{A}{2}]$ and whenever $Q_0 > \frac{A}{2}$, $Q^* \in [\frac{A}{2}, Q_0]$. Hence, the fact that a kink will generally occur at $Q = Q_0$ will play no role in the analysis, since on either side of $Q = Q_0$, the expected profit function will generally be concave in Q over $[\frac{A}{2}, Q_0]$ or $[Q_0, \frac{A}{2}]$.

Let $Q^*(A, Q_0)$ denote the optimal offer by the proposer given demand, A, and status quo quota, Q_0. Our primary interest is in characterising $P(A, Q^*(A, Q_0), Q_0)$ and, in particular, deriving conditions under which the probability of agreement is lower when the offer is to reduce the quota. We break the analysis into three parts.

Consider first the situation in which $Q_0 < \frac{A}{2}$. In this case, we know that $Q^*(A, Q_0) \in [Q_0, \frac{A}{2}]$. However, we may say more:

Result 1. *[1] The optimal proposal $Q^*(A, Q_0)$ may be increasing or decreasing in Q_0; however, there exists a threshold $\underline{Q}_0 \leq \frac{A}{2}$ such that $Q^*(A, Q_0)$ is increasing in Q_0 for all $Q_0 \geq \underline{Q}_0$. [2] Moreover, the probability that the optimal offer is accepted, $P(A, Q^*(A, Q_0), Q_0)$ is a (weakly) decreasing function of Q_0.*

The intuition for this result is straightforward. Consider Figure 3. When the quota is very low, one can see that the interests of the proposer and all responder types are aligned — in particular, all responder types are willing to accept *some* small increase in the quota. However, as the initial quota increases, the profit functions of the responder types begin to *fan out*, and the interests of the proposer and some responder types begin to diverge. Beyond this, as is apparent from the figure, as Q_0 increases, the acceptance region of each responder type is shrinking (and possibly becomes empty). Therefore, as Q_0 increases up to $\frac{A}{2}$, it becomes more and more difficult to make an acceptable offer. That $Q^*(A, Q_0)$ may be decreasing is easy to see. For some distributions, it may be that at $Q_0 = 0$, all types will accept a move to $\frac{A}{2}$; however, eventually as Q_0 increases, there are some high responder types that would reject such an offer. In this case, the proposer may *decrease* the offer to ensure that all responder types still accept the proposal. However, eventually, in order to ensure acceptance, the offer would have to be reduced too much. Therefore, at $\underline{Q}_0$, the proposer prefers to *increase* her offer and risk the possibility that some responder types may reject her proposal.

Next consider the case in which $Q_0 > \frac{A}{2}$, implying that the proposer offers $Q^*(A, Q_0) \le Q_0$. In this case, we obtain the following:

RESULT 2. *[1] The optimal proposal $Q^*(A, Q_0)$ may be increasing or decreasing in Q_0; however, there exists a threshold $\bar{Q}_0 \le A$ such that $Q^*(A, Q_0)$ is increasing in Q_0 for all $Q_0 \in [\frac{A}{2}, \bar{Q}_0]$. [2] The probability that the optimal offer is accepted, $P(A, Q^*(A, Q_0), Q_0)$ is a (weakly) increasing function of Q_0. [3] In general, $P(A, Q^*(A, A), A) \le P(A, Q^*(A, 0), 0)$.*

The intuition for parts [1] and [2] are similar to Result 1. The intuition for [3] is simple. When $Q_0 = 0$, all responder types are willing to accept some increase in the quota, since by rejecting, their status quo profits would be 0. In contrast, due to the *fanning out* as b increases of profit functions, when $Q_0 = A$, all responder types with $b < 1$ have strictly positive status quo profits; moreover, any responder types with $b \le \frac{1}{2}$ would like a *larger* quota, while the proposer always offers a *smaller* quota. Therefore, these responder types will surely reject the offer. Thus the asymmetry is apparent: when Q_0 is low, all responder types and the proposer have their interests in alignment, while when Q_0 is large, there may be responder types whose interests are opposite to the proposer, leading to a greater probability of disagreement when Q_0 is high.

The above two cases were quite easy because either the proposer wanted to strictly increase the quota or strictly decrease the quota. However, at $Q_0 = \frac{A}{2}$, the proposer is at her ideal point and optimally chooses $Q^*(A, \frac{A}{2}) = \frac{A}{2}$. Recalling that $P(A, Q, Q_0)$ is discontinuous whenever $Q = Q_0$, it suggests the possibility of a discontinuity. Indeed, we have the following:

$$\begin{aligned} P^- &= \lim_{Q_0 \nearrow \frac{A}{2}} P(A, Q^*(A, Q_0), Q_0) = F(1) \\ P^+ &= \lim_{Q_0 \searrow \frac{A}{2}} P(A, Q^*(A, Q_0), Q_0) = 1 - F(1) \end{aligned} \tag{6}$$

Therefore, we will have a discontinuity at $Q = Q_0 = \frac{A}{2}$ whenever the median of the distribution of responder types is different from 1. Say that the game is *balanced* if $F(1) = 1 - F(1) = \frac{1}{2}$ and *unbalanced* otherwise. In a balanced game, exactly half of the responder types prefer a quota larger than the proposer's ideal quota and exactly half prefer a smaller quota — hence there is no

Figure 4: Equilibrium Behaviour: Static Game, Uniform Dist.

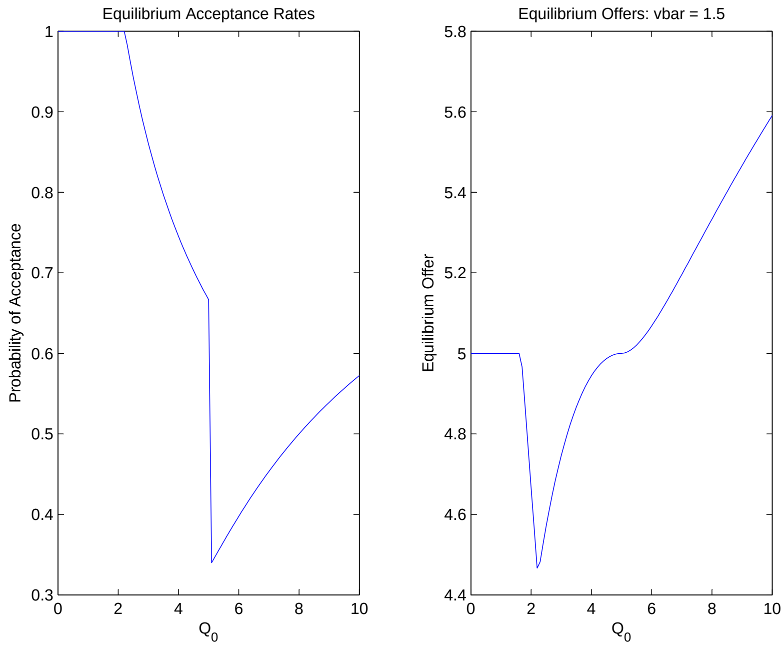

discontinuity. Consider now unbalanced games. One can see that there are two types of unbalanced games depending on whether $P^+ > P^-$ or $P^+ < P^-$. In the former case, we have that very close to $Q_0 = \frac{A}{2}$, it is easier to agree to reduce the quota, while in the latter case it is easier to agree to increase the quota. The intuition for this is also very easy. Consider the case in which $P^+ < P^-$ so that the mass of responder types with $b < 1$ is greater than $\frac{1}{2}$. Therefore, more than half of the responder types prefer a quota *strictly* larger than $\frac{A}{2}$, making it difficult to agree on the way down and leading to a downward discontinuity in the equilibrium acceptance probability function.

Example 1 (Uniform Distribution). The equilibrium acceptance probability, $P(A, Q^*(A, Q_0), Q_0)$, and the equilibrium offer function, $Q^*(A, Q_0)$, for the case in which F is uniform on $[0, 1.5]$ are on display in Figure 4. Note that this is an example of an unbalanced game with $P^+ < P^-$, meaning that disagreement is substantially more likely when reducing the quota.[25]

5.2 The Dynamic Game: Myopic Responder

Here we extend the basic model two allow for two things: First we introduce demand fluctuations in the form of a two state Markov process. Second, we assume that the proposer is far-sighted, in that she maximises the expected discounted sum of payoffs, with discount factor $\delta \in [0, 1)$. More specifically, suppose that demand can either be high, A_h, or low, A_l and that there is a Markov

[25]If F were exponential with parameter $\lambda = \ln 2$, the game would be balanced since the median of the exponential is 1 when $\lambda = \ln 2$.

transition matrix:

$$\Pi = \begin{bmatrix} 1-\mu & \mu \\ \gamma & 1-\gamma \end{bmatrix} \tag{7}$$

We assume that the responder draws, *i.i.d.* over time, a new realisation of b. When useful, we will denote p_{ij} to be the transition probability from state i to state j. For example, $p_{lh} = \mu$.

Under the assumption that the responder is myopic, she will accept or reject according to (3), meaning that (4) describes the relevant acceptance probability for the proposer. The proposer's value function, $V(A, Q_0)$ is defined by the following:

$$\begin{aligned} V(A_i, Q_0) &= \max_{q\in[0,A_l]} Q_0(A_i - Q_0) + P(A_i, q, Q_0)\left[q(A_i - q) - Q_0(A_i - Q_0)\right] \\ &+ \delta P(A_i, q, Q_0)\left[p_{ii}(V(A_i, q) - V(A_i, Q_0)) + p_{ij}(V(A_j, q) - V(A_j, Q_0))\right] \\ &+ \delta\left[p_{ii}V(A_i, Q_0) + p_{ij}V(A_j, Q_0)\right] \end{aligned}$$

for $i = l, h$ and $j \neq i$. We make the assumption that $A_l < A_h < 2A_l$ so that we can assume that the quota, in the high demand state, will never be above A_l; *i.e.,* we are assuming that demand shocks are not too large. Upon noting that the reward function is continuous in Q, standard arguments imply that a solution to the dynamic programming problem exists.[26]

Given that a solution to the DP problem exists, we now wish to characterise its properties. In fact, it is not at all difficult to show that a similar result to that derived in the stage game also holds here. In both the high demand and low demand states, there will generally exist numbers Q_{0h} and Q_{0l} such that it is optimal to propose the status quo at these points, while if Q_0 is larger, the proposal is to decrease the quota and if Q_0 is smaller, the proposal is to increase the quota. The only difference is that Q_{0j} will generally depend upon δ, μ and γ and so need not equal $\frac{A_j}{2}$. Intuitively, it must be that $Q_{0l} \geq \frac{A_l}{2}$ and $Q_{0h} \leq \frac{A_h}{2}$. In a world with no demand fluctuations the proposer would eventually move the quota to $\frac{A}{2}$ and indeed, from a static perspective, she would like to do so even in a world of demand fluctuations. However, since the proposer is forward looking, she will optimally choose $Q_{0l} \geq \frac{A_l}{2}$, since she recognises that demand is likely to increase in the future and it will be relatively difficult to increase the quota around $\frac{A_l}{2}$; therefore, she sacrifices some static payoffs in the low state, so that she may experience higher profits upon a transition to the high demand state. A similar intuition holds for why $Q_{0h} < \frac{A_h}{2}$. Therefore, we see that it will generally be *most* difficult to lower the quota when demand in high.

Remark 5. Therefore, we see that in the dynamic setting considered here, balancedness is generally endogenous. The possibility of transitions to a new demand state force the proposer to deviate from what would be here statically optimal behaviour. Therefore, the discontinuities in the acceptance probability function depend upon the underlying parameters of the model. Figure 5 shows the optimal offers and the equilibrium acceptance probability functions for the uniform $[0, 2]$ distribution

[26]Define the appropriate mapping $\Phi : \mathbb{C}(\{A_l, A_h\} \times [0, A_l]) \mapsto \mathbb{R}^{\{A_l, A_h\}\times[0,A_l]}$ and use the Maximum Theorem to show that it is a self-map. Finally, it is apparent that the conditions of Blackwell's Contraction Lemma are satisfied, implying that Φ is a contraction and a solution to the DP exists.

Figure 5: Equilibrium Behaviour: Myopic Responder

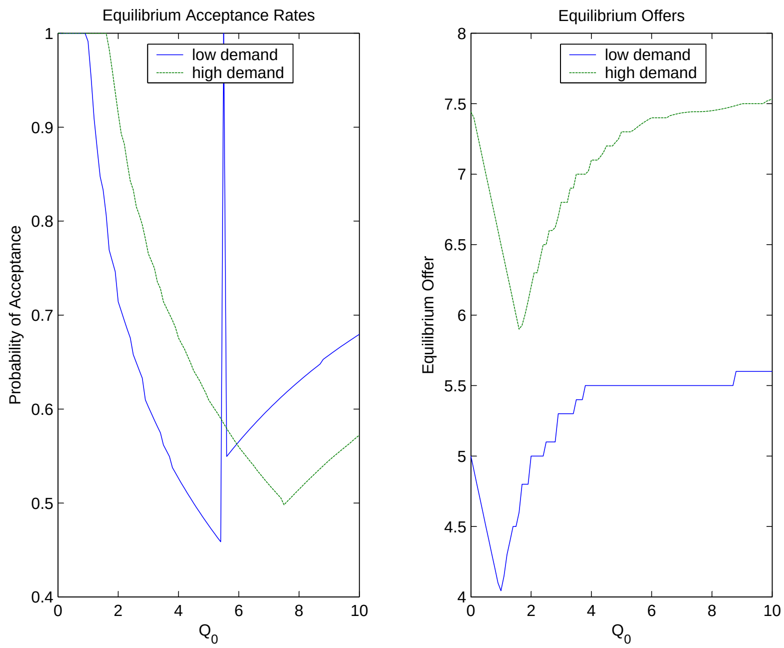

with relatively frequent transitions from low to high demand, $\mu = \frac{1}{2}$. While in the static version, this game would be balanced, the aforementioned intuition is at play here and we see that discontinuities arise. In particular, because of the relatively frequent transitions from low to high demand, the proposer prefers a quota $Q > \frac{A_l}{2} = 5$, leading to the upward discontinuity, even though, in a static sense, the model is balanced. Indeed, here for a given range of initial quotas, it may actually be more difficult to increase the quota.[27]

5.3 Discussion

In the introduction, we argued against taking the standard repeated game approach in the current context. There, we argued, that Athey *et al* and other repeated game models fail because they *necessarily* rely on changes in relative quota shares. However, as shown in Figure 1, shares have been remarkably constant over time. In contrast, we feel that the model of Section 5 does a good job of explaining the main empirical findings reported in Section 3; however, there may be other explanations that are also capable of rationalising the data. We examine two such possibilities. The first that immediately comes to mind is based on aspirations and entitlements. Suppose that OPEC members come to view their allocated quota as an entitlement, with the implication being that they would *always* agree to expand their entitlements; however, if times turn sour and

[27] We solved the dynamic programming problem through a value iteration process with a grid size of 0.1 and a given convergence criterion: $\sum_j \left[|V^{t+1}(A_h, 0.1j) - V^t(A_h, 0.1j)| + |V^{t+1}(A_l, 0.1j) - V^t(A_l, 0.1j)|\right] < 0.0025$.

entitlements must be taken away, they find it much more difficult to agree.[28] In this way, there is a *ratchet effect* at work which makes the quota easy to increase but difficult to decrease. While intuitively appealing, this explanation fails on a two grounds. First, if agents are far-sighted, there will likely be disagreements on the way up: even in a good state, agents may not allow the quota to increase, because in doing so, entitlements would increase, leading to disagreement when demand later decreases. Second, the entitlements *must* be on each member's quota. If, instead, entitlements were based on profits or relative shares, agreements would still be possible in both good and bad states; in the latter case, they may simply adjust the aggregate quota, leaving relative shares unchanged.

Returning to the rational paradigm, in some situations, a model of adjustment costs may create the necessary status quo bias, leading to an increased likelihood of disagreement when demand decreases. Under this view, which also takes a bargaining approach, upon the realisation of demand, the proposer offers a new quota to a responder who faces a random cost of adjustment. If this adjustment cost is not too large, the responder accepts, while if she experiences a very large adjustment cost, she rejects and the status quo obtains. For this model to replicate our empirical findings, one needs to assume that the gains from agreeing to the proposed quota are *larger* following an increase in demand than a decrease. In this way, when demand increases, the gains to reaching an agreement are large and so, for a larger set of realised adjustment costs, an agreement can actually be reached, while when demand decreases, the gains are smaller, leading to a higher probability of disagreement. To be sure, such an explanation has merit; however, the model of random shocks presented above is capable of generating the stark predictions that our empirical findings would suggest. In contrast, for many demand and cost functions, the gains from agreeing following a demand increase are not substantially larger than the gains from agreeing following a negative demand shock — if, in fact, the gains are even larger at all. Therefore, we would not expect the probability of agreement to differ much depending upon whether demand has increased or decreased.

However, where we believe that our explanation truly succeeds is its ability to make accurate predictions beyond the relative likelihood of agreement given the state of nature in a way that the aforementioned models cannot. In a simple (yet more realistic) extension of the basic model, one might imagine that, following a proposal to reduce the quota, there are some responder types who agree, but only very grudgingly or who, despite a contrary vote, have a new quota imposed on them by the agreement of some super-majority of the remaining cartel members. In these cases, it seems likely that such agents will be substantially more likely to cheat on their (now lower) quota. Therefore, this extended model also predicts a negative relationship between cheating and the size of one's quota. This is something that we can empirically test. Define $C_t^i = \frac{P_t^i - Q_t^i}{Q_t^i}$ to be the fraction of production over quota by country i at time t. For each country, we estimate the

[28] Indeed, in an experimental setting, Gächter and Riedl (2003) show that such entitlements can lead to a much different set of outcomes when these entitlements are infeasible.

Figure 6: Predicted Relationship Between Cheating and the Quota

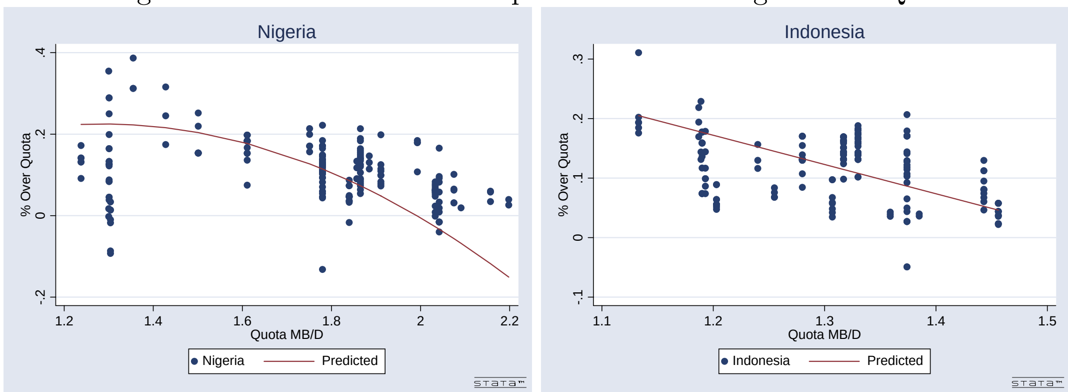

following equation:

$$C_t^i = \alpha + \beta_1 Q_t^i + \beta_2 Q_t^i D_t + \beta_3 D_t + \beta_4 NO_t + \beta_5 OIL_t + X_t\gamma + \epsilon_t^i \tag{8}$$

where D_t is crude oil consumption by OECD countries, NO_t is crude oil production by non-OPEC countries, OIL_t is the two-month forward price of oil and X_t is vector of other control variables, including quarter dummies, a dummy for the Person Gulf War, which takes value 1 from August 1990 to March 1991 and zero outside this interval, and a time index. The estimated results for each country are provided in Table 9.

The effect on cheating for a change in the quota is then given by $\frac{\partial C_t^i}{\partial Q_t^i} = \beta_1 + \beta_2 D_t$. If we evaluate this expression at the average level of OECD consumption in our sample, one can see that the result is negative for all countries and significant for 7 of 10 countries, with Nigeria, Qatar and the United Arab Emirates being the three exceptions.[29] However, if we allow for a potentially non-linear relationship between cheating and the quota by adding a quadratic term, we find a significant relationship even for Nigeria and Qatar. One can see this predicted relationship for the cases of Nigeria and Indonesia in Figure 6. The solid line is the predicted relationship between cheating and the quota, evaluating all other variables at their means, while the dots are the actual data. Thus, our model explains not only the main empirical finding that agreements are more likely during good times, but also accurately predicts a negative relationship between the size of the quota (controlling for other demand and supply variables). The other models that we have discussed either fail as an explanation on some other fundamental ground (*e.g.,* Athey *et al* and aspirations) or lack the predictive power to explain observed patterns of cheating (*e.g.,* adjustment costs).

Remark 6. Note that the negative relationship between cheating (as measured by the fraction of production over the quota) is not simply because all OPEC members produce X barrels more than their quota, regardless. The absolute amount of cheating also varies with the size of the quota.

[29]The average level of OECD crude oil consumption was 41.6 million barrels per day.

6 Concluding Remarks

In this paper we have presented the results of an event study analysis of OPEC behaviour and its effects on the stock returns in the energy sector. The results are quite clear. When OPEC reduces the quota, positive and often significant abnormal returns accrue in the three sub-sectors considered. Next, when the quota is increased, there is no significant pattern of abnormal returns. Finally, when OPEC takes no action negative and significant abnormal return accrue across all sectors. Moreover, these abnormal returns appear to be economically meaningful, accumulating in some cases to $\pm 5\%$ by the end of the event window.

We believe strongly that these results offer us an interesting insight into OPEC's decision-making behaviour. Specifically, the results indicate an asymmetric ability of OPEC to secure agreements. That is, when demand is increasing, OPEC is far more likely to increase the quota than they are to decrease it when demand is decreasing. In part, this interpretation is confirmed when one examines the news reports surrounding these status quo announcements. Robustness checks reported throughout the paper provide further support for this interpretation.

In an attempt to explain our findings, it is clear that demand uncertainty is not enough. Often in repeated games with only demand uncertainty, there is a precise mapping between the state of demand and the maximal level of collusion sustainable. Therefore, a market forecasting future behaviour can never be surprised. The few extensions of Rotemberg and Saloner (1986) such as Staiger and Wolak (1992) and Athey, Bagwell and Sanchirico (2004), which can generate surprise often do so in an uncoordinated or asymmetric way and predict fluctuating market shares. However, as Figure 1 showed, quota shares have been remarkably constant over time. It is this constancy of shares which takes us away from repeated game models of collusion to the bargaining model considered above.

As we briefly mentioned in section 4.3, the conventional wisdom in much of the popular press is that there has been a dramatic shift in OPEC behaviour since the Asian Financial Crisis. Just as many Asian countries entered into a state of crisis, OPEC actually increased production, further exacerbating the decline in prices due to the adverse financial shocks experienced. Many have since argued that oil prices near $10 per barrel brought a truce amongst long-standing rivals Saudi Arabia and Iran. Thus, post-Financial Crisis, OPEC has been much better at counteracting downturns in oil prices; indeed, as *The Economist* pointed out OPEC, and specifically Saudi Arabia, "worked out that watching OECD countries' stocks, and slashing OPEC output if they rose, was a superb way of propping up the oil price."[30]

However, at present, with oil prices still at historically high levels, OPEC's ability to influence prices is seen as being greatly diminished for a variety of reasons. First, strong growth in oil demand from the United States, China and now India means that many OPEC countries are already operating at near capacity. Given that what little excess capacity remains is not evenly spread amongst OPEC members, it has been argued that this is a current source of disagreement

[30] See the article, "The Central Bank Takes Stock; OPEC." *The Economist,* 19 March 2005.

— to have any meaningful impact on oil prices, any quota increase must involve a reallocation of quota shares, something which we have argued OPEC members are loath to do (See footnote 6.). It seems possible to provide an explanation for this phenomenon with our model: at present the support of the distribution of types has shifted right as production approaches capacity. Such a rightward shift implies that disagreements to increase the quota are more likely, since the interests of smaller OPEC countries and Saudi Arabia (the country with the most spare capacity) would now seem to diverge. Second, and related to the first, western oil refineries are also operating at nearly full capacity, making it unlikely that they could accommodate increased OPEC supply if it were to materialise. Third, concerns persist about the stability of both Iraq, as it faces a continued insurgency, and Saudi Arabia, which has seen relations with western nations strained recently. Fourth, political instability in Russia, Nigeria and Venezuela, all large exporters, has led to decreased and unstable production, further adding upward pressure on oil prices.

In the static version of our model, we easily obtained that disagreements were much more likely to occur when the proposer wishes to lower the quota. In moving to a more dynamic version in which the proposer cares about current and future payoffs (though the responder is still myopic), many of the results continue to go through. However, the results now depend on the underlying parameters of the model — in particular, the Markov transition matrix. For example, as Figure 5 depicts, when transitions from the low to high demand state are relatively more frequent, disagreements may be more likely when raising the quota. We have not yet extended the model to allow for far-sighted responder types. We conjecture that in a stationary equilibrium, the qualitative results will closely resemble those obtained in the static analysis. Where as a myopic responder may be willing to lower the quota, a far-sighted responder will realise that demand may increase in the future and, therefore, resist the proposer's efforts to lower the quota. However, we leave this exercise for future work.

References

[1] Alnasrawi, Abbas, *OPEC in a Changing World Economy,* The Johns Hopkins University Press, 1984.

[2] Athey, Susan, Kyle Bagwell and Chris Sanchirico, *Collusion and Price Rigidity.* Review of Economic Studies 71 (2), April 2004.

[3] Bagwell, Kyle and Robert W. Staiger. *Collusion Over The Business Cycle.* RAND Journal of Economics Vol. 28, No. 1 (1997), 82-106.

[4] Campbell, John Y., Andrew W. Lo and A Craig MacKinley, *The Econometrics of Financial Markets* (Princeton University Press, 1997)

[5] Card, David. *Strikes and Wages: A Test of an Asymmetric Information Model.* The Quarterly Journal of Economics, Vol. 105, No. 3. (Aug., 1990), pp. 625-659.

[6] Card, David., *Longitudinal Analysis of Strike Activity.* Journal of Labor Economics, Vol. 6, No. 2. (Apr., 1988), pp. 147-176.

[7] Cyree, Ken B and Ramon P. Degennaro., *A Generalized Method for Detecting Abnormal Returns and Changes in Systematic Risk.* Review of Quantitative Finance and Accounting, 19:399-416, 2002.

[8] Gächter, Simon and Arno Riedl., *Moral Property Rights In Bargaining With Infeasible Claims.* Tinbergen Institute Discussion Paper, No 03-055/1.

[9] Griffin, James M., *OPEC Behavior: A Test of Alternative Hypotheses.* American Economic Review 75 (Dec 1985)

[10] Griffin, James M. and Weiwen Xiong., *The Incentive To Cheat: An Empirical Analysis of OPEC.* Journal of Law and Economics 40 (Oct. 1997), pp. 289-316.

[11] Guidolin, Massimo and Eliana La Ferrara., *Diamonds Are Forever, Wars Are Not. Is Conflict Bad For Private Firms?* CEPR Discussion Paper 4668, 2004.

[12] Haltiwanger, John and Joseph E. Harrington Jr., *The Impact Of Cyclical Demand Movements on Collusive Behaviour.* RAND Journal of Economics 22 (1991)

[13] Harrison, Alan and Mark Stewart., *Cyclical Fluctuations in Strike Durations.* The American Economic Review, Vol. 79, No. 4. (Sep., 1989), pp. 827-841.

[14] Hart, Oliver., *Bargaining and Strikes.* The Quarterly Journal of Economics, Vol. 104, No. 1 (Feb. 1989), pp. 25-43.

[15] Hotelling, Harold., *The Economics of Exhaustible Resources.* Journal of Political Economy 39 (April 1931) pp. 137-175.

[16] International Energy Agency, *Energy Price Volatility: Trends and Consequences.* (2002)

[17] Kaufmann, Robert, Stephane Dees, Pavlos Karadeloglou, and Marcelo Sánchez (2004) *Does OPEC Matter? An Econometric Analysis of Oil Prices.* The Energy Journal 25(4):67-90

[18] Kennan, John., *The Duration of Contract Strikes in U.S. Manufacturing.* Journal of Econometrics, April 1985, 28: 5-28.

[19] Kennan, John and Robert Wilson., *Strategic Bargaining Models and Interpretation of Strike Data.* Journal of Applied Econometrics, Vol. 4, Supplement: Special Issue on Topics in Applied Econometrics. (Dec., 1989), pp. S87-S130.

[20] Koyama, Ken, *Crude Oil Prices After the OPEC Meeting.* The Institute of Energy Economics, Japan. 12 April 2000.

[21] Levenstein, Margaret and Valerie Suslow, *What Determines Cartel Success?"* Working Paper (2002).

[22] Lewis, Matt., *Asymmetric Price Adjustment and Consumer Search: An Examination of the Retail Gasoline Market.* University of California, Berkeley. (2003)

[23] Lockwood, L and K. Rao Kadiyala *Risk Measurement of Event Dependent Security Returns.* Journal of Business and Economic Statistics, 6, 1, January 1988

[24] Mabro, Robert, *Mabro On The Futures Markets And The Price Of Oil.* Middle East Economic Survey XLVII(No. 22) (2004).

[25] Mabro, Robert, *Does Oil Price Volatility Matter?* Oxford Institution for Energy Studies Monthly Comment, June 2001.

[26] Peltzman, Sam., *Prices Rise Faster Than They Fall.* Journal of Political Economy 108(3) (2000), pp. 466-502.

[27] Ramcharran, Harri, *Oil production reponses to price changes: an empirical application of the competitive model to OPEC and non-OPEC countries.* Energy Economics 24 (2002)

[28] Reynolds, Stanley and Bart Wilson. *Market Power and Price Movements Over The Business Cycle*. Journal of Industrial Economics 53(2) (2005), pp. 145-174.

[29] Romer, Paul M., *Preferences, Promises, and the Politics of Entitlement,* in Victor Fuchs (ed.): "Individual and Social Responsibility - Child Care, Education, Medical Care, and Long-Term Care in America." Chicago: Chicago University Press 1996.

[30] Rotemberg, Julio J. and Garth Saloner. *A Supergame-Theoretic Model of Price Wars During Booms.* American Economic Review 76 (1986)

[31] Scherer, F. M. and David Ross, *Industrial Market Structure and Economic Performance.* 3rd Edition (Houghton Mifflin Company, Boston) 1990.

[32] Staiger, Robert and Frank Wolak, *Collusive Pricing With Capacity Constraints in the Presence of Demand Uncertainty.* RAND Journal of Economics 23(2) (1992) pp. 203-220.

A Figures

Figure 7: Cumulative Abnormal Return Functions

B Extraordinary Meetings

One may argue that the there is a potential problem because, as mentioned in the main body of the text, OPEC sometimes calls extraordinary meetings, beyond those which are regularly scheduled. It is plausible that if an extraordinary meeting is called, the market, strongly expects OPEC to implement some change in policy, possibly affecting our empirical results.

As was mentioned in the main body of the text, there is a possible endogeneity regarding the timing of OPEC meetings. As was mentioned in the main body of the text, there were nine extraordinary meetings, with three fitting into each category. We replicate our event study, excluding all extraordinary meetings.

The patterns are largely identical to those reported earlier in which all meetings were included in the analysis. In particular, status quo announcements are still met with negative abnormal returns of approximately -2 to -3%, while quota reductions are met with positive abnormal returns of between 4 and 7%. In contrast, for the oil industry as a whole and for exploration and drilling companies, there appears to be very little impact for quota increases. However, there does seem to be a positive impact of quota increases on integrated companies.

The most striking difference between the current analysis and that of the main body is that there appears to be some evidence that information leaks out before the announcement. In particular, look at the period 9 days before the announcement and specifically look at the CAR function for reduction and status quo announcements. In each case, there is a significantly positive abnormal return for quota reductions and a negative (though generally not significant) abnormal return for status quo announcements. This indicates to us that approximately 9 days before the event, at least some of the uncertainty surrounding whether or not there will be a reduction is resolved before the actual announcement. The remaining uncertainty is then resolved in the days immediately surrounding the announcement. Note also that for integrated companies there appears to be a positive impact (on the announcement day) of an announcement to increase the quota. This does appear to be somewhat at odds with our previous analysis and interpretation, which says that increases are largely predicted — and hence, there should not be significant abnormal returns. However, in defence of our interpretation, it should be noted that the positive abnormal return on the announcement day and, indeed, over the entire event window is much smaller than those which occur for status quo and reduction announcements. It is, therefore, plausible to argue that any surprise arises due to the absolute size of the change in the quota, rather than whether the quota would be raised at all. Moreover, if there was true uncertainty over whether the quota would actually be raised, we would expect to see significant abnormal returns at the same time for status quo announcements, something which we do not see. Finally, if one looks at the CAR functions for the two-month forward price of oil, the significant announcement-day effect for quota increases is gone — providing even more evidence that quota increases are predicted. Thus, we believe that our main empirical findings continue to hold even when we exclude extraordinary meetings.

Figure 8:

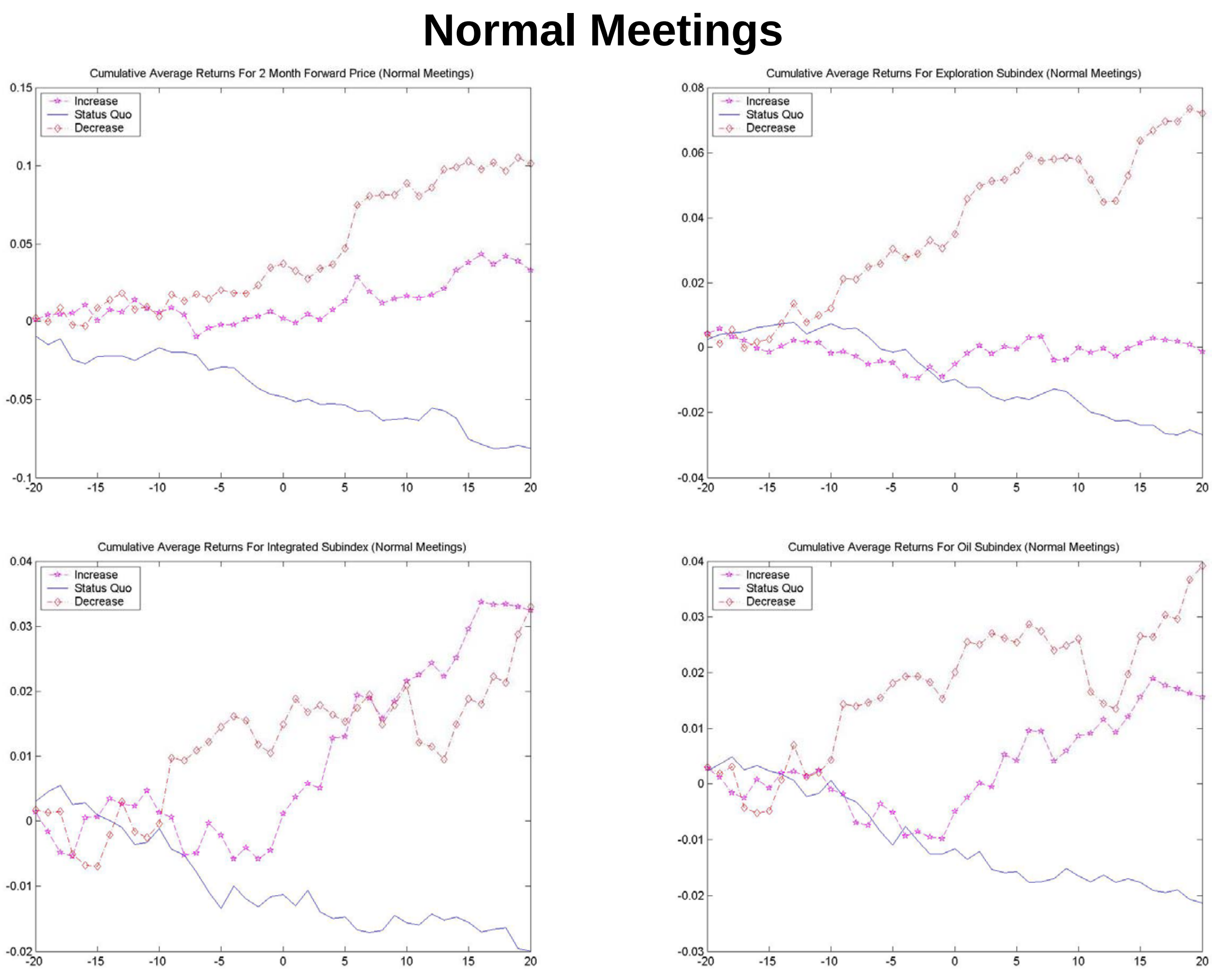
Normal Meetings
Cumulative Average Returns For 2 Month Forward Price (Normal Meetings)
Increase
Status Quo
Decrease
Cumulative Average Returns For Exploration Subindex (Normal Meetings)
Increase
Status Quo
Decrease
Cumulative Average Returns For Integrated Subindex (Normal Meetings)
Increase
Status Quo
Decrease
Cumulative Average Returns For Oil Subindex (Normal Meetings)
Increase
Status Quo
Decrease

Table 3: **Cumulative Abnormal Returns For Oil Companies Sub-Index**

Event Day	Q↑ $\overline{CAR}$	Q↑ s.e.	Q↔ $\overline{CAR}$	Q↔ s.e.	Q↓ $\overline{CAR}$	Q↓ s.e.
-20	0.0018	0.0024	0.0021	0.0020	0.0007	0.0034
-19	-0.0002	0.0035	0.0021	0.0028	-0.0014	0.0048
-18	-0.0025	0.0043	0.0036	0.0034	0.0012	0.0058
-17	-0.0034	0.0049	0.0018	0.0039	-0.0026	0.0068
-16	-0.0012	0.0055	0.0014	0.0044	-0.0023	0.0076
-15	-0.0020	0.0061	0.0008	0.0049	-0.0040	0.0083
-14	-0.0009	0.0066	0.0001	0.0053	0.0027	0.0090
-13	-0.0036	0.0071	-0.0016	0.0057	0.0082	0.0097
-12	-0.0052	0.0075	-0.0044	0.0060	0.0039	0.0103
-11	-0.0019	0.0080	-0.0047	0.0064	0.0041	0.0109
-10	-0.0053	0.0084	-0.0021	0.0067	0.0018	0.0115
-9	-0.0062	0.0088	-0.0047	0.0070	0.0061	0.0120
-8	-0.0119	0.0092	-0.0061	0.0073	0.0068	0.0125
-7	-0.0109	0.0095	-0.0065	0.0076	0.0064	0.0130
-6	-0.0077	0.0099	-0.0095	0.0079	0.0076	0.0135
-5	-0.0090	0.0103	-0.0112	0.0082	0.0122	0.0140
-4	-0.0139	0.0106	-0.0086	0.0085	0.0134	0.0145
-3	-0.0111	0.0110	-0.0117	0.0088	0.0153	0.0150
-2	-0.0128	0.0113	-0.0122	0.0090	0.0149	0.0154
-1	-0.0131	0.0116	-0.0122	0.0093	0.0129	0.0159
0	-0.0061	0.0119	-0.0114	0.0095	0.0105	0.0163
1	-0.0049	0.0122	-0.0134	0.0098	0.0070	0.0168
2	-0.0017	0.0125	-0.0115	0.0100	0.0088	0.0172
3	-0.0033	0.0128	-0.0143	0.0103	0.0104	0.0176
4	0.0014	0.0131	-0.0160	0.0105	0.0133	0.0180
5	-0.0004	0.0134	-0.0162	0.0108	0.0122	0.0184
6	0.0041	0.0137	-0.0183	0.0110	0.0170	0.0188
7	0.0024	0.0140	-0.0192	0.0112	0.0150	0.0192
8	-0.0019	0.0143	-0.0187	0.0115	0.0097	0.0196
9	-0.0002	0.0146	-0.0180	0.0117	0.0116	0.0200
10	-0.0008	0.0149	-0.0205	0.0119	0.0134	0.0204
11	-0.0001	0.0152	-0.0233	0.0121	0.0061	0.0208
12	0.0016	0.0154	-0.0212	0.0124	0.0062	0.0212
13	-0.0007	0.0157	-0.0228	0.0126	0.0102	0.0216
14	0.0039	0.0160	-0.0223	0.0128	0.0163	0.0219
15	0.0054	0.0163	-0.0238	0.0130	0.0227	0.0223
16	0.0075	0.0165	-0.0235	0.0133	0.0209	0.0227
17	0.0059	0.0168	-0.0249	0.0135	0.0248	0.0230
18	0.0050	0.0171	-0.0252	0.0137	0.0242	0.0234
19	0.0047	0.0173	-0.0263	0.0139	0.0281	0.0237
20	0.0044	0.0176	-0.0255	0.0141	0.0318	0.0241

Table 4: **Cumulative Abnormal Returns For Exploration Companies Sub-Index**

Event Day	Q↑ $\overline{CAR}$	Q↑ s.e.	Q↔ $\overline{CAR}$	Q↔ s.e.	Q↓ $\overline{CAR}$	Q↓ s.e.
-20	0.0034	0.0024	0.0016	0.0023	0.0020	0.0037
-19	0.0047	0.0033	0.0016	0.0033	-0.0014	0.0053
-18	0.0021	0.0041	0.0022	0.0041	0.0039	0.0065
-17	0.0012	0.0048	0.0030	0.0047	0.0006	0.0075
-16	-0.0024	0.0053	0.0033	0.0053	0.0035	0.0084
-15	-0.0031	0.0059	0.0037	0.0058	0.0055	0.0092
-14	-0.0015	0.0064	0.0035	0.0063	0.0123	0.0100
-13	-0.0021	0.0068	0.0030	0.0068	0.0185	0.0107
-12	-0.0037	0.0073	-0.0015	0.0072	0.0150	0.0114
-11	-0.0033	0.0077	-0.0010	0.0076	0.0142	0.0120
-10	-0.0058	0.0081	0.0010	0.0080	0.0133	0.0127
-9	-0.0051	0.0085	-0.0011	0.0084	0.0182	0.0133
-8	-0.0071	0.0088	-0.0015	0.0088	0.0192	0.0138
-7	-0.0086	0.0092	-0.0020	0.0091	0.0219	0.0144
-6	-0.0078	0.0095	-0.0055	0.0095	0.0226	0.0150
-5	-0.0087	0.0099	-0.0058	0.0098	0.0289	0.0155
-4	-0.0139	0.0102	-0.0055	0.0102	0.0258	0.0160
-3	-0.0126	0.0105	-0.0101	0.0105	0.0279	0.0166
-2	-0.0109	0.0109	-0.0105	0.0108	0.0311	0.0171
-1	-0.0133	0.0112	-0.0133	0.0111	0.0307	0.0176
0	-0.0068	0.0115	-0.0119	0.0114	0.0268	0.0180
1	-0.0041	0.0118	-0.0145	0.0117	0.0258	0.0185
2	-0.0011	0.0121	-0.0135	0.0120	0.0305	0.0190
3	-0.0036	0.0124	-0.0162	0.0123	0.0333	0.0195
4	-0.0022	0.0127	-0.0201	0.0126	0.0372	0.0199
5	-0.0036	0.0129	-0.0192	0.0129	0.0381	0.0204
6	-0.0012	0.0132	-0.0205	0.0132	0.0451	0.0208
7	-0.0021	0.0135	-0.0200	0.0134	0.0438	0.0212
8	-0.0075	0.0138	-0.0188	0.0137	0.0416	0.0217
9	-0.0076	0.0141	-0.0211	0.0140	0.0420	0.0221
10	-0.0079	0.0143	-0.0251	0.0142	0.0409	0.0225
11	-0.0095	0.0146	-0.0303	0.0145	0.0372	0.0230
12	-0.0094	0.0149	-0.0303	0.0148	0.0353	0.0234
13	-0.0123	0.0151	-0.0318	0.0150	0.0381	0.0238
14	-0.0098	0.0154	-0.0314	0.0153	0.0453	0.0242
15	-0.0101	0.0157	-0.0335	0.0156	0.0553	0.0246
16	-0.0100	0.0159	-0.0318	0.0158	0.0554	0.0250
17	-0.0101	0.0162	-0.0344	0.0161	0.0591	0.0254
18	-0.0117	0.0164	-0.0349	0.0164	0.0584	0.0258
19	-0.0131	0.0167	-0.0334	0.0166	0.0608	0.0262
20	-0.0147	0.0169	-0.0340	0.0169	0.0592	0.0266

Table 5: **Cumulative Abnormal Returns For Integrated Companies Sub-Index**

Event Day	Q↑ $\overline{CAR}$	Q↑ s.e.	Q↔ $\overline{CAR}$	Q↔ s.e.	Q↓ $\overline{CAR}$	Q↓ s.e.
-20	0.0009	0.0027	0.0030	0.0019	-0.0008	0.0032
-19	-0.0028	0.0038	0.0033	0.0027	-0.0026	0.0046
-18	-0.0047	0.0046	0.0046	0.0034	-0.0006	0.0056
-17	-0.0052	0.0054	0.0025	0.0039	-0.0039	0.0065
-16	-0.0005	0.0060	0.0016	0.0044	-0.0051	0.0073
-15	-0.0001	0.0066	0.0008	0.0048	-0.0075	0.0080
-14	0.0010	0.0072	-0.0001	0.0052	-0.0027	0.0087
-13	-0.0023	0.0077	-0.0018	0.0056	0.0018	0.0093
-12	-0.0027	0.0082	-0.0039	0.0059	-0.0017	0.0099
-11	0.0019	0.0086	-0.0042	0.0063	-0.0018	0.0105
-10	-0.0017	0.0091	-0.0019	0.0066	-0.0039	0.0110
-9	-0.0031	0.0095	-0.0045	0.0069	0.0005	0.0116
-8	-0.0091	0.0099	-0.0055	0.0072	0.0009	0.0121
-7	-0.0076	0.0103	-0.0064	0.0075	0.0009	0.0126
-6	-0.0037	0.0107	-0.0091	0.0078	0.0021	0.0130
-5	-0.0048	0.0111	-0.0115	0.0081	0.0045	0.0135
-4	-0.0089	0.0115	-0.0084	0.0083	0.0063	0.0140
-3	-0.0055	0.0119	-0.0108	0.0086	0.0074	0.0144
-2	-0.0075	0.0122	-0.0104	0.0089	0.0049	0.0149
-1	-0.0064	0.0126	-0.0086	0.0091	0.0034	0.0153
0	0.0004	0.0129	-0.0086	0.0094	0.0024	0.0157
1	0.0016	0.0133	-0.0100	0.0096	-0.0002	0.0161
2	0.0044	0.0136	-0.0072	0.0099	-0.0002	0.0166
3	0.0028	0.0139	-0.0097	0.0101	0.0005	0.0170
4	0.0088	0.0143	-0.0109	0.0104	0.0025	0.0174
5	0.0087	0.0146	-0.0115	0.0106	0.0008	0.0178
6	0.0138	0.0149	-0.0138	0.0108	0.0041	0.0181
7	0.0114	0.0152	-0.0144	0.0110	0.0038	0.0185
8	0.0091	0.0155	-0.0141	0.0113	-0.0017	0.0189
9	0.0118	0.0158	-0.0124	0.0115	0.0016	0.0193
10	0.0120	0.0161	-0.0147	0.0117	0.0053	0.0196
11	0.0134	0.0164	-0.0168	0.0119	-0.0019	0.0200
12	0.0150	0.0167	-0.0142	0.0122	-0.0019	0.0204
13	0.0132	0.0170	-0.0159	0.0124	0.0021	0.0208
14	0.0176	0.0173	-0.0160	0.0126	0.0070	0.0211
15	0.0199	0.0176	-0.0177	0.0128	0.0106	0.0215
16	0.0229	0.0179	-0.0180	0.0130	0.0081	0.0218
17	0.0218	0.0182	-0.0186	0.0132	0.0131	0.0222
18	0.0222	0.0185	-0.0196	0.0135	0.0135	0.0225
19	0.0219	0.0188	-0.0219	0.0137	0.0175	0.0228
20	0.0218	0.0191	-0.0207	0.0139	0.0231	0.0232

Table 6: **Abnormal Returns For Oil Companies Sub-Index**

Event Day	Q↑ $\overline{\epsilon}^*$	Q↑ s.e.	Q↔ $\overline{\epsilon}^*$	Q↔ s.e.	Q↓ $\overline{\epsilon}^*$	Q↓ s.e.
-20	0.0018	0.0024	0.0021	0.0020	0.0007	0.0034
-19	-0.0021	0.0024	-0.0001	0.0020	-0.0021	0.0033
-18	-0.0023	0.0024	0.0015	0.0020	0.0026	0.0034
-17	-0.0008	0.0024	-0.0017	0.0020	-0.0038	0.0034
-16	0.0022	0.0025	-0.0005	0.0020	0.0003	0.0033
-15	-0.0008	0.0024	-0.0006	0.0020	-0.0017	0.0033
-14	0.0011	0.0024	-0.0007	0.0020	0.0067	0.0034
-13	-0.0027	0.0024	-0.0016	0.0020	0.0055	0.0034
-12	-0.0016	0.0024	-0.0029	0.0020	-0.0043	0.0034
-11	0.0033	0.0025	-0.0003	0.0020	0.0002	0.0034
-10	-0.0034	0.0024	0.0027	0.0020	-0.0024	0.0034
-9	-0.0009	0.0025	-0.0026	0.0020	0.0043	0.0034
-8	-0.0058	0.0024	-0.0014	0.0020	0.0007	0.0034
-7	0.0010	0.0025	-0.0004	0.0020	-0.0003	0.0034
-6	0.0032	0.0024	-0.0030	0.0020	0.0011	0.0033
-5	-0.0013	0.0024	-0.0017	0.0020	0.0046	0.0034
-4	-0.0049	0.0024	0.0026	0.0020	0.0012	0.0034
-3	0.0028	0.0024	-0.0031	0.0020	0.0019	0.0033
-2	-0.0017	0.0024	-0.0005	0.0020	-0.0004	0.0034
-1	-0.0003	0.0024	0.0000	0.0020	-0.0020	0.0034
0	0.0069	0.0024	0.0008	0.0020	-0.0024	0.0034
1	0.0013	0.0024	-0.0020	0.0020	-0.0035	0.0034
2	0.0032	0.0024	0.0018	0.0020	0.0018	0.0034
3	-0.0016	0.0024	-0.0028	0.0020	0.0016	0.0033
4	0.0047	0.0024	-0.0017	0.0020	0.0029	0.0033
5	-0.0018	0.0024	-0.0001	0.0020	-0.0011	0.0034
6	0.0046	0.0024	-0.0022	0.0020	0.0048	0.0033
7	-0.0018	0.0024	-0.0008	0.0020	-0.0019	0.0034
8	-0.0042	0.0024	0.0004	0.0020	-0.0053	0.0034
9	0.0017	0.0024	0.0007	0.0020	0.0018	0.0033
10	-0.0007	0.0024	-0.0025	0.0020	0.0018	0.0034
11	0.0007	0.0025	-0.0029	0.0020	-0.0073	0.0033
12	0.0017	0.0025	0.0022	0.0020	0.0001	0.0034
13	-0.0022	0.0025	-0.0017	0.0020	0.0040	0.0033
14	0.0046	0.0025	0.0006	0.0020	0.0060	0.0034
15	0.0015	0.0024	-0.0016	0.0020	0.0064	0.0034
16	0.0021	0.0024	0.0003	0.0020	-0.0018	0.0033
17	-0.0016	0.0024	-0.0013	0.0020	0.0040	0.0034
18	-0.0009	0.0024	-0.0004	0.0020	-0.0006	0.0034
19	-0.0003	0.0025	-0.0010	0.0020	0.0039	0.0034
20	-0.0003	0.0024	0.0007	0.0020	0.0037	0.0033

Table 7: **Abnormal Returns For Exploration Companies Sub-Index**

Event Day	Q↑ $\bar{\epsilon}^*$	Q↑ s.e.	Q↔ $\bar{\epsilon}^*$	Q↔ s.e.	Q↓ $\bar{\epsilon}^*$	Q↓ s.e.
-20	0.0034	0.0024	0.0016	0.0023	0.0020	0.0037
-19	0.0013	0.0024	0.0000	0.0023	-0.0034	0.0037
-18	-0.0026	0.0024	0.0005	0.0023	0.0053	0.0037
-17	-0.0010	0.0024	0.0009	0.0023	-0.0033	0.0037
-16	-0.0036	0.0024	0.0003	0.0023	0.0029	0.0037
-15	-0.0007	0.0024	0.0004	0.0023	0.0020	0.0037
-14	0.0016	0.0024	-0.0003	0.0023	0.0068	0.0037
-13	-0.0006	0.0024	-0.0004	0.0023	0.0062	0.0037
-12	-0.0016	0.0024	-0.0046	0.0023	-0.0034	0.0037
-11	0.0004	0.0024	0.0005	0.0023	-0.0009	0.0037
-10	-0.0025	0.0024	0.0020	0.0023	-0.0009	0.0037
-9	0.0007	0.0024	-0.0021	0.0023	0.0049	0.0037
-8	-0.0020	0.0024	-0.0004	0.0023	0.0010	0.0037
-7	-0.0015	0.0024	-0.0005	0.0023	0.0027	0.0037
-6	0.0008	0.0024	-0.0035	0.0023	0.0007	0.0037
-5	-0.0009	0.0024	-0.0003	0.0023	0.0063	0.0037
-4	-0.0052	0.0024	0.0002	0.0023	-0.0031	0.0037
-3	0.0013	0.0024	-0.0045	0.0023	0.0021	0.0037
-2	0.0016	0.0024	-0.0005	0.0023	0.0032	0.0037
-1	-0.0024	0.0024	-0.0028	0.0023	-0.0004	0.0037
0	0.0065	0.0024	0.0014	0.0023	-0.0040	0.0037
1	0.0027	0.0024	-0.0026	0.0023	-0.0009	0.0037
2	0.0030	0.0024	0.0010	0.0023	0.0046	0.0037
3	-0.0025	0.0024	-0.0027	0.0023	0.0028	0.0037
4	0.0014	0.0024	-0.0039	0.0023	0.0039	0.0037
5	-0.0014	0.0024	0.0008	0.0023	0.0009	0.0037
6	0.0024	0.0024	-0.0013	0.0023	0.0070	0.0037
7	-0.0009	0.0024	0.0005	0.0023	-0.0013	0.0037
8	-0.0054	0.0024	0.0012	0.0023	-0.0022	0.0037
9	-0.0001	0.0024	-0.0023	0.0023	0.0005	0.0037
10	-0.0003	0.0024	-0.0041	0.0023	-0.0011	0.0037
11	-0.0016	0.0024	-0.0051	0.0023	-0.0037	0.0037
12	0.0002	0.0024	-0.0001	0.0023	-0.0019	0.0037
13	-0.0029	0.0024	-0.0014	0.0023	0.0028	0.0037
14	0.0025	0.0024	0.0004	0.0023	0.0072	0.0037
15	-0.0003	0.0024	-0.0022	0.0023	0.0100	0.0037
16	0.0001	0.0024	0.0018	0.0023	0.0002	0.0037
17	-0.0001	0.0024	-0.0027	0.0023	0.0036	0.0037
18	-0.0016	0.0024	-0.0004	0.0023	-0.0006	0.0037
19	-0.0014	0.0024	0.0014	0.0023	0.0024	0.0037
20	-0.0016	0.0024	-0.0005	0.0023	-0.0016	0.0037

Table 8: **Abnormal Returns For Integrated Companies Sub-Index**

Event Day	Q↑ $\overline{\epsilon}^*$	Q↑ s.e.	Q↔ $\overline{\epsilon}^*$	Q↔ s.e.	Q↓ $\overline{\epsilon}^*$	Q↓ s.e.
-20	0.0009	0.0027	0.0030	0.0019	-0.0008	0.0032
-19	-0.0037	0.0027	0.0003	0.0019	-0.0018	0.0032
-18	-0.0019	0.0027	0.0013	0.0019	0.0020	0.0032
-17	-0.0005	0.0026	-0.0022	0.0019	-0.0032	0.0032
-16	0.0048	0.0027	-0.0009	0.0019	-0.0012	0.0032
-15	0.0003	0.0026	-0.0008	0.0019	-0.0025	0.0032
-14	0.0011	0.0027	-0.0009	0.0019	0.0049	0.0032
-13	-0.0033	0.0027	-0.0017	0.0019	0.0045	0.0032
-12	-0.0004	0.0027	-0.0021	0.0019	-0.0035	0.0032
-11	0.0046	0.0027	-0.0004	0.0019	-0.0001	0.0032
-10	-0.0036	0.0027	0.0024	0.0019	-0.0022	0.0032
-9	-0.0014	0.0027	-0.0026	0.0019	0.0045	0.0032
-8	-0.0060	0.0026	-0.0010	0.0019	0.0003	0.0032
-7	0.0016	0.0027	-0.0009	0.0019	0.0001	0.0032
-6	0.0039	0.0027	-0.0027	0.0019	0.0012	0.0032
-5	-0.0011	0.0026	-0.0023	0.0019	0.0024	0.0032
-4	-0.0041	0.0027	0.0030	0.0019	0.0018	0.0032
-3	0.0033	0.0027	-0.0024	0.0019	0.0011	0.0032
-2	-0.0020	0.0027	0.0004	0.0019	-0.0025	0.0032
-1	0.0011	0.0026	0.0018	0.0019	-0.0014	0.0032
0	0.0068	0.0026	0.0000	0.0019	-0.0010	0.0032
1	0.0012	0.0027	-0.0015	0.0019	-0.0026	0.0032
2	0.0028	0.0026	0.0029	0.0019	0.0000	0.0032
3	-0.0016	0.0026	-0.0025	0.0019	0.0007	0.0032
4	0.0060	0.0026	-0.0012	0.0019	0.0020	0.0032
5	-0.0001	0.0026	-0.0006	0.0019	-0.0017	0.0032
6	0.0051	0.0027	-0.0023	0.0019	0.0033	0.0032
7	-0.0024	0.0027	-0.0007	0.0019	-0.0003	0.0032
8	-0.0022	0.0026	0.0004	0.0019	-0.0055	0.0032
9	0.0027	0.0027	0.0017	0.0019	0.0033	0.0032
10	0.0001	0.0027	-0.0024	0.0019	0.0036	0.0032
11	0.0015	0.0027	-0.0020	0.0019	-0.0072	0.0032
12	0.0016	0.0027	0.0025	0.0019	0.0000	0.0032
13	-0.0018	0.0027	-0.0017	0.0019	0.0040	0.0032
14	0.0044	0.0027	-0.0001	0.0019	0.0049	0.0032
15	0.0023	0.0027	-0.0017	0.0019	0.0036	0.0032
16	0.0030	0.0027	-0.0003	0.0019	-0.0025	0.0032
17	-0.0010	0.0027	-0.0007	0.0019	0.0050	0.0032
18	0.0004	0.0026	-0.0009	0.0019	0.0004	0.0032
19	-0.0003	0.0027	-0.0023	0.0019	0.0040	0.0032
20	-0.0001	0.0027	0.0012	0.0019	0.0056	0.0032

Table 9: Estimated Coefficients: Dependent Variable is C_t^i

	Algeria	Indonesia	Iran	Kuwait	Libya	Nigeria	Qatar	Saudi Arabia	UAE	Venezuela
Quota	1.63	-1.15	0.19	-2.08	1.55	0.05	-3.45	0.08	-0.06	0.05
	(2.84)	(-3.26)	(1.29)	(-2.61)	(3.90)	(0.17)	(-1.79)	(0.92)	(-0.14)	(0.18)
Quota*Demand	-0.07	0.02	-0.01	0.03	-0.04	-0.01	0.02	-0.0038	-0.02	-0.01
	(-5.52)	(1.88)	(-2.40)	(1.57)	(-4.65)	(-1.55)	(0.54)	(-1.89)	(-1.50)	(-2.52)
Demand	0.06	-0.02	0.03	-0.04	0.06	0.02	0.03	0.03	0.05	0.06
	(5.34)	(-1.52)	(2.49)	(-1.05)	(4.57)	(1.75)	(1.56)	(2.07)	(2.06)	(4.36)
Non-OPEC	-0.01	-0.03	-0.04	-0.03	0.0030	-0.02	0.02	-0.04	-0.05	-0.03
	(-3.27)	(-9.22)	(-6.56)	(-1.27)	(0.44)	(-2.64)	(1.93)	(-2.92)	(-2.53)	(-3.33)
OIL	0.0024	4.7e-4	0.0040	-0.0023	0.0034	0.0015	-2.9e-5	0.0026	0.01	0.002
	(2.49)	(0.58)	(3.37)	(-0.44)	(2.82)	(1.12)	(-0.01)	(1.32)	(2.32)	(1.16)
Time	0.0012	5.8e-5	0.0012	0.0020	-7.1e-4	0.0013	0.0028	6.6e-4	6.2e-5	0.002
	(5.91)	(0.34)	(3.65)	(1.55)	(-2.34)	(3.31)	(4.36)	(1.11)	(0.06)	(3.56)
Q1	0.0044	-0.01	0.03	-0.05	-0.002	-0.02	0.0043	-0.06	-0.08	-0.02
	(0.45)	(-1.32)	(2.36)	(-0.90)	(-0.17)	(-1.10)	(0.16)	(-2.88)	(-2.25)	(-1.07)
Q2	1.3e-4	-0.01	0.0039	-0.10	0.01	-0.0015	0.08	-0.04	-0.05	0.02
	(0.01)	(-0.56)	(0.28)	(-1.70)	(0.71)	(-0.10)	(2.66)	(-1.78)	(-1.11)	(1.07)
Q3	0.0034	-0.01	-0.0041	-0.02	0.01	-0.01	0.09	-0.02	-0.03	0.01
	(0.33)	(-0.96)	(-0.32)	(-0.40)	(0.87)	(-0.72)	(3.14)	(-0.97)	(-0.65)	(0.59)
WAR	-0.0024	0.07	0.01	-1.10	0.11	0.07	0.08	0.28	0.29	0.05
	(-0.11)	(4.12)	(0.23)	(-9.50)	(4.57)	(2.62)	(1.44)	(6.87)	(3.82)	(1.36)
Const	-0.44	2.62	0.71	4.06	-2.24	0.44	-0.93	0.51	1.34	-0.16
	(-0.87)	(5.10)	(1.07)	(2.14)	(-3.30)	(0.58)	(-0.97)	(0.55)	(0.99)	(-0.18)

*t-statistics are below the estimated coefficients.

**We do not estimate the model on Iraqi data given its unique relationship with OPEC since the end of the Persian Gulf War in 1991.