Disagreement In Bargaining: An Empirical Analysis of OPEC*

Kyle Hyndman
Southern Methodist University

June 2006, revised June 2007

We consider a stylised model in which two cartel members bargain over the aggregate-production quota in a world of asymmetric information. We show that when the two cartel members are sufficiently different, the probability of agreement depends on both the current state of demand and initial production. Specifically, the probability of agreement is much lower when demand is low (and initial production is relatively high) than when demand is high (and initial production is relatively low). We also find that, regardless of the current demand state, the more extreme is initial production, the higher is the probability of agreement. Using an event study, where we take as events OPEC production quota announcements, we demonstrate empirically that the predictions of the model are borne out.

JEL Classification: D70, D82, L13, L71

Keywords: Status Quo, OPEC, Collusion, Disagreement, Bargaining

1 Introduction

In this paper, we study whether agreements are easier to achieve when times are good or when times are bad. In the industrial organisation literature, this is an issue that has received careful attention starting with the work of Rotemberg and Saloner (1986) who,

*I have received helpful comments from many people but would like to single out, in particular, my advisor Debraj Ray as well as Luís Cabral, Giovanni Serio, an Associate Editor and two anonymous referees. Financial support from NSF grant #0241070 (Debraj Ray) is also gratefully acknowledged. This is a substantially revised version of the first chapter of my Ph.D. thesis as well as an earlier draft circulated under the title, "Status Quo Effects in Bargaining: An Empirical Analysis of OPEC."
Address: Department of Economics, Southern Methodist University, 3300 Dyer Street, Suite 301R, Dallas, TX 75275. Tel: (214) 768-2755. email: hyndman@smu.edu, url: http://faculty.smu.edu/hyndman.

more precisely, ask when is it easier for firms to collude. In a world of i.i.d. demand shocks, Rotemberg and Saloner show that collusion is most difficult when demand is high. Subsequent work has provided conditions for which it is hardest to collude during recessions (see *e.g.,* Athey, Bagwell and Sanchirico – henceforth ABS – (2004), Bagwell and Staiger (1997) and Satiger and Wolak (1992)). There is also some debate at the empirical level as to whether collusion is easier or harder to sustain during good times. Scherer and Ross (1990, Ch. 8) have argued that collusion is more difficult to sustain during recessions. Wilson and Reynolds (2005) provide empirical evidence consistent with the view that successful collusion is more difficult during booms, though they caution that other macroeconomic factors may be at play which are not captured by their oligopoly model (see p. 165).[1]

We contribute to this debate by providing a model in which the probability of agreement depends on the current state of nature, and by providing conditions under which agreements are more likely in good times than in bad times. Our model eschews the traditional oligopoly models in favour of a bargaining approach. This is because our interest lies in examining the behaviour of the Organization of Petroleum Exporting Countries (OPEC) and we feel that the bargaining problem its members confront at each meeting cannot be ignored. OPEC is obviously a cartel that restricts output in order to obtain super-competitive profits and must be concerned with the incentives each of its members has to overproduce. 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, and the problem of splitting these gains involves a great deal of closed-door negotiations. At times individual members' posturing for market share leads to extended periods of inaction, causing lower (though still super-competitive) profits for all.

Before discussing the precise details of our bargaining model, we first motivate the necessary departure from the standard oligopoly models of collusion. Consider ABS (2004) who study optimal collusive behaviour when firms' marginal costs are subject to random shocks.[2] While they show that collusion is more difficult to sustain during bad times, their model makes other predictions which are not borne out in our data. In particular, market share instability is a key feature of ABS (and other models of collusion), since low-cost firms are allowed to undercut the monopoly price without fear of punishment in the low-demand state. In contrast, at least since the early 1980s, OPEC has been playing a quantity game, setting shares of an aggregate quota and allowing prices to fluctuate.

If one looks at Figure 1, which displays each member's share of the aggregate quota over

[1] Although Wilson and Reynolds's model is non-collusive, many of its testable implications are identical to those of Staiger and Wolak (1992), which is a model of collusion. We will discuss Staiger and Wolak (1992) below.

[2] In the context of OPEC, it may strain reality to assume that marginal costs are unknown and subject to random shocks; however, it is reasonable to assume that there are other economic and political factors, which are unknown and random, that affect a country's "opportunity" marginal cost.

time, it appears that the shares of most countries have been relatively stable over time. There are a couple of exceptions but we feel that they have little to do with OPEC. For example, during the first Persian Gulf War, Saudi Arabia's share of the aggregate quota increased dramatically, while Iraq's and Kuwait's shares dropped to zero. Following the war, Kuwait's share recovered to its pre-war level, while Iraq's share has fluctuated widely for obvious political reasons. To a lesser extent Indonesia's share has declined, while Qatar and Venezuela have seen their shares increase over time. For Indonesia, this is due mainly to declining reserves, while the increased prominence of Qatar and Venezuela can partly be explained by the drop in Iraq's production after 1998. Hence, our point is not that OPEC shares have not changed, but that they are much more stable than received theory would predict.[3]

FIGURE 1: OPEC-Member Quota Shares

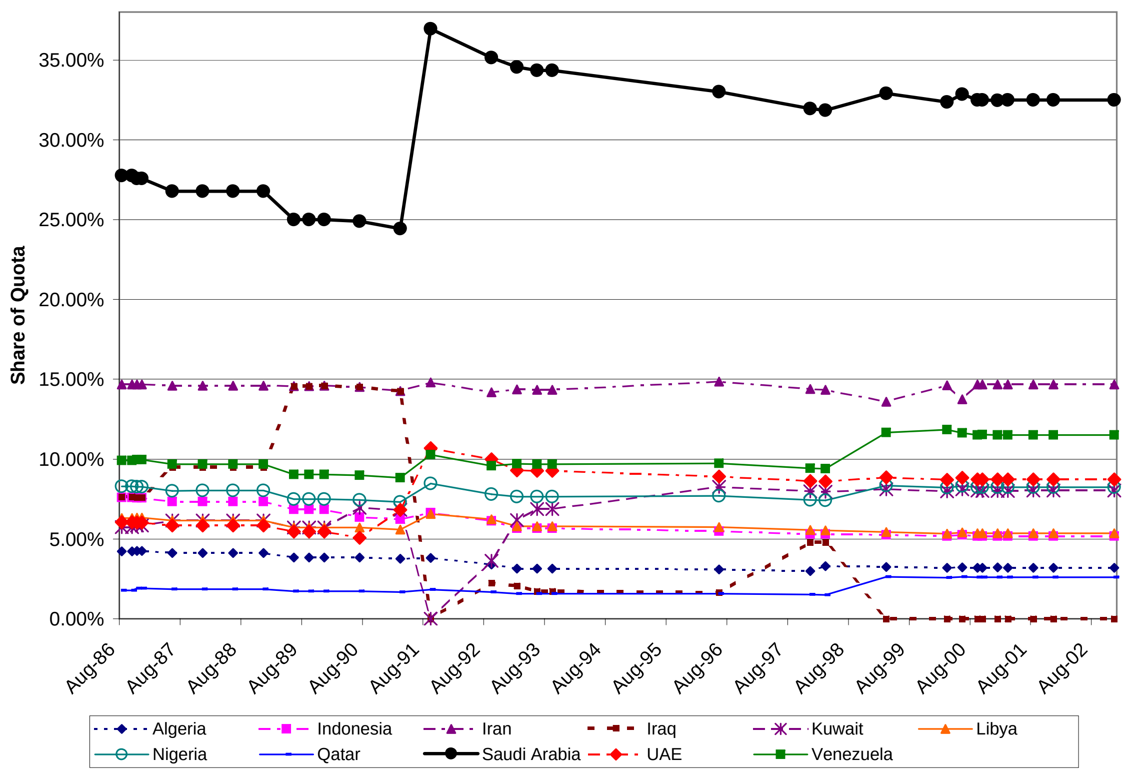

Another important stylised fact which motivates our bargaining approach concerns the size of shocks and the probability of agreement. In Section 3.4, we show that for large shocks (positive or negative) OPEC is more likely to reach agreement than for small shocks. In the model of Staiger and Wolak (1992), for small negative shocks, the maximal level of collusion

[3]One may also be concerned that while quota shares have been stable over time, actual production has not. Cheating is a problem in OPEC; however, Kaufmann *et al* (2004) have shown in a regression analysis that OPEC influences oil prices through two distinct channels: its announced quota and by how much it cheats on the quota. We focus on the former channel.

decreases in a continuous manner, while for large negative shocks, collusion breaks down and players employ mixed strategies in the quantity-setting subgame. As with ABS, market shares are unstable but, beyond that, large shocks lead to unsuccessful collusion. In our bargaining model, the presence of private information creates a wedge between the interests of the proposer and the responder. Importantly, the size of this wedge is independent of the size of the demand shock. Therefore, for a larger demand shock, the private information is relatively less burdensome, leading to an increased likelihood of agreement.

We are not the first to discuss collusion and cartels in a bargaining framework. For example, Ray and Vohra (1997) use their model to characterise stable cartels, while Seidmann and Winter (1998) provide a brief discussion of cartels in the context of gradual coalition formation. In papers more closely related to ours, Cramton and Palfrey (1990, 1995) discuss cartel formation through the lens of mechanism design. The connection to bargaining is particularly strong in the latter paper where the authors include an explicit ratification stage. Finally, and perhaps most valuable to us, there is a survey by Levenstein and Suslow (2002) on what determines cartel success. In it, they argue that, "[b]argaining problems were much more likely to undermine collusion than was secret cheating" (p. 16). In addition they state that, "bargaining issues may arise as a result of a decline in demand" (p. 18). Scott Morton (1997) provides further evidence along these lines in her study of British shipping cartels arguing that, "It was much easier for a [cartel] to allocate six sailings a year to an entrant if the original members could keep their current schedules . . . Hence, increasing trade on a route made negotiating entry easier" (p. 702).

In Section 2 we provide our model of a quota-setting cartel and develop some testable implications, which we then take to our empirical study of OPEC. Our model completely abstracts from the problem of sustaining collusion, assuming that agreements between agents are binding and enforceable. In the model's static incarnation, there are two cartel members who negotiate over the aggregate quota, leaving individual shares fixed. The basic ingredients of the model are as follows: given demand and initial production, the proposer (firm 1) offers a potentially different aggregate quota, which is either accepted or rejected by the responder (firm 2). If the proposal is accepted, the new production level is implemented, while if it is rejected, the status quo obtains. In order to introduce the possibility of disagreement, we assume that there is some uncertainty over the respondent's ideal point. This uncertainty may arise, for example, due to randomness in its marginal cost of production, or other economic or political considerations that influence the respondent's ideal point.[4] Firms are also asymmetric in the sense that one of them, on average, has a higher ideal point

[4]For example, within the context of OPEC, one could argue that the state of relations between Saudi Arabia and the United States is an important factor that is subject to (private) fluctuations and likely influences Saudi Arabia's negotiating stance. There may also be idiosyncratic uncertainty regarding the effectiveness of sanctions on Iran, or insurgents' ability to disrupt Iraqi oils shipments.

than the other. We provide conditions under which disagreements are more likely when initial production is high relative to demand (and so should be reduced) than when initial production is low relative to demand (and so should be increased). We also show that the more extreme is initial production (either too high or too low), the more likely is an agreement to be reached.

Of course, the real-world in which OPEC operates is not static — demand fluctuates over time, as do costs and other economic and political variables. In Appendix A we show that the model's two main predictions remain intact in a dynamic world with Markovian demand shocks.

In Section 3 we turn to our empirical study of OPEC. Our approach is at first indirect. In particular, we conduct an event study in order to determine how oil prices react to various OPEC production quota announcements. As is standard with this kind of analysis, 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 aggregate quota, (2) decreases in the aggregate quota and (3) no change in the aggregate quota (or status quo). We then examine how the price of oil responds to each announcement type and argue that the results imply that OPEC members find it more difficult to agree during bad times than during good times.[5]

In analysing how oil prices react to the three kinds of announcements, Section 3.3 presents three stylised facts. First, we find no evidence of abnormal returns, either positive or negative, when OPEC announces an increase in the aggregate quota. Second, we find significantly positive abnormal returns following an announced reduction in the aggregate quota. Third, we find significantly negative abnormal returns following status quo announcements. This suggests that the market is surprised by status quo announcements and aggregate quota reductions, but not by announcements of increases in the aggregate quota. All three stylised facts hold for both the spot and two-month forward price of oil. Our take on this is as follows: *status quo announcements represent a failure to reduce the quota during bad times.* That is, during bad times, the market holds intermediate beliefs and so when the aggregate quota is reduced, a positive reaction occurs, while when no change is announced a negative reaction occurs. In contrast, during good times, OPEC always increases the aggregate quota and so the market fully expects this, leading to no reaction to such announcements.

The remainder of Section 3 takes pains to show that the initial results and our interpretation of them are robust. In particular, we explicitly correlate OPEC announcements with variables pointing to strength or weakness in the oil market (specifically inventory growth

[5]Note that through its influence on oil prices, OPEC also influences the profitability of oil companies because they all control substantial reserves whose value is tied to the world price of oil. Therefore, we would expect the same pattern of results to hold for stock market indices of oil companies. We will discuss this point in more detail later.

by western nations and residual demand growth). Beyond merely reinforcing our earlier interpretation, these results support the model's second main prediction: the larger is the size of the shock, the more likely is OPEC to reach an agreement.

Finally, in Section 4 we provide some concluding remarks.

2 A Model of One-Sided Private Information

In this section we formulate our model of quota formation and derive its testable implications, which will then be taken to our empirical analysis of OPEC. In the main text we present only the static version of the model, leaving for Appendix A the generalisation to a dynamic framework.

As stated in the Introduction, assume that there are two firms in a cartel negotiating over the aggregate production Q. In order to keep things as simple as possible we make a number of assumptions. The timing is as follows: Given initial production Q_0, Firm 1, who faces an inverse demand curve $P_1(A, Q) = A - Q$, proposes a potentially new production level, Q' to firm 2 who either accepts it or rejects it. If the proposal is accepted, Q' is implemented, while if it is rejected, the status quo obtains. We also assume that firm 2 faces an inverse demand curve $P_2(A, Q, \eta) = A + \eta - Q$, where η is uniformly distributed on the interval $[-\alpha, \alpha + 2\beta]$, where α, $\beta \geq 0$.[6] At the time firm 1 makes its proposal, the realisation of η is known to firm 2 but unknown to firm 1.

We interpret α as capturing the degree of *symmetric* uncertainty in the model, whereas $\beta = \mathbb{E}[\eta]$ captures the expected degree of asymmetry between firms 1 and 2. When $\beta > 0$, on average, firm 2's ideal quantity is $\frac{A+\beta}{2}$, whereas, firm 1's ideal quantity is only $\frac{A}{2} < \frac{A+\beta}{2}$; therefore, the larger is β, the more potential for tension there is between the two firms.[7]

We are now ready to solve for the equilibrium of this model and find conditions under which agreement is more difficult during "bad" times than during "good" times. In a static setting, the terms "good times" and "bad times" do not have a precise meaning. We will take "good times" to mean that initial production is low relative to the proposer's ideal point, whereas "bad times" will indicate initial production that is high relative to the proposer's ideal point. In the dynamic model, good and bad times will have a more precise meaning, but the results will not change.

Given an offer of Q and initial production Q_0, the responder will accept if and only if

[6] While we have included the random component of firm 2's profit function as part of the inverse demand curve, this is for convenience only. Randomness could come from a variety of sources; for example, shocks to economic or political variable affecting the ideal point of the responder.

[7] One could easily re-write the model with $\beta < 0$, in which case firm 2 will generally prefer lower production than firm 1. Similar predictions are easily derived, though we focus on $\beta \geq 0$.

$(A+\eta-Q)Q \geq (A+\eta-Q_0)Q_0$, which can be rewritten as:

$$\eta \left\{ \begin{array}{c} \leq \\ \geq \end{array} \right\} Q_0 + Q - A \left\{ \begin{array}{l} \text{if } Q < Q_0 \\ \text{if } Q > Q_0 \end{array} \right\} \tag{1}$$

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

$$P(Q, Q_0) \in \begin{cases} 1, & \text{if } Q = Q_0 \\ \max\{0, \min\{1, \frac{Q+Q_0-A+\alpha}{2(\alpha+\beta)}\}\}, & \text{if } Q < Q_0 \\ \min\{1, \max\{0, 1 - \frac{Q+Q_0-A+\alpha}{2(\alpha+\beta)}\}\}, & \text{if } Q > Q_0 \end{cases} \tag{2}$$

Given (2), we may write the expected profit function of the proposer as:

$$\mathbb{E}[\pi_1(A, Q, Q_0)] = Q_0(A - Q_0) + P(A, Q, Q_0)\left[Q(A-Q) - Q_0(A-Q_0)\right] \tag{3}$$

Despite a discontinuity at $Q = Q_0$ for $\beta \neq 0$, it is easy to see that (3) is a continuous function of Q. Let $Q^*(Q_0)$ denote the optimal proposal by firm 1 to firm 2.[8]

Our main results can be summarised in the following two propositions:

Proposition 1. *Suppose that $Q_0 < \frac{A}{2}$. Then:*

(a) $Q^*(Q_0) \in [Q_0, \frac{A}{2}]$.

(b) The optimal proposal $Q^*(Q_0)$ may be increasing or decreasing in Q_0; however, there exists a threshold $\underline{Q} \leq \frac{A}{2}$ such that $Q^*(Q_0)$ is increasing in Q_0 for all $Q_0 \geq \underline{Q}$.

(c) The probability that the optimal offer is accepted, $P(Q^*(Q_0), Q_0)$ is a (weakly) decreasing function of Q_0.

On the other hand, if $Q_0 > \frac{A}{2}$:

(a′) $Q^*(Q_0) \in [\frac{A}{2}, Q_0]$.

(b′) The optimal proposal $Q^*(Q_0)$ may be increasing or decreasing in Q_0; however, there exists a threshold $\bar{Q} \leq A$ such that $Q^*(Q_0)$ is increasing in Q_0 for all $Q_0 \in [\frac{A}{2}, \bar{Q}]$.

(c′) The probability that the optimal offer is accepted, $P(Q^*(Q_0), Q_0)$ is a (weakly) increasing function of Q_0.

Proposition 2. *If $\beta > 0$, then for all $x \in [0, \frac{A}{2}]$, $P(Q^*(\frac{A}{2} - x), \frac{A}{2} - x) \geq P(Q^*(\frac{A}{2} + x), \frac{A}{2} + x)$, with a strict inequality for x sufficiently close to zero.*

[8] For ease of notation, we omit the obvious dependence of Q^* on both A and β.

Proposition 1 is essentially a computational exercise; however, some broad intuition can be given. When initial production is very low (and α is not too large), the interests of the proposer and all responder types are aligned: both would like higher production. In such situations, agreements occur with probability 1 on firm 1's ideal point. Eventually, as the status quo increases the interests of the proposer and *low* responder types (*i.e.,* $\eta \approx -\alpha$) will begin to diverge. At this stage the proposer has two options. Either she can continue to propose $\frac{A}{2}$ and have the offer rejected with positive probability, or she can lower the proposal slightly in order to ensure "unanimous" consent. Initially, she will prefer the latter option — thereby leading proposals to decrease. However, eventually, the quantity that would ensure unanimous consent becomes too low and so firm 1 prefers to increase her proposal and risk rejection by low responder types. This intuition also makes it apparent that the further away is initial production from firm 1's ideal point, the easier it is for agreement to be reached. There are two related reasons for this. First, the further initial production is from $\frac{A}{2}$, the more there is to be gained by adjusting output (either up or down). Second, with extreme initial production levels the interests of the proposer and many responder types are aligned. In contrast, for more moderate levels of initial production, there is less to be gained from a change in output and the interests of more and more responder types will diverge from those of the proposer. Figure 2 depicts equilibrium offers for the particular example of $(\alpha, \beta, A) = (2, 1, 10)$, while Figure 3 shows the equilibrium acceptance probabilities for the same parameter values.

The logic behind Proposition 2 is also relatively straightforward. Firm 1's profit function, $(A - Q)Q$ is symmetric about $\frac{A}{2}$; however, when $\beta > 0$, on average firm 2 prefers higher quantities. Therefore, when initial production is high it is substantially more likely that the interests of the proposer and responder will diverge than it is when it is low. Importantly, as can be seen in Figure 3, there is a discontinuity at $Q_0 = \frac{A}{2}$. This arises due to the asymmetry of the distribution of responder types. Moreover, it is easy to see that as β increases, it becomes ever easier to increase production and ever more difficult to reduce it.

3 EMPIRICAL ANALYSIS OF OPEC

In this section, we turn our attention to OPEC. In particular, we seek empirical support for the two main predictions of the model:

1. Agreements are more likely in high-demand states than in low-demand states; and

2. Agreement is more likely the greater is the need to change the quota.

We begin by briefly describing the data and our empirical methodology.

FIGURE 2: Equilibrium Offers

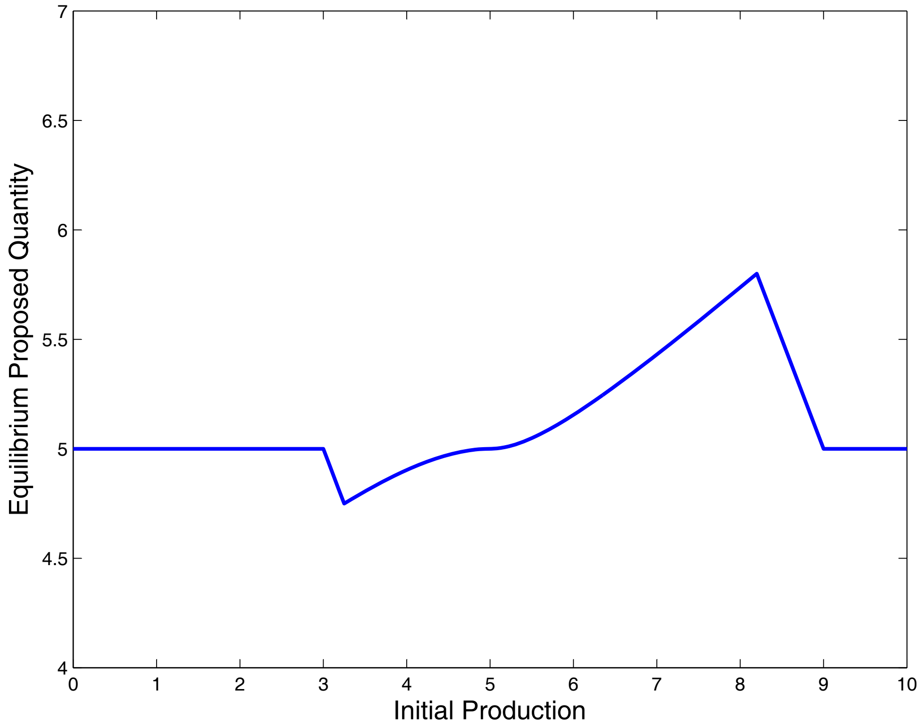

3.1 Data

Our primary data consists of OPEC quota announcements over the period August 1986 to September 2002. These announcements come in three forms: aggregate-quota reductions, aggregate-quota expansions and status quo announcements. For each of these three announcement types we will examine how certain certain economic variables react through our event study. In particular, we will mostly concern ourselves with the reaction of oil prices to these announcements; however, we will also examine these announcements' effects on the stock market returns of groups of oil companies.

More precisely, our sample contains 52 announcements by OPEC, of which 51 can be unambiguously classified.[9] Eleven of these announcements are for aggregate-quota reductions, 17 for aggregate-quota increases and 23 maintain the status quo. We also have daily data on both the spot price and the two-month forward price of West Texas Intermediate crude

[9] The meeting of December 28, 2001 formally announced that the quota reduction agreed to on November 14, 2001, but which was conditional on reductions by Norway, Russia and other producers, would be implemented. Therefore, it is not clear whether both of these announcements should be considered reductions, or if one should be considered a status quo announcement. Our analysis presumes that the November 14 announcement is a reduction, while the December 28 announcement is not. None of the results presented are sensitive to this.

Figure 3: Equilibrium Acceptance Probability

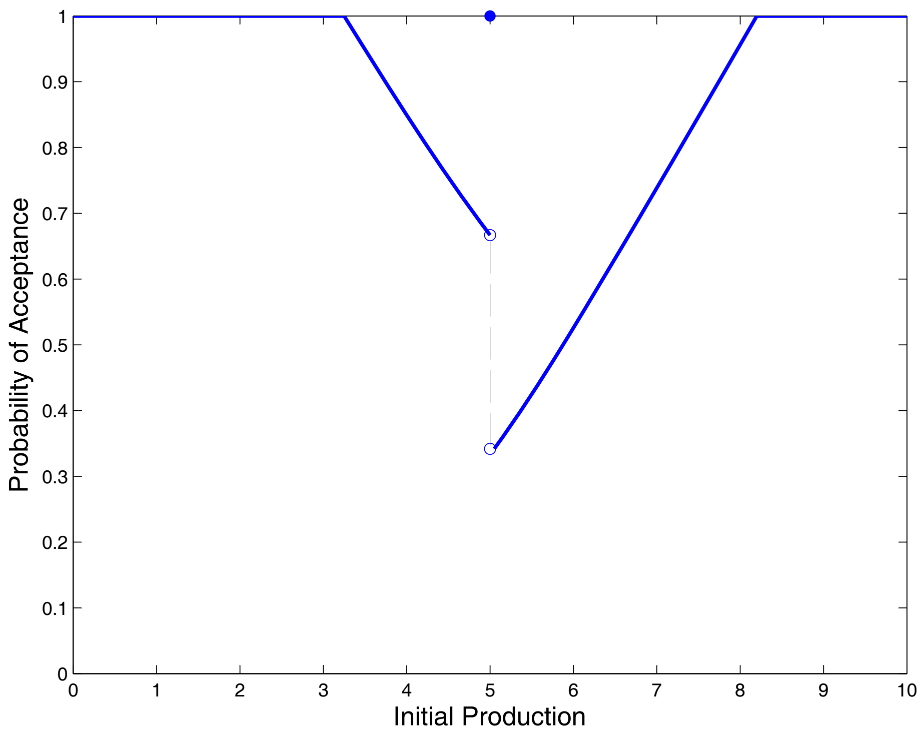

oil. While OPEC has its most immediate impact on oil prices through its actions, it also indirectly influences the profitability of oil companies because they control substantial oil reserves whose value is tied to the world market price of oil. Therefore, to the extent that we see a reaction in oil prices to OPEC announcements, we should see a similar reaction in the stock prices of these oil companies. In order to capture this, we also conduct our event study on an index of all oil companies, as well on a sub-index of exploration and production companies, whose profitably is most directly tied to the price of oil amongst all the oil industry sub-sectors.[10]

Later, in examining the robustness of the results of our event study and in showing support for the second prediction of the theoretical model, we will also explicitly correlate the outcome of OPEC meetings to economic fundamentals which point to strength or weakness in the oil market from OPEC's perspective. In particular, we will look at the effect of changes in crude oil inventories by western nations and residual demand growth[11] on the likelihood of agreement.

[10]All of these, including the two series of oil prices, were obtained from Datastream. The specific oil company indices used were OILEPAM and OILGSAM. The former is an index of the exploration & drilling sub-sector, while the latter is an index for the overall oil sector for the Americas.

[11]That is, the demand for oil that is left over after accounting for non-OPEC oil production.

Before proceeding with the results, a word on OPEC is warranted. First, note that OPEC typically holds between two and four regularly scheduled meetings per year, but will sometimes call extraordinary meetings.[12] Often these extraordinary meetings are announced at the conclusion of a regularly scheduled meeting and usually occur between one and three months later. There has also been a downward trend in the duration of OPEC meetings. In the 1980s it was not uncommon for meetings to last several days. However, in recent years, OPEC meetings have become much more of a formality, with all of the tough negotiating taking place in the days and weeks before the actual gathering. It is this feature of OPEC which necessitates a special event window in our empirical analysis (see discussion below). Table 3 summarises the OPEC meetings in our sample. Specifically, it provides the start and end date of the meeting, the type of meeting, the announced change in the quota as well as the year on year change in each of inventories, residual demand and the spot price of oil.

3.2 Methodology

Our main tool of analysis is the by now well familiar method of event studies. In particular, we will look at the cumulative abnormal return functions for oil prices and stock-market indices of oil oil companies in a window around OPEC announcements. The reader is referred to the summary of event study techniques found in Campbell *et al* (1997).

Note that the timing of the OPEC announcements are all known well in advance, but the precise announcement (*i.e.,* whether OPEC announces an increase in the quota, a decrease in the quota or no change in the quota) is presumably not known until the announcement is actually made. The main question that arises from this methodology is the choice of the size of the event window. We take ours to be relatively large in comparison to much of the existing empirical literature — 15 to 20 days on either side of the announcement date — for a number of reasons. As can be seen from Table 3, at least in the early days, OPEC meetings tended to last many days; therefore, in order to capture any information leakages during the meeting, the event window should encompass the entire duration of the meeting. Of course, more recently, OPEC meetings have become more “summit-like”, with major negotiations taking place before the meeting itself. It would still seem prudent to take a sufficiently large window in order to include the negotiating phase, for as Mabro (2001) indicates, OPEC members are often in the habit of making public their conflicting negotiating stances. Similarly, after the announcement has occurred, OPEC members often make public declarations either in favour of or against the particular announcement, or about the level of discord within OPEC. All of these issues are relevant to our study, and so we feel, should be accommodated for when choosing the event window. Therefore, by choosing the window to be relatively large, we

[12] The current trend is for two regularly scheduled meetings per year but with more frequent extraordinary meetings.

seek to capture important information leakages both before and after the event that signal whether or not an agreement is likely to be reached, and if so, whether it is likely to be abided by.

3.3 Results: Probability of Agreement & Business Cycle

3.3.1 Normal & Extraordinary Meetings

Figure 4 plots the cumulative abnormal return functions for each of the three announcement types for the two month forward price of oil. In Appendix B we provide identical figures for the spot price of oil, as well as for an index of oil companies and an index of exploration & drilling companies. Consider first announcements to decrease the aggregate quota. As can be seen in Figure 4 and in Table 4, for the two-month forward price of oil, the spot price of oil and exploration companies, the cumulative abnormal returns become significantly positive at the end of the event window at the 5% level (one-sided test). Indeed, here the effect is quite strong; for both the spot and two month forward price of oil, at the end of the event window, the cumulative abnormal returns are nearly 10% while in for exploration and drilling companies they are approximately 6%. As can be seen from Figure 8 and Tables 5 – 7, while significance varies, the CAR functions for aggregate quota reductions are remarkably similar across the measures we used. In particular, in the days leading up to an announced reduction in the aggregate 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. Thus by announcing a reduction in the aggregate quota, OPEC has the ability to substantially influence oil prices, and also the stock market price of oil companies.

Next consider status quo announcements. We find negative cumulative abnormal returns which start to accrue in the days leading up to the announcement, continuing to further decline in the days 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; see Tables 7 and 6) towards the end of the event window. It is also important to note that these abnormal returns are quite large in percentage terms. At the end of the event window, the cumulative abnormal returns are between -2 and -3.5%. Thus, by its inaction, oil prices are significantly lower than they would have been but for the impact of OPEC's status quo announcement.

Finally we consider announcements of aggregate quota increases. In all cases, announcements of increases in the aggregate quota were met with no significant abnormal returns - either positive or negative.

Figure 4: Cumulative Abnormal Returns Functions for the Two-Month Forward Price of Oil

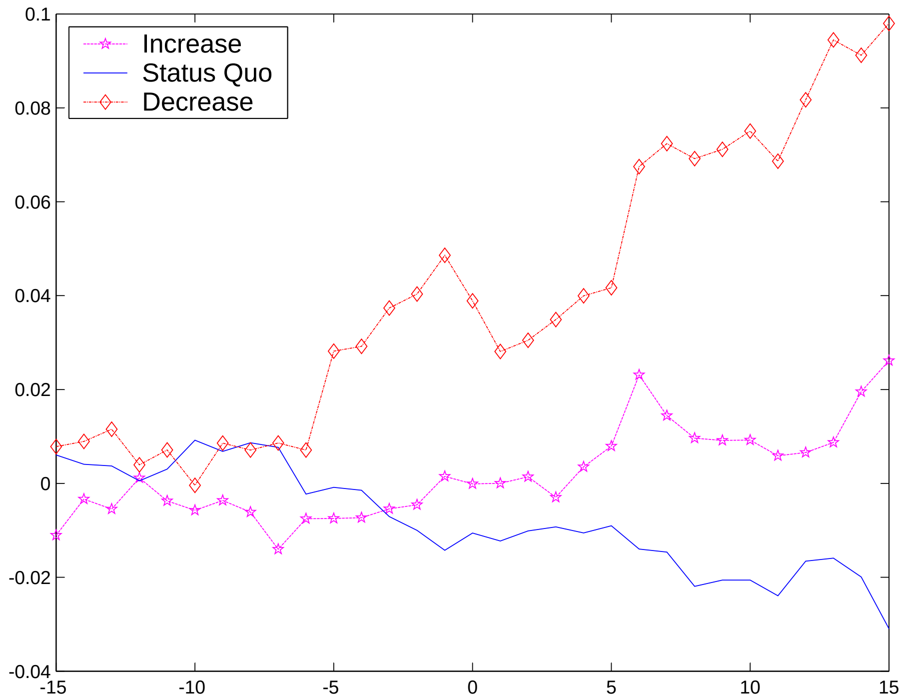

Note that, the cumulative abnormal return functions for the variables we considered (*i.e.,* spot price of oil, two-month forward price of oil, oil companies stock market index, exploration & production companies stock market index) were not all significantly different from zero at the end of the event window. However, in all cases, the *difference* between the CAR functions for aggregate-quota reductions and the CAR function for status quo announcements *is* significantly different in all cases. That is, the market reaction is very different for an announcement of an aggregate-quota reduction than it is for a status quo announcement.

3.3.2 Interpretation

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. The idea is very straightforward. When times are bad OPEC should reduce the aggregate quota, while when times are good OPEC should increase the aggregate quota. We have shown that when times are bad, sometimes they reduce the aggregate quota and sometimes they do not. Therefore, during bad times, the market holds intermediate beliefs and so when the aggregate quota is reduced, a positive reaction occurs, while when no change is announced a negative reaction occurs. In contrast, during good times, OPEC always increases the aggregate quota and so the market fully expects this, leading to no reaction to such announcements.

We believe this to be a compelling rationalisation of the empirical results and, moreover, one that is consistent with the main testable prediction of our theoretical model. While reasonable, our argument is but one interpretation and rests upon a logical argument. In Section 3.4 we will provide further support for our interpretation, and also show consistency with the other predictions of our model. 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 and not necessarily in the direction that we find. For example, upon observing an increase in the aggregate quota, it may signal that demand is actually high and lead to positive returns. However, we find this argument fundamentally implausible — timely information on all aspects of the oil market is widely available by the U.S. Department of Energy's Energy Information Administration 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 aggregate quota, either up or down, signals that OPEC as a cartel is strong but that status quo announcements signal weakness in the cartel. Suppose that this argument 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 aggregate quota.

3.3.3 Normal vs. Extraordinary Meetings

The results discussed above were based on the entire sample of events. However, we must recognise that 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. As we have said, OPEC typically has between two and four regularly scheduled meetings per year, but will sometimes hold extraordinary meetings.[13] Typically OPEC gives substantial advanced notice of an extraordinary meeting, but this still raises some concerns about the causal relationship between the state of demand and the outcome of OPEC meetings. For example, rather than being called in response to a demand shock (as we argue), it may be that extraordinary meetings are called at times of extraordinary discord within the cartel. In this case, one might imag-

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

ine that disagreement would be even more likely, and we would expect announcements of increases or decreases in the quota to be *more* surprising, while status quo announcements would be *less* surprising. Thus by restricting attention to normal meetings only, the CAR functions for reductions and increases should be larger, while the CAR functions for status quo announcements should be closer to zero.

Of course, it may well 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. Under this scenario, announcements of increases or decreases in the aggregate quota would be *less* surprising, while status quo announcements, signaling disagreement, would be *more* surprising. Therefore, if we look at only 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.

We replicate our event study, excluding all extraordinary meetings, the relevant figures being relegated to Appendix B. 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 aggregate-quota reductions are met with positive abnormal returns of between 4 and 7%. For the oil industry as a whole and for exploration and drilling companies, there appears to be very little impact for aggregate-quota increases.

The most striking difference between the current and previous analysis is that more information appears to leak 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 aggregate-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. We also note that there is some support for the conjecture that extraordinary meetings signal that OPEC is *more* likely to agree.

3.4 Results: The Size of the Shock Matters

Having shown empirically via our event study that agreements are more likely during good times than bad, we seek now to examine the second prediction of our theoretical model: the larger is the size of the shock, the more likely is an agreement to be reached. If we adopt the view of OPEC as a residual supplier of oil on the world market then we define shocks in at least a couple of different ways. First, if we look at the year-on-year change in residual demand, then positive residual demand growth indicates a positive shock.[14] Second, we can look at the year-on-year growth of inventories held by western nations. Here things go in the opposite direction with positive growth in inventories representing a negative shock from OPEC's perspective.

Consider the following latent variable model:

$$y_i^* = \beta_1 \Delta \text{stocks}_i + \beta_2 (\Delta \text{stocks}_i)^2 + \beta_3 \Delta \text{resid}_i + \epsilon_i \tag{4}$$

but we observe:

$$y_i = \begin{cases} -1, & \text{if } \Delta \text{quota} < 0 \iff y_i^* \leq \mu_1 \\ 0, & \text{if } \Delta \text{quota} = 0 \iff \mu_1 < y_i^* \leq \mu_2 \\ 1, & \text{if } \Delta \text{quota} > 0 \iff \mu_2 < y_i^* \end{cases} \tag{5}$$

In this specification, one can think of y^* as an unobserved measure of the actual aggregate quota in relation to its optimal value. The other terms are as follows: Δstock_i is the year-on-year change in crude oil stocks at the time of event i, Δresid_i is the year-on-year change in residual demand for oil, $\epsilon_i \sim N(0, \sigma^2)$ is a random error term and μ_1 and μ_2 are threshold parameters which must also be estimated.

If, as we have argued above agreements are easier during good times, then we expect $\beta_1 < 0$ and $\beta_3 > 0$. While we have no prediction for β_2, we include it both because it improves the fit and because its sign may be of independent interest. Indeed, as Table 1 indicates, our predictions for β_1 and β_3 are borne out, though only in the former case is the result significant.[15] If we restrict attention (not shown) to the binary decision of whether or not to increase the total production quota, then β_3 is also significantly positive. This suggests that changes in demand are only salient to OPEC when it comes time to increasing the aggregate quota, while inventory changes have salience on both the upside and the downside. However, the results of this estimation tells us more. Specifically, the larger (smaller) is the year-on-year increase (decrease) in residual demand (inventories), the more likely is OPEC to increase the aggregate quota. Therefore, Table 1 provides support for both of the model's testable implications.

We also plot the results of our estimation in Figure 5. In this figure, the solid line plots the estimated probability that OPEC will increase the aggregate quota as a function of the year-on-year change in inventories (*i.e.,* $\Pr[\Delta \text{quota} > 0 | \Delta \text{stocks}]$); similarly, the dashed line is the estimated probability that OPEC will reduce the aggregate quota as a function of the year-on-year change in inventories (*i.e.,* $\Pr[\Delta \text{quota} < 0 | \Delta \text{stocks}]$), while the dash-

[14] For our purposes we define residual demand to be the difference between oil consumption and non-OPEC oil production. That is, after all non-OPEC production is accounted for, it is the demand that is left over for OPEC to supply.

[15] We have only 51 observations, rather than 52 because there were two meetings held in July 2001. See Table 3 for more details; none of the results reported here depend on this.

TABLE 1: **Estimated Parameter Values: Ordered Probit Regression**

	Coef.	Std. Err.	z
Δstocks	-0.00777	0.00307	-2.53
$(\Delta\text{stocks})^2$	0.000042	0.00002	2.33
Δresid	0.000055	0.00011	0.52
μ_1	-0.721	0.24498	
μ_2	0.638	0.24012	
	$n = 51$	$LR(3) = 15.08$	

dotted line is the estimated probability that OPEC will leave the aggregate quota unchanged (*i.e.,* $\Pr[\Delta\text{quota} = 0|\Delta\text{stocks}]$). Importantly, this figure corresponds quite closely to the predictions of our theoretical model. First, for Δstocks < 0 and large, OPEC is very likely to increase the aggregate quota. Second, as Δstocks increases it becomes less and less likely that OPEC will increase the aggregate quota. Third, for Δstocks > 0 and large, it is almost equally likely that OPEC will reduce or leave the aggregate quota unchanged. Fourth, as Δstocks decreases, the probability that OPEC reduces the aggregate quota decreases, while the probability of no change in the aggregate quota increases. Finally, for comparably sized shocks (but in the opposite direction) it is always strictly more likely that the quota will be increased than that it will be decreased.

3.5 Potential Issues With Empirical Methodology and Interpretation

In addition to the issue of the potential endogeneity of events that we discussed in Section 3.3.3, there is at least one other issue that we must address in order to lend further support to our results and the interpretation that we have given them. 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 aggregate-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 aggregate 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.

Figure 5: Estimated Probabilities of Announcement Types as a Function of Δstocks

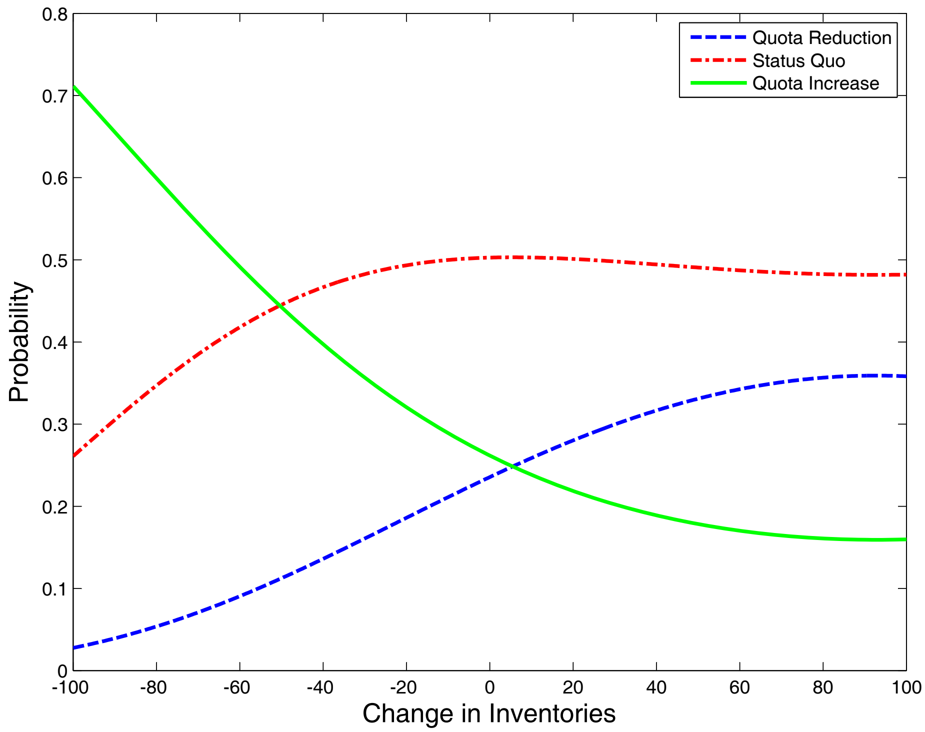

In our opinion, it is difficult to think of a variable which, unrelated to OPEC's bargaining problem, satisfies all of these criteria. Three possibilities emerge, though we believe that all are lacking in at least one dimension. The most obvious is that of outside political pressure. Following a negative 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 changes in behaviour by consuming nations or western oil companies. However, rather than exposing a flaw in our argument, we believe that this possibility points to a specific dimension over which OPEC must bargain — how accommodative to be towards consuming nations in any quota decisions.

Consider next non-OPEC supply shocks. 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. Finally, consider cost shocks to OPEC producers. Such shocks may well be unobservable but given the discussion in the introduction, such a model of random costs would lead to very different predictions about market share fluctuations.

3.6 Anecdotal Evidence

There is plenty of anecdotal evidence which suggests that OPEC is not only particularly well-placed to provide a test of our theoretical model, but also that the model and the empirical evidence suggested thus far actually have merit. For example, regarding OPEC's June 1988 production announcement maintaining the status quo, the New York Times reported[16] that both Iraq and the United Arab Emirates exempted themselves from their individual quota obligations and that they did so despite the fact that oil prices were low and overproduction was widespread. As another example, 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.[17] This example also illustrates that OPEC's decision making is often all-or-nothing; OPEC could have both lowered production minimally and accommodated Kuwait's desire for greater post-war production, which would have been surplus maximising; however, OPEC chose instead to maintain the (inefficient) status quo, rather than let one country gain disproportionately.

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. In November 1997, OPEC apparently misjudged the severity of the developing crisis and actually increased production 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."[18] However, even after the Asian Financial Crisis, this paper's central thesis is still supported in the popular press, with one article noting that "[t]he uncertainty surrounding OPEC policies underscores the cartel's difficulties in coping with a slowing market after three years of rising prices and surging demand. Analysts view the group as too divided to be effective".[19]

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

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

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

[19] See the article, "With Markets Doubting OPEC Cut, Oil Is Lowest in a Year," New York Times, 17 November 2006, C4.

4 Discussion & Concluding Remarks

In this paper we have presented the results of an event study analysis of OPEC behaviour and its effects on both oil prices and stock returns in the oil industry. The results are quite clear. When OPEC reduces the aggregate quota, positive and often significant abnormal returns accrue to oil prices and in the three sub-sectors considered. Next, when the aggregate 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, during good times, OPEC is far more likely to increase the aggregate quota than they are to decrease it during bad times. This interpretation is given further credence when we correlate the outcome of OPEC meetings with changes in residual demand for oil and changes in inventories. In addition, that analysis also showed that OPEC is more likely to agree the larger is the shock (either positive or negative). Finally, our interpretation is further supported when one examines the news reports surrounding these status quo announcements.

As we argued in the Introduction, the standard oligopoly models must be cast aside. Therefore, to explain our results we adopt a bargaining approach and develop a simple model of bargaining with one-sided private information that neatly captures our empirical results. In the model, the private information creates a wedge between the interests of the proposer and the responder. Disagreement then occurs for two distinct reasons: symmetric uncertainty, as captured by the parameter α, and asymmetric uncertainty, as captured by the parameter β. The former simply captures the idea that cartel members face shocks to their cost of production or to other economic and political factors which affect their ideal quantity, while the latter introduces an asymmetry between the two cartel members such that one prefers, on average, higher production than the other. When $\beta > 0$, disagreements are much more likely during periods of low demand than during periods of high demand.

References

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

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

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

[4] Cramton, Peter and Thomas Palfrey, *Ratifiable Mechanisms: Learning From Disagreement.* Games and Economic Behavior 10 (1995), 255283.

[5] Cramton, Peter and Thomas Palfrey, *Cartel Enforcement with Uncertainty about Costs.* International Economic Review 31 (1990), 17-47.

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

[7] Hyndman, Kyle and Debraj Ray, *Coalition Formation With Binding Agreements.* The Review of Economic Studies, *forthcoming*.

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

[9] 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

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

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

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

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

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

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

[16] Scott Morton, Fiona, *Entry and Predation: British Shipping Cartels 1879 – 1929.* Journal of Economics & Management Strategy 6(4) (1997) pp. 679-724.

[17] 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.

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

APPENDIX A: THE DYNAMIC MODEL

In this appendix, we show that the central intuition and predictions of the static model continue to hold in a dynamic model with Markovian demand shocks.

We take the static game of Section 2 as our stage game. To it we add the following. Demand in period t, indexed by A_t, takes one of two values: $\overline{A} > \underline{A}$ and follows a symmetric Markov process where the probability of a state transition in each period is given by $\mu \in (0,1)$. We make three further assumptions on the structure of the model. First, we assume that both firms have a common discount factor $\delta \in (0,1)$. Second, we assume that firm 2's type η_t in period t is privately drawn i.i.d. each period from a discrete version of F_η. Finally, we assume that production can only be adjusted in increments of $\Delta > 0$ on the bounded interval $[0, \overline{A}]$. Thus we have a finite number of states, which can be written as $S_t = (A_t, Q_{0t}, \eta_t)$, where A_t is the level of demand, Q_{0t} is the initial quota and η_t is the type of firm 2.[20]

With the basic structure now at hand, we can now describe strategies and equilibrium; we restrict attention to stationary Markov strategies and also Markov perfect equilibria. A strategy for firm 1 is a feasible proposal $Q(Q_{0t}, A_t)$, while a strategy for firm 2 is a behaviour strategy $\Theta(A_t, \eta_t, Q_{0t}, Q) \in [0,1]$, indicating whether the proposal is accepted or rejected. Using arguments similar to those found in Hyndman and Ray (2007), a stationary Markov equilibrium can be shown to exist.

Remark 1. *While we are assured that an equilibrium exists, explicitly characterising it is exceedingly difficult. Instead, we solve for the stationary equilibrium numerically. In particular, we will solve a finite version of the model and take T sufficiently high in order to approximate the true model. Specifically, we take the smallest T such that $\delta^T < 10^{-4}$ so that the final period has a negligible effect on payoffs.*

To see what happens, first take a look at Figures 6 and 7, which depicts the specific case $(\alpha, \beta, \delta, \mu, \overline{A}, \underline{A}) = (2, 1.5, 0.4, 0.1, 15, 10)$. Notice that for both the high and low-demand states, for some initial production levels there is a strictly positive probability that the offer will be rejected, and so, on the surface, it is not clear that agreements will be more likely when an increase is proposed than when a decrease is proposed. However, we can say more. In the high-demand state, for many initial production levels, but in particular for initial quantities, Q_{0t} in the neighbourhood of 5 (*i.e.,* $\frac{\underline{A}}{2}$), firm 1 proposes that production be increased to 7.5 (*i.e.,* $\frac{\overline{A}}{2}$) and the offer is accepted by all responder types. In contrast, in the low-demand state, when $Q_{0t} \approx \frac{\overline{A}}{2}$ there are two distortions. First, firm 1 recognises that the likelihood of a proposal all the way to $\frac{\underline{A}}{2}$ would be rejected with high probability. Therefore,

[20] Of course, only $S'_t = (A_t, Q_{0t})$ is observable to firm 1.

her proposal is less ambitious — proposing only a reduction to 5.5. Second, even with a higher proposal, some responder types will reject the offer. Therefore, the firms will reach a steady state in which production always falls between $[\frac{\underline{A}}{2}, \frac{\overline{A}}{2}]$. In this steady state, when demand is low a proposal to reduce output is probabilistically rejected, while when demand is high the proposal is immediately accepted. Note that in both high- and low-demand states, both firms earn greater than competitive profits; however, in the high-demand state they are able to coordinate on the monopoly quantity, while in the low-demand state, they tend to over-produce, as well as sometimes fail to reduce the output.

FIGURE 6: Equilibrium Proposals Given Demand & Initial Quota

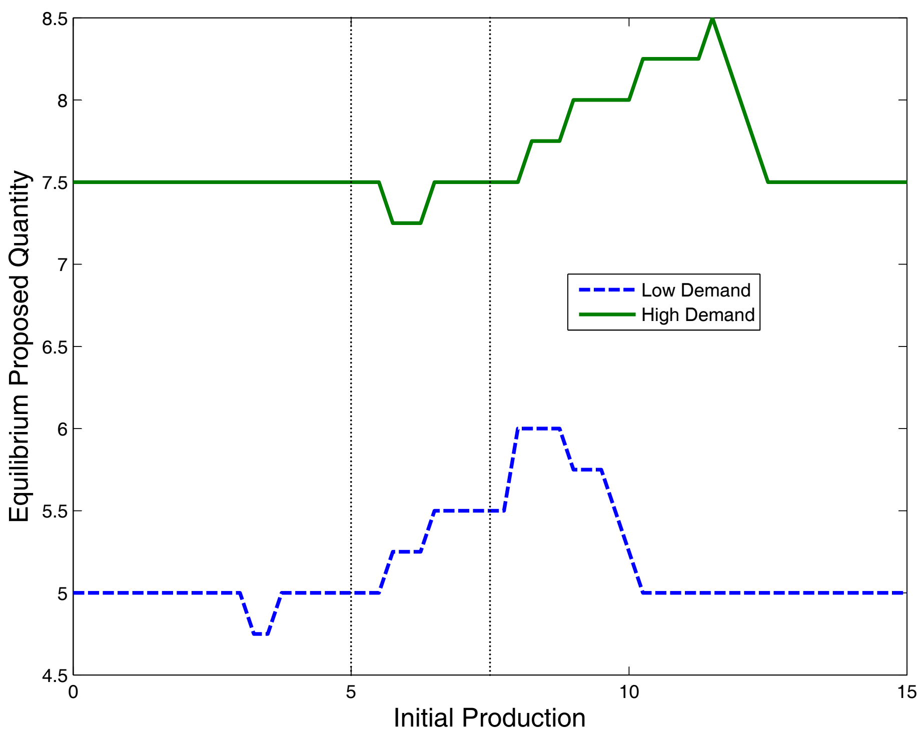

Of course, when moving to a dynamic framework, things do change, as can be seen from Table 2 where we examine equilibrium offers and agreement rates for many different values of the exogenous parameters. The first four columns contain specific values of exogenous parameters. The next two show the equilibrium probability of agreement for two important initial production levels. In particular, $\Pr[\frac{\overline{A}}{2}|\underline{A}]$ is the probability that the equilibrium offer of $\frac{\overline{A}}{2}$ will be accepted conditional on being at the low-demand state with an initial quantity of $Q_{0t} = \frac{\underline{A}}{2}$. Similarly, $\Pr[\frac{\underline{A}}{2}|\overline{A}]$ is the probability that an offer of the equilibrium offer will be accepted conditional on being at the high-demand state with an initial quantity of $Q_{0t} = \frac{\overline{A}}{2}$. The final two columns show equilibrium offers in the low and high-demand states at those

Figure 7: Equilibrium Acceptance Probabilities Given Demand & Initial Quota

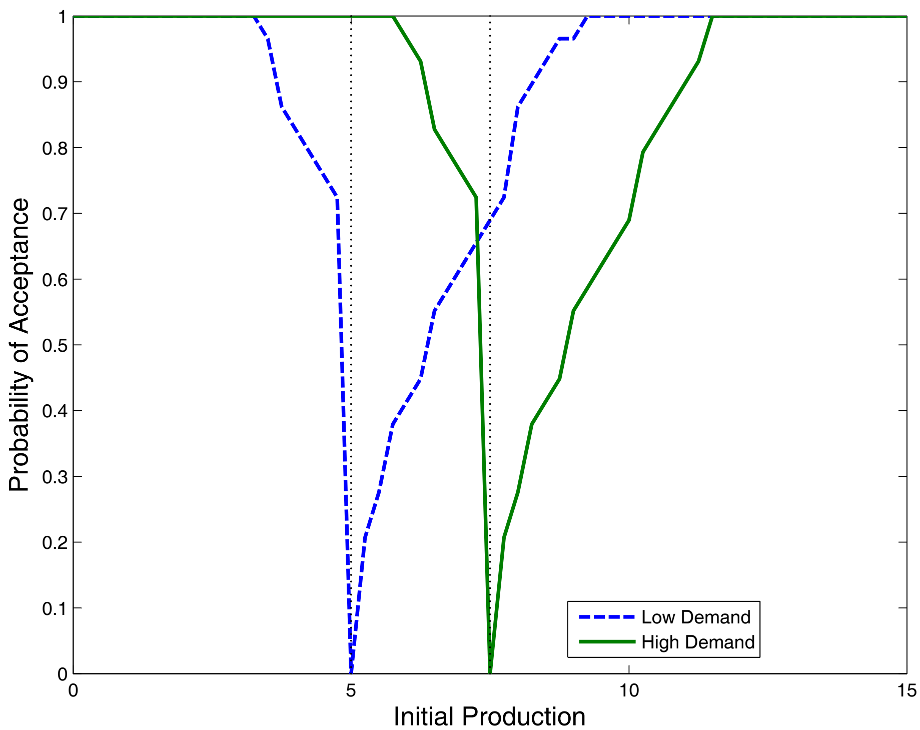

same two initial quantities. Importantly, as with the case depicted in Figures 6 and 7, there will be a steady state distribution of initial quantities, which is a strict subset of the interval $[\frac{\underline{A}}{2}, \frac{\overline{A}}{2}]$.

Except for the final three rows, all cases depicted in the table have $\overline{A} = 15 > 10 = \underline{A}$. In the final three rows, we leave $\underline{A}$ unchanged but reduce $\overline{A}$ to 12, thereby making the demand fluctuations smaller. Lowering the size of the demand shock has two main effects. First, in the low-demand state, firm 1 is more demanding in her offer. Second, and as a consequence of the first, the probability of agreement is now substantially lower in both the low- and high-demand states. The intuition for this is quite simple: with a smaller demand shock, the gain from agreement (or the cost to disagreement) is much smaller and the interests of the proposer and many responder types are more likely to diverge. While having a simple intuition, it is an important testable implication of our model: *the larger the size of the demand shock, the more likely are agreements to change output.*

Table 2 also makes clear that disagreements are always strictly more likely to occur in the low-demand state than in the high-demand state. The model makes many other predictions, but we will not dwell too much on them here. First, as μ increases, so that the likelihood of

TABLE 2: **Acceptance Probabilities With Different Parameter Values**

Exogenous Parameters				Acceptance Prob.		Equilibrium Offers	
α	β	δ	μ	$\Pr[\frac{\overline{A}}{2}\|\underline{A}]$	$\Pr[\frac{\underline{A}}{2}\|\overline{A}]$	$Q(\frac{\overline{A}}{2},\underline{A})$	$Q(\frac{\underline{A}}{2},\overline{A})$
1	1.5	0.7	0.5	0.857	1	5.75	7.50
1	2.0	0.7	0.5	0.640	1	5.75	7.25
1	2.5	0.7	0.5	0.552	1	5.75	7.25
2	1.5	0.7	0.5	0.690	1	5.50	7.25
2	2.0	0.7	0.5	0.576	1	5.50	7.25
2	2.5	0.7	0.5	0.487	1	5.50	7.25
3	1.5	0.7	0.5	0.676	1	5.75	7.25
3	2.0	0.7	0.5	0.537	1	5.50	7.25
3	2.5	0.7	0.5	0.444	1	5.50	7.00
2	2.0	0.4	0.5	0.606	1	5.50	7.25
2	2.0	0.9	0.5	0.515	1	5.75	7.00
2	2.0	0.7	0.1	0.576	1	5.50	7.50
2	2.0	0.7	0.9	0.576	1	5.75	7.00
2	1.5	0.4	0.1	0.690	1	5.50	7.50
2	1.5	0.4	0.9	0.690	1	5.50	7.25
2	0	0.7	0.1	1	1	5.00	7.50
2	0.25	0.7	0.1	1	1	5.00	7.50
2*	1.0	0.7	0.1	0.440	0.84	5.25	6.00
2*	1.0	0.7	0.9	0.480	0.96	5.25	5.75
2*	1.0	0.4	0.1	0.520	0.84	5.25	6.00

* Here $\overline{A} = 12$ and $\underline{A} = 10$. In all other cases, $\overline{A} = 15$ and $\underline{A} = 10$.

a state transition increases, firm 1 becomes strictly less demanding in her proposals, and the probability of disagreement increases in the low-demand states. Thus volatility is efficiency destroying on two fronts. Second, as the degree of uncertainty (represented by α) or the degree of asymmetry (represented by β) increases, disagreements also become more likely. Finally, as firms become more patient, the probability of disagreement also rises.

APPENDIX B: FIGURES & TABLES

FIGURE 8: Cumulative Abnormal Return Functions (All Meetings)

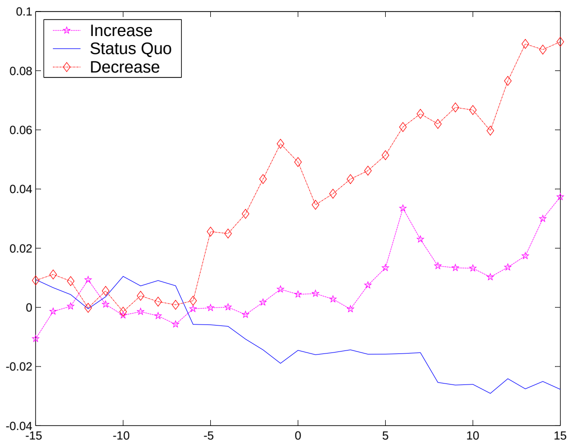

(a) The Spot Price of Oil

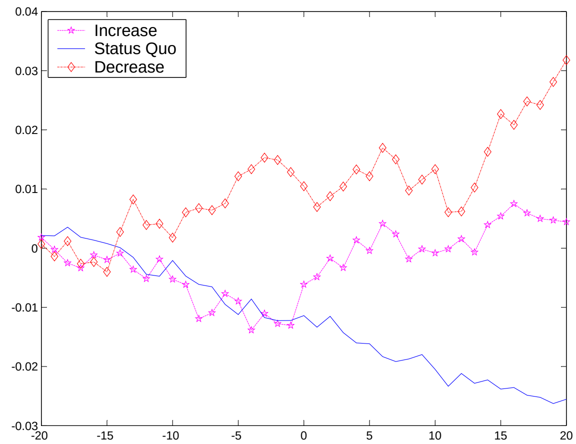

(b) All Oil Companies

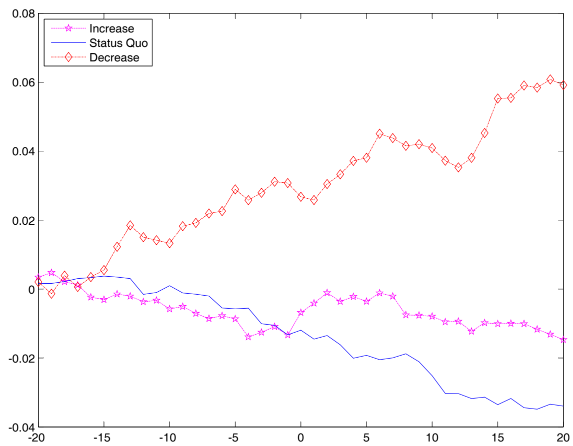

(c) Exploration & Drilling Companies

Figure 9: Cumulative Abnormal Return Functions (Normal Meetings)

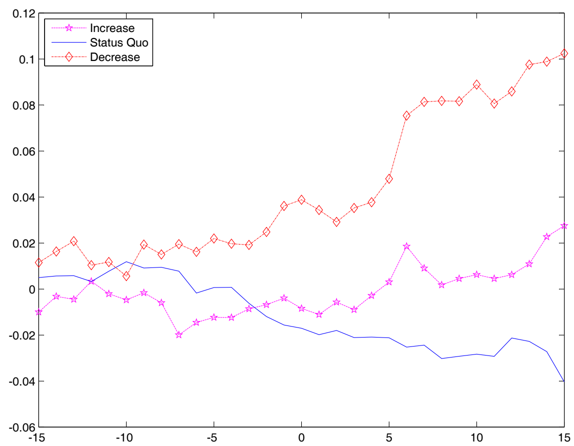

(a) The Two-Month Forward Price of Oil

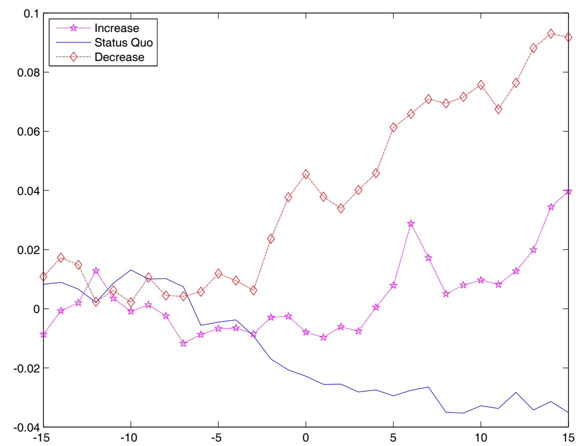

(b) The Spot Price of Oil

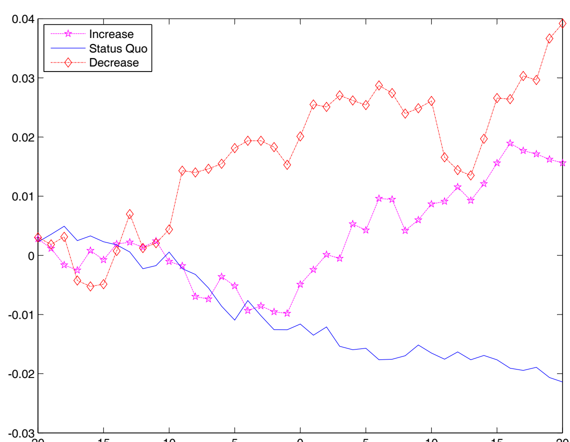

(c) All Oil Companies

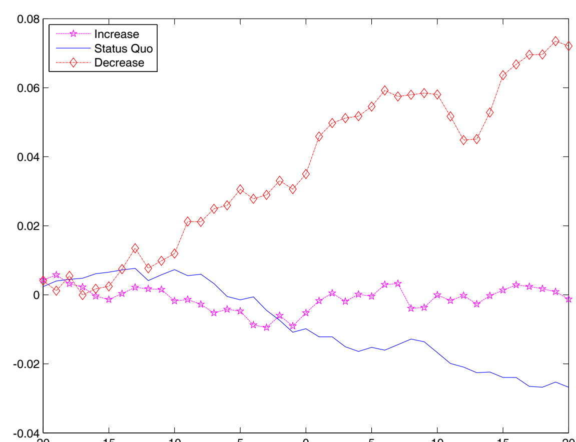

(d) Exploration & Drilling Companies

TABLE 3: **Summary of OPEC Meetings; August 1986 – September 2002**

Start Date	End Date	Meeting[♭]	Δ Quota[†]	Δ_{12} Invent[‡]	Δ_{12} Resid Dem[♯]	Δ_{12} Price
28-Jul-86	5-Aug-86	N	0	211.0	233.0	-15.60
6-Oct-86	22-Oct-86	E	161 - 239	237.0	1169.0	-13.80
10-Dec-86	20-Dec-86	N	-705	134.0	1701.0	-14.55
25-Jun-87	27-Jun-87	N	800	45.0	1740.0	5.80
9-Dec-87	14-Dec-87	N	0	56.0	414.0	3.20
11-Jun-88	14-Jun-88	N	0	135.0	339.0	-2.00
21-Nov-88	28-Nov-88	N	800	-32.0	3650.0	-6.15
5-Jun-89	7-Jun-89	N	1000	0.0	1465.0	2.20
22-Sep-89	28-Sep-89	MMC	1000	59.0	0.0	3.75
21-Nov-89	28-Nov-89	N	1586	89.0	-297.0	6.55
	4-May-90	?	0	157.7	990.7	-1.80
26-Jul-90	27-Jul-90	N	405	124.4	3356.0	-3.40
12-Dec-90	13-Dec-90	N	0	92.4	-2889.9	9.40
11-Mar-91	12-Mar-91	MMC	-193	-21.2	-566.0	-1.85
4-Jun-91	4-Jun-91	N	0	-73.3	940.4	3.65
24-Sep-91	25-Sep-91	MMC	1352	26.4	946.3	-5.20
26-Nov-91	27-Nov-91	N	0	55.9	2086.2	-11.35
12-Feb-92	15-Feb-92	MMC	-668	47.2	2197.2	-2.40
21-May-92	22-May-92	N	0	-1.2	850.6	-0.40
16-Sep-92	17-Sep-92	MMC	1218	-56.0	431.5	-0.55
25-Nov-92	27-Nov-92	N	382	-10.2	1854.0	-3.02
13-Feb-93	16-Feb-93	MMC	-1000	66.0	1329.9	1.34
8-Jun-93	10-Jun-93	N	0	89.2	701.5	-1.81
25-Sep-93	29-Sep-93	E	932	114.7	931.9	-3.79
23-Nov-93	24-Nov-93	N	0	72.9	1218.0	-3.31
25-Mar-94	26-Mar-94	MMC	0	-22.6	-2.6	-5.82
21-Nov-94	22-Nov-94	N	0	32.2	-845.6	1.21
19-Jun-95	20-Jun-95	N	0	5.5	326.1	0.57
21-Nov-95	22-Nov-95	N	0	-76.7	827.5	-0.91
5-Jun-96	7-Jun-96	N	513	-61.2	-1726.5	0.93
27-Nov-96	28-Nov-96	N	0	-69.4	-684.1	5.28
25-Jun-97	26-Jun-97	N	0	233.3	5568.8	1.25
26-Nov-97	1-Dec-97	N	2467	284.2	6037.6	-5.99
30-Mar-98	30-Mar-98	E	-1245	105.5	1656.1	-4.91
24-Jun-98	24-Jun-98	N	-1868	173.6	197.8	-6.10
	30-Oct-98	?	0	155.9	-104.8	-5.53
25-Nov-98	26-Nov-98	N	0	124.4	875.5	-6.54
23-Mar-99	23-Mar-99	N	-1411	70.9	2646.9	-3.11
27-Mar-00	29-Mar-00	N	1500-1650	-279.3	-2212.1	19.43
21-Jun-00	21-Jun-00	E	800-1000	-233.6	-1170.8	13.85
10-Sep-00	11-Sep-00	N	800	-171.8	-273.5	11.64
30-Oct-00	30-Oct-00	N	500	-165.6	-399.6	7.50
17-Jan-01	17-Jan-01	E	-1499	23.2	1468.0	-0.34
16-Mar-01	17-Mar-01	N	-1000	71.8	-878.4	-3.86
3-Jul-01	3-Jul-01	E	0	9.7	115.0	-7.05
25-Jul-01	25-Jul-01	NM	-1000	9.7	115.0	-7.05
26-Sep-01	27-Sep-01	N	0	90.6	-2296.2	-6.18

Continued on Next Page

Table 3 – **continued from previous page**

Start Date	End Date	Meeting[♭]	Δ Quota[†]	Δ_{12} Invent[‡]	Δ_{12} Resid Dem[♯]	Δ_{12} Price
14-Nov-01	14-Nov-01	E	-1501	85.9	-532.7	-12.72
28-Dec-01	28-Dec-01	E	0*	113.0	-2842.7	-11.87
15-Mar-02	15-Mar-02	N	0	106.4	-1348.2	-5.49
26-Jun-02	26-Jun-02	E	0	92.5	-1576.7	-2.67
19-Sep-02	19-Sep-02	N	0	-44.7	114.0	1.80

[♭] N: normal meeting; E: extraordinary meeting; MMC: Ministerial Monitoring Committee Meeting.
[†] Δ Quota measured in thousands of barrels per day.
[‡] Δ_{12} Inventories measured in millions of barrels.
[♯] Δ_{12} Residual Demand measured in thousands of barrels per day.
[*] Special meeting where previously announced *conditional* reduction is implemented.

TABLE 4: **Cumulative Abnormal Returns For The Two-Month Forward Price of Oil**

Event Day	Q↑ $\overline{CAR}$	Q↑ s.e.	Q↔ $\overline{CAR}$	Q↔ s.e.	Q↓ $\overline{CAR}$	Q↓ s.e.
-15	-0.0109	0.0048	0.0078	0.0047	0.008	0.0078
-14	-0.0035	0.0068	0.0059	0.0067	0.0107	0.0111
-13	-0.0052	0.0083	0.0045	0.0082	0.0128	0.0136
-12	0.0014	0.0096	0.0011	0.0096	0.0057	0.0157
-11	-0.0032	0.0108	0.0041	0.0107	0.0086	0.0176
-10	-0.0054	0.0119	0.0098	0.0118	0.0031	0.0194
-9	-0.0037	0.0128	0.0058	0.0128	0.01	0.021
-8	-0.0062	0.0138	0.0079	0.0137	0.009	0.0225
-7	-0.0144	0.0147	0.0065	0.0146	0.0091	0.024
-6	-0.0081	0.0155	-0.0041	0.0154	0.0082	0.0254
-5	-0.008	0.0163	-0.0041	0.0162	0.0253	0.0267
-4	-0.0081	0.0171	-0.0045	0.017	0.0274	0.0279
-3	-0.0065	0.0178	-0.011	0.0177	0.0359	0.0292
-2	-0.0051	0.0186	-0.0141	0.0184	0.0396	0.0304
-1	0.0008	0.0193	-0.0186	0.0191	0.0483	0.0315
0	-0.0011	0.02	-0.0139	0.0198	0.0398	0.0327
1	-0.0011	0.0207	-0.0136	0.0205	0.0259	0.0338
2	0.0001	0.0213	-0.0106	0.0212	0.0301	0.0349
3	-0.0042	0.022	-0.0096	0.0218	0.0346	0.0359
4	0.0023	0.0226	-0.0108	0.0224	0.0405	0.037
5	0.0066	0.0232	-0.0098	0.0231	0.0445	0.038
6	0.022	0.0238	-0.0134	0.0237	0.0698	0.039
7	0.0132	0.0244	-0.0136	0.0243	0.0752	0.04
8	0.0087	0.025	-0.0222	0.0249	0.0713	0.041
9	0.0083	0.0256	-0.0212	0.0255	0.0725	0.0419
10	0.0088	0.0262	-0.0203	0.026	0.0723	0.0429
11	0.0058	0.0268	-0.0229	0.0266	0.0666	0.0438
12	0.0068	0.0274	-0.016	0.0272	0.0807	0.0448
13	0.0089	0.0279	-0.0156	0.0277	0.0936	0.0457
14	0.0195	0.0285	-0.0176	0.0283	0.0916	0.0466
15	0.026	0.029	-0.0269	0.0288	0.0946	0.0475

Table 5: **Cumulative Abnormal Returns For The Spot Price of Oil**

Event Day	Q↑ $\overline{CAR}$	Q↑ s.e.	Q↔ $\overline{CAR}$	Q↔ s.e.	Q↓ $\overline{CAR}$	Q↓ s.e.
-15	-0.0106	0.0052	0.0094	0.0053	0.0091	0.0084
-14	-0.0014	0.0074	0.0066	0.0075	0.0111	0.0119
-13	0.0003	0.0091	0.0043	0.0092	0.0089	0.0147
-12	0.0093	0.0105	-0.0005	0.0106	-0.0002	0.017
-11	0.001	0.0118	0.0036	0.0119	0.0056	0.019
-10	-0.0027	0.0129	0.0105	0.0131	-0.0015	0.0209
-9	-0.0015	0.014	0.0072	0.0142	0.0039	0.0227
-8	-0.0029	0.015	0.0091	0.0152	0.0019	0.0243
-7	-0.0057	0.016	0.0073	0.0162	0.0009	0.0259
-6	-0.0005	0.0169	-0.0058	0.0171	0.0022	0.0274
-5	-0.0002	0.0178	-0.0059	0.018	0.0256	0.0288
-4	0.0001	0.0186	-0.0064	0.0189	0.0249	0.0302
-3	-0.0025	0.0195	-0.0107	0.0197	0.0316	0.0315
-2	0.0017	0.0202	-0.0144	0.0205	0.0434	0.0328
-1	0.0061	0.021	-0.0189	0.0213	0.0553	0.034
0	0.0044	0.0218	-0.0145	0.0221	0.0491	0.0352
1	0.0047	0.0225	-0.016	0.0228	0.0347	0.0364
2	0.0028	0.0232	-0.0153	0.0236	0.0384	0.0376
3	-0.0006	0.0239	-0.0144	0.0243	0.0433	0.0388
4	0.0075	0.0246	-0.0158	0.025	0.0462	0.0399
5	0.0134	0.0253	-0.0158	0.0257	0.0514	0.041
6	0.0335	0.026	-0.0156	0.0263	0.0609	0.0421
7	0.023	0.0267	-0.0153	0.027	0.0654	0.0431
8	0.014	0.0273	-0.0254	0.0277	0.062	0.0442
9	0.0134	0.028	-0.0263	0.0283	0.0676	0.0452
10	0.0132	0.0286	-0.026	0.029	0.0667	0.0463
11	0.0102	0.0292	-0.0291	0.0296	0.0597	0.0473
12	0.0136	0.0298	-0.0241	0.0302	0.0766	0.0483
13	0.0174	0.0304	-0.0276	0.0309	0.0891	0.0493
14	0.03	0.0311	-0.025	0.0315	0.0871	0.0503
15	0.0373	0.0317	-0.0277	0.0321	0.0898	0.0512

Table 6: **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 7: **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