

DID IRAQ CHEAT THE UNITED NATIONS? UNDERPRICING, BRIBES, AND THE OIL FOR FOOD PROGRAM*

CHANG-TAI HSIEH AND ENRICO MORETTI

From 1997 through 2003, the UN Oil for Food Program allowed Iraq to export oil for humanitarian supplies. We hypothesize that Iraq deliberately set the price of its oil below market prices to solicit bribes from oil buyers. By comparing the price gap between Iraqi oil and its close substitutes during the Program to the gap prior to the Program, we find evidence of significant underpricing. Our central estimate suggests that Iraq collected \$1.3 billion in bribes from underpricing its oil, or 2 percent of oil revenues. Underpricing is higher during periods of high volatility in oil markets—when detection is more difficult—but declines after the UN limited Iraq's ability to set the price of its oil.

I. INTRODUCTION

In response to Iraq's invasion of Kuwait, the United Nations imposed an embargo on exports of Iraqi oil. To alleviate the economic hardship caused by the embargo, the UN created the Oil for Food program that allowed Iraq to export oil in exchange for humanitarian goods, largely food and medicine. The proceeds from the oil exports had to be deposited in an escrow account controlled by the UN and could only be used to purchase humanitarian supplies.

By some measures, the Oil for Food program was the largest humanitarian relief program in world history. By the time the program ended in 2003, the Oil for Food program had provided \$32 billion in humanitarian supplies to Iraq. This makes the Oil for Food program about 40 percent of the size of the Marshall Plan in absolute terms, and almost five times larger than the Marshall Plan on a per capita basis. There is evidence that the humanitarian supplies provided by the program significantly improved the standards of living of Iraq's civilian population (especially children) during a period of severe need.¹

* We thank Steven Davis, Stefano DellaVigna, Raymond Fisman, Richard Gilbert, John Fawcett, Frank Hydoski, Alan Krueger, Charles Jones, Maurice Lorenz, Edward Miguel, Mary Jane Schirber, and three referees for helpful comments. We are grateful to Eric Chaney, Emily Conover, Daniel Hartley, Bing-ru Teh, and Matthew Wiswall for superb research assistance. Hsieh thanks the Alfred P. Sloan Foundation for financial support.

1. For example, during the years of the program, chronic malnutrition among children and the overall incidence of vaccine-preventable diseases declined sharply [FAO and WFP 2003].

© 2006 by the President and Fellows of Harvard College and the Massachusetts Institute of Technology.

The Quarterly Journal of Economics, November 2006

While an important goal of the Oil for Food program was to provide humanitarian assistance, the UN was also keen to prevent Iraq from obtaining resources that might be used to purchase weapons or luxury goods. However, two key features of the Oil for Food program potentially made it possible for Iraq to extract cash bribes and political favors from the oil buyers. First, Iraq could freely choose the buyers of Iraqi oil. Second, until September 2001, Iraq had some discretion over the selling price of its oil. Therefore, particularly during periods of high oil market volatility, Iraq potentially was able to deliberately set the price of its oil below market prices and then choose buyers that were willing to pay bribes to obtain the underpriced oil.

The objective of this paper is to measure these bribes. Clearly, because of its illicit nature, finding direct evidence of the alleged cash payments and political favors—a paper trail, for example—is difficult. On the basis of documents provided by Iraq’s oil company, the CIA, and a commission headed by Paul Volcker estimate that Iraq received \$230 million in *cash* bribes through the strategic underpricing of its oil [CIA 2004; IIC 2005b]. However, this figure may be an underestimate since it does not capture the monetary value of the political favors provided by the oil traders (or their associates) in exchange for the underpriced oil, nor does it include bribes that might have been paid to Iraqi officials outside the oil company (for example, to Saddam Hussein).

We take a different approach. Instead of looking for a paper trail, we use publicly available data to uncover patterns suggestive of bribery.² Specifically, we measure the potential bribe by the gap between the official selling price of Iraqi oil (i.e., the price paid to the Oil for Food Program) and our estimates of the “market” price of Iraqi oil (i.e., the price of Iraqi oil at the destination port, typically in Europe and the United States, after subtracting shipping and insurance costs). We then compare this gap with that observed before the program. If Iraqi oil was underpriced in an effort to collect bribes, we expect to see two features in the data. First, we expect to see a larger gap between the “market” price and the official price of Iraqi oil during the Oil for Food program years than before the Program. In the years

2. Our analysis follows the approach taken by Di Tella and Schargrodsky [2003], Fisman and Wei [2004], Olken [2004b], and Reinikka and Svensson [2004] in using patterns in prices and quantities to make inferences about corruption.

before the Oil for Food program, Iraq had full control over its oil revenues and thus did not have any incentive to underprice. Second, we expect the price difference to become smaller and return to its historical levels after September 2001. The reason is that in September 2001, the UN adopted a “retroactive” pricing system under which Iraq was forced to set its selling price retroactively based on the price actually observed in the market.

To illustrate our evidence, consider the gap between the official selling price of Basrah Light—one of the two Iraqi oils—and the price of Arabian Light in the top panel of Figure I. Arabian Light and Basrah Light are chemically very similar and are considered close substitutes. The price gap indicates that underpricing of Basrah Light was modest in the first few years of the program, reached a peak in the 2000–2001 period, and declined after the adoption of retroactive pricing in the fall of 2001. Notably, we find a qualitatively similar pattern when we measure the gap between official selling price of Iraqi oil and the spot market price of Iraqi oil itself (after deducting shipping costs). This gap reflects the difference between the price of Iraqi oil when the oil is first sold at the Iraqi port under UN supervision and the spot price of the *same* oil when it reaches its destination and is sold to a refiner in Europe or the United States. The price paid to oil traders in these spot market transactions presumably includes the cost of the bribe paid to the Iraqi regime. Unlike the evidence based on competitors, this gap cannot be explained by unobserved shocks to the relative demand or relative supply of Iraqi oil.

Furthermore, consistent with the notion that price volatility in oil markets made it more difficult for the UN to detect underpricing, we find that underpricing was positively correlated with volatility in the period before retroactive pricing. In contrast, in the period after retroactive pricing, when volatility should not matter, we find no relationship between underpricing and price volatility. In addition, we find that underpricing was associated with changes in the composition of oil buyers. While a multinational oil company might be reluctant to pay bribes, a small oil trader operating in Russia or Switzerland might be more pliant. We find a positive correlation between the degree of underpricing and the fraction of oil sold to individual oil traders (rather than to large oil multinationals).

On the basis of this evidence, we interpret the product of the change in the price gap (relative to the price gap prior to the program) and the quantity of oil sold as an estimate of the rent

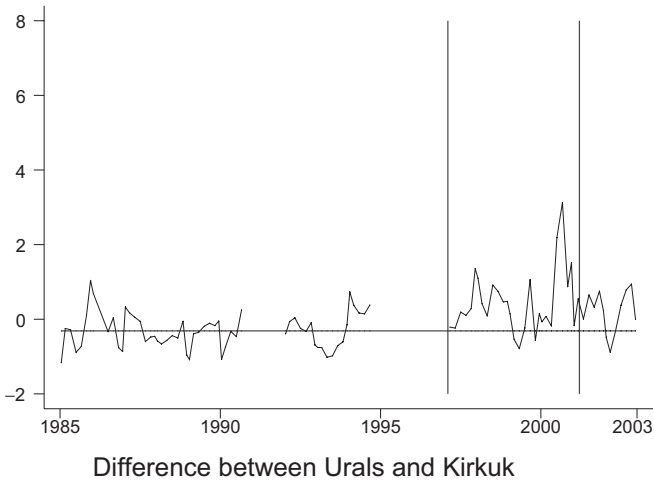



FIGURE I
Difference between the Market Price of Close Substitutes
and the Official Selling Price of Iraqi Oils

The top panel shows the difference between the market price of Arabian Light and the official selling price of Basrah. Arabian Light is the closest substitute of Basrah. The bottom panel shows the difference between the market price of Urals and the official selling price of Kirkuk. Urals is the closest substitute of Kirkuk. The first vertical line marks the beginning of the Oil for Food Program. The second vertical line indicates the beginning of retroactive pricing. The horizontal line is the average difference for the years before the Oil for Food Program. Iraqi oil was not traded in 1991–1993 due to the first Gulf War, and in 1995–1997 due to delays in the Oil for Food Program.

generated by this strategic underpricing scheme. Our best estimate is that Iraq generated \$3.5 billion in rents by deliberately underpricing its oil. These rents were presumably shared between the oil buyers and the Iraqi government (in the form of favors and cash bribes). We use a simple model of monopolistic competition among oil traders and data on the standard markup of oil traders to estimate the amount of bribes collected by Iraq. Our best estimate suggests that Iraq collected 1.3 billion dollars in bribes. This amounts to 2 percent of the total value of oil sold under the Oil for Food program, which does not appear to be very large when compared with the typical amount of corruption that we observe in many projects in developing countries.³

The rent-seeking mechanism we examine in this paper is not limited to Iraq, but it is something we see in other countries, including the United States. For example, there is some evidence that energy companies extracting oil and natural resources from public land reported to the U. S. government that they received lower prices than the actual market price, thus depriving the U. S. Treasury of \$700 million in royalties in 2005 (*New York Times*, 1/23/2006). This underpricing of oil is remarkably similar to what we saw in Iraq during the Oil for Food program. Historically, there have also been many cases in which public resources were sold by U. S. government officials to private firms at below market prices in exchange for bribes in the early twentieth century [Glaeser 2004].

The experience of the Oil for Food program also holds important lessons for other humanitarian assistance programs. For example, Chad recently reached an agreement with a consortium of oil companies led by ExxonMobil to develop its oil industry. To prevent the oil revenues from being misused, Chad pledged that its revenues from the project were to be deposited into an escrow account controlled by the World Bank. The program is remarkably similar to the Oil for Food Program, and the World Bank is considering extending the same model to other developing countries. The experience of the Oil for Food program highlights the importance of institutional design in preventing initiatives such as the ExxonMobil-Chad program from resulting in rent-seeking behavior. For example, in the case of the Oil for Food Program, a

3. It is of course possible that Iraq was also obtaining illegal resources by exploiting the program in other ways. For example, Iraq could have collected bribes by paying above market prices for the humanitarian goods purchased by the Oil for Food program.

relatively minor change in design (the introduction of retroactive pricing) appears to have dramatically curtailed the scope for bribes.

Finally, we want to stress that there is evidence that the Oil for Food program provided food and medicines to Iraqi civilians in a period of acute need, and reduced malnutrition and malaria rates among children. Our paper does not measure the benefits of the humanitarian assistance provided by the Oil for Food Program, and therefore does not speak to the issue of whether the program overall was successful.

II. THEORY: BRIBES AND UNDERPRICING OF STATE-SOLD GOODS

In this section we sketch a simple model to illustrate the incentive for Iraq to underprice its oil. The model is borrowed from Shleifer and Vishny's [1992] analysis of bribes and underpricing in state owned firms in the former Soviet Union. Glaeser [2004] has a similar model, applied to the sale of public assets at below-market prices by local governments in the United States. Suppose that the inverse demand curve for Iraqi oil is represented by $D(Q)$, the cost of producing oil by $C(Q)$, and the official price of Iraqi oil by P^{official} . The key assumption is that Iraq does not keep any of the official revenues and places no value on the food and medicine purchased by the Oil for Food Program. Therefore, instead of maximizing its profits, Iraq maximizes profits *net of the official revenues*, or $D(Q) \times Q - C(Q) - P^{\text{official}} \times Q$. Assuming for now that P^{official} is exogenous and that bribery is not penalized, the first-order condition is that the marginal revenue is equal to the sum of the marginal cost and the official price: $R'(Q^*) = C'(Q^*) + P^{\text{official}}$. This equilibrium is represented in Figure II, where P^{market} represents the market price, Q^* the equilibrium quantity, and the shaded area represents the rent created by underpricing.

What happens if the official price is not exogenous, but is endogenously set by Iraq? It is easy to see that Iraq would benefit from having the official price set as low as possible. In the absence of any constraints on the official price, Iraq would set $P^{\text{official}} = 0$. In reality, the constraint that prevented Iraq from setting this price as low as it would have liked was that the official price had to be approved by the UN at the beginning of each month. The UN tried to make sure that the price was set as close as possible to the expected market price over the month.

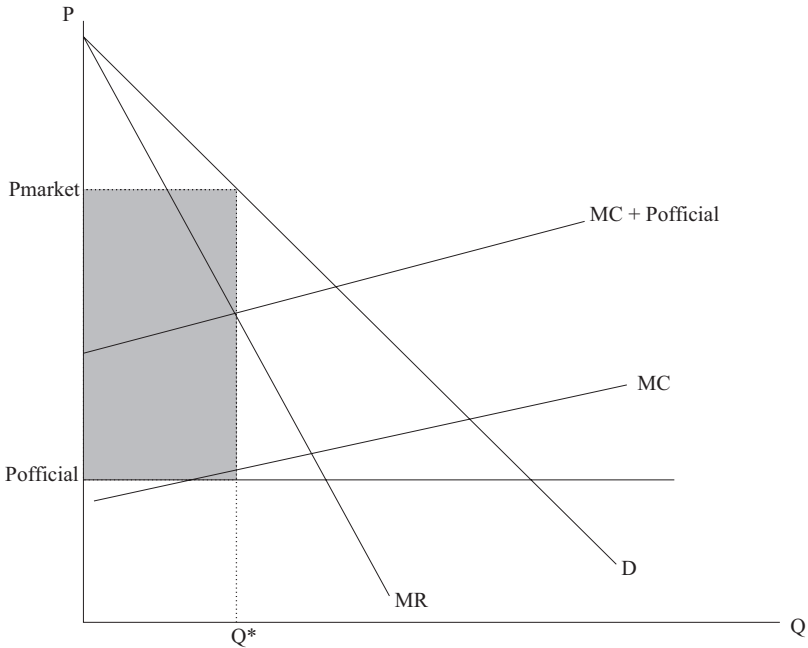


FIGURE II
Underpricing and Bribes

The equilibrium spot price and quantity are P^{market} and Q^* . The official price paid by the oil traders is P^{official} . The shaded area is the total amount of illegal rents. These rents are presumably split between Iraq and the traders.

We therefore assume that Iraq can choose the official price, but that there is uncertainty over the future market price and that Iraq is punished if the UN is reasonably certain that the official price is deliberately set below market prices. Specifically, we assume that the expected punishment is an increasing convex function of the gap between the market price and the official price divided by the standard deviation of the market price:

$$(1) \quad f\left(\frac{P^{\text{market}} - P^{\text{official}}}{SD}\right).$$

We scale the price gap by the standard deviation of the market price to capture the idea that it was more difficult for the UN to determine whether the proposed official price was abnormally low in periods of high oil market volatility. In other words, the ex post realization of a positive price gap could reflect deliberate under-

pricing, or could simply reflect ex ante uncertainty over future prices. (More volatility implies more ex ante uncertainty.) With this assumption, the "optimal" official price chosen by Iraq is determined by $Q^* = f'(P^{\text{market}} - P^{\text{official}}/SD)$. Intuitively, Iraq chooses the official price such that the marginal gain from lowering the official price is equal to the marginal increase in the expected punishment from lowering the official price.

There are two straightforward implications of this model. A first implication is that increased volatility lowers the marginal probability of detection for any given price gap. Therefore, Iraq will choose to lower the official price in response to an increase in price volatility. Intuitively, higher price volatility implies greater uncertainty over whether a given price gap reflects cheating or whether it simply reflects the higher volatility in the market price. Therefore, if our story of underpricing and bribes is correct, we expect that Iraq will take advantage of the greater uncertainty due to higher volatility to lower the official price of its oil. We will provide empirical evidence in support of this hypothesis in subsection V.A.

A second implication has to do with buyer heterogeneity. While the cost of detection was probably low for many obscure individual traders, it was high for large reputable multinationals. One would expect the first group to be more likely to pay bribes than the second. Under this scenario, some of the rent was appropriated by reputable buyers in the form of low oil prices. Obviously, Iraq would prefer to deal only with individual traders willing to pay bribes. However, it is possible that, in some periods, the supply of Iraqi oil exceeded the demand from nonreputable buyers or that Iraq had to deal at least in part with reputable oil companies to minimize international scrutiny. We will examine this hypothesis in subsection V.B.

III. EMPIRICAL EVIDENCE ON UNDERPRICING

In this section we document the extent by which Iraqi oil was underpriced when compared with the market price of several comparison crude oils. First, we compare the official selling price of Iraqi crude oils with the price of its closest substitutes: Arabian Light, Arabian Medium, and Urals. Second, we compare the official selling price of Iraqi crude oils with the spot market price of Iraqi crude oil itself. The data used are described in the Appendix.

III.A. Close Competitors

Iraq produces two types of crude oil, known as Basrah Light and Kirkuk. The closest substitute of Basrah Light is Arabian Light, a crude oil produced in Saudi Arabia. In turn, the Russian Urals crude oil is widely regarded as the closest substitute for Kirkuk. Because these oils are physically so similar, their prices are almost perfectly correlated. In the years before the program, the correlation between the price of Arabian Light and Basrah Light is 0.993, while the correlation between the price of Kirkuk and Urals is 0.985.⁴

Another way to assess the comparability of Arabian Light and Urals to the two Iraqi oils is to plot the price gap between the comparison oils and the Iraqi official prices. The top panel in Figure I plots the difference between the price of Arabian Light and the official selling price of Basrah. The bottom panel shows the difference between the market price of Urals and the official selling price of Kirkuk. The first vertical line marks the beginning of trading under the Oil for Food Program (tenth week of 1997) and the horizontal line is the average difference for the years before the Oil for Food Program. As can be seen, the price gap between Basrah and Arabian Light averages zero in the years prior to the Oil for Food program. The price gap between Kirkuk and Urals in the 1980s is also zero. Together with the almost perfect correlation in the price level, the absence of a price gap provides some reassurance that Arabian Light and Urals are reasonably close substitutes for the Iraqi oils.

After the Oil for Food Program is introduced, this difference is clearly larger for Basrah: the gap averages \$2 in 1997 and 1998, and climbs above \$5 between 2000 and 2001. Notably, after retroactive pricing is introduced in September 2001, the difference drops quickly to its historical level (the second vertical line in Figure I indicates the beginning of retroactive pricing). There

4. The market price of Arabian Light and official prices of the two Iraqi oils are quoted as a fob price, while the market price of Urals is typically quoted as a delivered price in the port of Augusta (Italy). To compare the market price of Urals with the official selling price of Kirkuk, we convert the delivered price of Urals in Italy to a fob price by subtracting oil tanker rates from the Black Sea to Italy. The average transportation cost between 1989 and 2002 is \$.87. During the program years, transportation costs fluctuate between .58 in 1999 and 1.22 in 2000. Transportation costs are slightly higher during the Oil for Food program years. A regression of transportation cost on a dummy for the Oil for Food years yields .10 (.04). Transportation costs are also positively correlated with the price of oil.

is also some evidence that Kirkuk was underpriced from 2000–2001, albeit by not as much as Basrah.

We can quantify the average price gap by estimating the following regression:

$$(2) \quad \Delta P_t = \alpha + \beta \text{Program}_t + \epsilon_t,$$

where t indexes week, ΔP_t is the difference between the market price of the comparison crude and the official selling price of the relevant Iraqi crude, and Program_t is an indicator variable for periods of the Oil for Food Program in which sales of Iraqi oil took place (March 1997–December 2002). The excluded category, captured by the intercept, is the period before the Oil for Food Program (in most cases it includes years 1980–1996). The coefficient β measures the average *change* in the price *difference* during the Program relative to the historical baseline before the program. Because the figures indicate that the price gaps vary significantly between 1997 and 2002, we also estimate a model that allows the gap to differ in different periods:

$$(3) \quad \Delta P_t = \alpha + \beta_1 \text{Program}_{1t} + \beta_2 \text{Program}_{2t} + \beta_3 \text{Program}_{3t} + \epsilon_t,$$

where Program_{jt} is an indicator variable for the j th subperiod of the Oil for Food Program (for example, $j = 1997\text{--}1999, 2000\text{--}2001, 2002$).

One concern is that the error term ϵ_t might be autocorrelated. The residual is the *difference* between shocks to Iraqi oil and shocks to its competitors. Serial correlation might occur if the shocks to the demand and supply of Iraqi oil relative to its competitors are serially correlated. To account for possible serial correlation, we estimate models where the error structure is assumed to be autocorrelated up to five lags and heteroskedastic [Newey and West 1987].⁵

Table I quantifies the visual impression one gets from Figure I by presenting estimates of β [from (2)]. The first column presents the estimates of the change in the gap between Arabian Light and Basrah. The estimate in the top panel indicates that

5. We have also experimented with models where the error structure is assumed to be autocorrelated up to ten lags, and found that results do not change significantly. Because of the gaps in our series, we have also experimented with the Baltagi and Wu [1999] estimator for cases where the observations that are unequally spaced over time and the residual follows an AR(1) process. Results are similar. We have also experimented with models where we cluster the standard errors at the trimester-year level. This is not ideal, because it assumes no correlation between shocks in different trimesters.

TABLE I
DIFFERENCE BETWEEN THE MARKET PRICE OF CLOSE COMPETITORS AND THE
OFFICIAL SELLING PRICE OF IRAQI OILS

	Arabian Light– Basrah (1)	Arabian Medium– Basrah (2)	Urals– Kirkuk (3)
Model 1: Overall estimates			
Difference for 1997–2002 – difference for 1980–1995	2.44** (0.25)	1.24** (0.34)	0.69** (0.18)
Model 2: Estimates by period			
Difference for 1997–1999 – difference for 1980–1995	2.07** (0.18)	0.84** (0.24)	0.53** (0.15)
Difference for 2000–2001 – difference for 1980–1995	3.91** (0.41)	2.67** (0.61)	1.07** (0.40)
Difference for 2002 – difference for 1980–1995	0.68 (0.36)	–0.40 (0.49)	0.45 (0.29)
Total revenue from underpricing (billion)	4.12	2.28	0.76

Standard errors are in parentheses. The dependent variable is the difference between the price of the closest competitor and the official selling price of the relevant Iraqi crude oil, in dollars per barrel. The entry in row 1 is the coefficient on a dummy equal 1 for observations during the Oil for Food Program (between the tenth week of 1997 and the last week of 2002). The excluded category includes observations before the Oil for Food Program (between 1980 and 1996). The level of observation is a week. Entries in rows 2–4 are from one regression, and are coefficients on dummies equal 1 for observations during the stated period. The error structure is assumed to be autocorrelated up to five lags and heteroskedastic. The entry in the bottom row is obtained by multiplying the amount of the relevant Iraqi oil sold in each week between 1997 and 2002 by the difference between the price difference observed in each week between 1997 and 2002 and the average price difference for the period 1980–1995. (See (4).) Sample size in columns (1)–(3) is, respectively, 785, 667, and 690. **Significant at the 5 percent level.

the underpricing averaged \$2.44 a barrel during the Oil for Food Program years (relative to the years before the program).⁶ Because the price gaps appear to vary over time, in the bottom panel we show separate estimates for three subperiods (3). These estimates indicate that the price gap was significantly different from zero in the 1997–1999 period, reaches a peak in the 2000–2001 period, and declines in 2002.

An alternative substitute for Basrah Light is the Saudi crude oil Arabian Medium. While the Arabian Light has historically been the closest substitute of Basrah light, there is some evidence that the quality of Basrah Light became much closer to that of

6. Specifically, this is the gap from the tenth week of 1997 through the last week of 2002 relative to the fifteen years prior to the program.

Arabian Medium by the end of the program.⁷ For this reason, in column (2) we present a comparison based on the price of Arabian Medium. We find that the estimated underpricing averaged \$1.24 a barrel during the Oil for Food Program years (relative to the years before the program). This estimate is smaller than the comparison based on Arabian Light (column (1)), which suggests that part of the increase in the price gap between Arabian Light and the official selling price of Basrah Light after 1997 may be due to the deteriorating quality of the Iraqi oil. Finally, the estimate based on the comparison between Urals and Kirkuk (column (3)) suggests an average underpricing of only \$0.69 a barrel.⁸

On the basis of the estimates of these price gaps, we can calculate the rents created by the underpricing of Iraqi oil. Specifically, we estimate these rents as the product of quantity sold in each week and the difference between that week's price gap and the baseline price gap:

$$(4) \quad I = \sum_t [(\Delta P_t - \hat{\alpha}) \times Q_t],$$

where ΔP_t is the actual price gap in a given week, $\hat{\alpha}$ is the estimate of the average price gap in the years before the Program (the intercept in (2)), and Q_t is the quantity of the relevant Iraqi oil sold that week. This estimate indicates that the underpricing of Iraqi oil created \$2.28 billion to \$4.12 billion in rents for buyers of Basrah Light (depending on whether Arabian Light or Medium is used a comparison) and \$0.76 billion in rents for buyers of Kirkuk. These rents were presumably split between the Iraqi regime and oil buyers.⁹

III.B. Market Price of Iraqi Oil

We now turn to our second measure of the market value of Iraqi oil: the spot market price of Iraqi oil itself. That is, we compare the official selling price of Iraqi oil with the spot market

7. Authors' personal communication with Mr. Maurice Lorenz (American oil overseer for the Oil for Food program from 1997 to 1998).

8. The estimated AR(1) parameters in columns (1) and (3) are 0.88 and 0.62.

9. For completeness, we have also investigated the gap between the official price of the Iraqi oils and the market price of West-Texas Intermediate (WTI) and Brent. WTI and Brent are two crude oils widely used as benchmarks for the world price of oil. The patterns of the price gap are shown are generally similar to the patterns in Figure I (see Hsieh and Moretti [2005]). The patterns of the price gap are generally similar to the patterns in Figure I.

price of the same oil. Oil traders that received Iraqi oil contracts sold these contracts in spot markets in Europe and the United States. The price paid to oil traders in these spot market transactions presumably includes the cost of the bribe paid to the Iraqi regime. After an adjustment for transportation costs, the gap between the spot market price and the official selling price should be the most direct measure of underpricing. However, the spot market data of Iraqi oil are of lower quality than the data on the spot market prices of the other four comparison crude oils.¹⁰

Figure III shows the price difference between the market price of Iraqi crude and its own official selling price. The top panel is for Basrah. As can be seen, the overall picture is qualitatively similar to Figure I. The figure suggests that underpricing was somewhat limited from 1997 to 1999, large from 2000 to 2001, and declined after September 2001. The bottom panel is for Kirkuk. The graph is noisier than the previous figures, and the evidence for underpricing for Kirkuk appears therefore weaker than that for Basrah. We do not know exactly why the documented price gap appears to be larger for Basrah than for Kirkuk.¹¹

Table II quantifies the magnitude of the price gap. Looking at the overall estimates, the point estimates are 1.64 and 0.29 for Basrah and Kirkuk, respectively. The estimate for Basrah is remarkably close to the estimate based on the comparison between Arabian Medium and Basrah (column (2) in Table I), while the estimate for Kirkuk is lower than the previous estimate.¹²

Estimates by period show a pattern similar to the one shown in Table I, with the peak of the underpricing reached in the

10. In particular, we have data on the spot market price of Kirkuk and Basrah from three independent sources (ICISLOR, Platts, and Petroleum Argus; see the Appendix). But no data source covers the entire period. To construct a series for the spot market price of the two Iraqi crude oils over the entire time period, we simply take the average of the price gaps from the data sources available in each week.

11. We offer a potential explanation in Section IV.

12. The estimated AR(1) parameters are .78 and .54. Clustered standard errors are, respectively, .28 and .15. We have also reestimated all the models assuming that the error structure is autocorrelated up to ten lags. The standard errors increased slightly, without changing the qualitative conclusions on significance of the coefficients. For example, the standard errors in column (1)–(3) in Table I become .32, .32, and .16. Standard errors in column (1) and (2) of Table II become .25 and .13. We have reestimated the models in Tables I and II limiting the preprogram samples to be the same in both tables. The estimates do not change significantly. To check robustness, we reestimated all the models controlling for the price of Brent. Results do not change very much. For example, estimates in column (1)–(3) in Table I become 2.47 (.24), 1.19 (.22), and .65 (.12). Estimates in column (1) and (2) of Table II became 1.25 (.16) and .27 (.13).

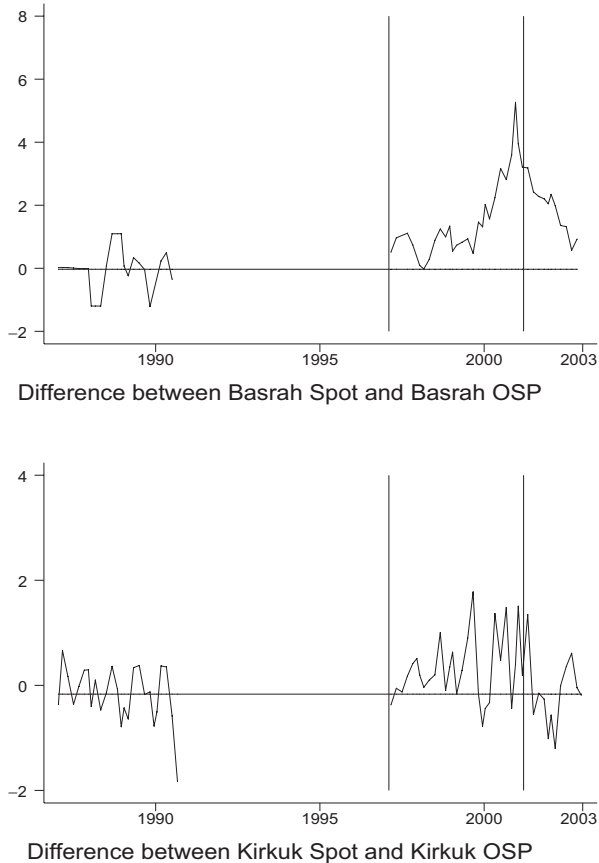


FIGURE III
Difference between the Market Price of Iraqi Oils and the
Official Selling Price of Iraqi Oils

The top panel shows the difference between the market price of Basrah and the official selling price of Basrah. The bottom panel shows the difference between the market price of Kirkuk and the official selling price of Kirkuk. The first vertical line marks the beginning of the Oil for Food Program. The second vertical line indicates the beginning of retroactive pricing. The horizontal line is the average difference for the years before the Oil for Food Program. Iraqi oil was not traded in 1991–1993 due to the first Gulf War, and in 1995–1997 due to delays in the Oil for Food Program.

2000–2001 period. Although the patterns over time are generally similar, it appears that the findings in column (1) and (3) of Table I may be an overestimate of the true magnitude of the price gap.

TABLE II
DIFFERENCE BETWEEN THE MARKET PRICE OF IRAQI OILS AND THE OFFICIAL
SELLING PRICE OF IRAQI OILS

	Basrah market price – Basrah official selling price (1)	Kirkuk market price – Kirkuk official selling price (2)
Model 1: Overall estimates		
Difference for 1997–2002 – difference for 1980–1995	1.64** (0.20)	0.29** (0.12)
Model 2: Estimates by period		
Difference for 1997–1999 – difference for 1980–1995	0.85** (0.15)	0.33** (0.13)
Difference for 2000–2001 – difference for 1980–1995	2.82** (0.24)	0.40* (0.21)
Difference for 2002 – difference for 1980–1995	1.46** (0.22)	–0.03 (0.21)
Total revenue from underpricing (billion)	3.03	0.48

Standard errors are in parentheses. The dependent variable is the difference between the market price of an Iraqi crude oil and the official selling price of the same Iraqi crude, in dollars per barrel. The entry in row 1 is the coefficient on a dummy equal to 1 for observations during the Oil for Food Program (between the tenth week of 1997 and the last week of 2002). The excluded category includes observations before the Oil for Food Program (between 1980 and 1996). The level of observation is a week. Entries in rows 2–4 are from one regression, and are coefficients on dummies equal to 1 for observations during the stated period. The error structure is assumed to be autocorrelated up to five lags and heteroskedastic. The entry in the bottom row is obtained by multiplying the amount of the relevant Iraqi oil sold in each week between 1997 and 2002 by the difference between the price difference observed in each week between 1997 and 2002 and the average price difference for the period 1980–1995. (See (4).) Sample size in columns (1) and (2) is, respectively, 458 and 487. **Significant at the 5 percent level; * significant at the 10 percent level.

However, the estimate of underpricing based on the comparison of Basrah Light with Arabian Medium is remarkably similar to the estimate obtained from the spot market price of Basrah Light. Estimates of the illegal revenues are \$3.03 billion and \$0.48 billion, respectively.

Our estimates of underpricing match fairly well the only existing piece of documentary evidence of underpricing. This documentary evidence is from a report by the commission headed by Paul Volcker, and refers to five purchases of Kirkuk crude oil (for a total of 7.3 million barrels) by an oil trading company called the Africa Middle East Petroleum Company (AMEP). Based on written documentation on the price paid by AMEP to Iraq and on the price received by AMEP on the spot market by refiners (mostly by Shell), the Volcker commission calculates that Kirkuk was un-

derpriced by an average of \$.25 in these five cases [IIC 2005a]. Our estimates based on the comparison between the OSP price for Kirkuk and the market price for Kirkuk in the five three-week windows centered around the date of the AMEP sales is \$.27, which is remarkably consistent with the evidence provided by the Volcker commission.

In an ideal world we would want to have the kind of documentary evidence compiled by the Volcker commission for the AMEP oil sales for all 1300 oil contracts during the Oil for Food Program, but this would be prohibitively costly. As an alternative, our estimates of underpricing have the advantage of being easy to compute and for relying only on publicly available market data. In addition, it is worth noting that the five transactions examined in detail by the Volcker commission do not appear to be representative of the underpricing during the Program. First, the AMEP contracts were for Kirkuk, and we have shown that underpricing was much less prevalent with Kirkuk than with Basrah. Second, three of the sales are from the 1998 to 1999 period, where underpricing was generally low.

IV. ALTERNATIVE EXPLANATIONS

In the previous section we have presented evidence that there seems to have been significant underpricing of Iraqi oil (especially of Basrah). We interpret the underpricing as a deliberate attempt by Iraq to extract bribes from the oil buyers. Of course, this is not the only possible interpretation. In this section we analyze whether alternative interpretations are consistent with the evidence and ultimately conclude that they are not.

(1) *Stigma Associated with Buying Iraqi Oil.* It is possible that in the years after the first Gulf War, dealing with the Iraqi regime carried a stigma in the eyes of the world public opinion. Large multinational oil companies interested in preserving their reputation might have been particularly sensitive to this, and thus chose not to purchase Iraqi oil. If this is true, the demand for Iraqi oil might have declined relative to the demand of its close competitors. This decline in relative demand—not strategic underpricing aimed at obtaining bribes—could in theory explain the lower prices of Iraqi oil relative to its competitors.

However, this hypothesis is not consistent with three pieces of evidence. First, it is difficult to reconcile with the existence of a gap between the market price of Iraqi oil and its *own* official selling price. Any stigma associated with purchasing Iraqi oil should affect both the official selling price and the market price of Iraqi oil. If the stigma effect was the only explanation of underpricing, we should not see underpricing when we compare the market price of Iraqi oil with its own official selling price. Second, to explain the decline in underpricing after the introduction of retroactive pricing, the stigma hypothesis would require a significant improvement in the perception of the Iraqi regime in the eyes of the world public opinion after September 2001.

Third, underpricing does not appear to be more prevalent during periods when world public attention was focused on Iraq.¹³ One way to directly quantify such attention is to measure the coverage of Iraq in the media. Figure IV plots the total number of articles that appeared in the *New York Times* and the *Wall Street Journal* between 1996 and 2002 that contain the word “Iraq” (top panel), “Iraq” and “Weapons” (middle panel), and “Iraq” and “Human Rights” (bottom panel). As can be seen, there was a spike in news coverage about Iraq in 1998 and 1999 (largely on weapon inspections), but this was a period of relatively low underpricing. There also does not appear to be any particular decline in the number of articles after September 2001, nor does it show any particular increase in the peak underpricing period in 2000–2001.¹⁴

(2) *Decline in Iraq Shipping Facilities.* An alternative explanation for the lower prices commanded by Iraqi oil is the decline

13. In particular, the UN was focused on Iraq in 1998 and 1999 due to the struggle over the weapons inspections regime. Yet, this was a period when underpricing was relatively low. In contrast, the peak underpricing period apparently began in the May 2000 when the UN did not have its eyes on Iraq, but rather was focused on two crises in Africa: the kidnapping of UN peacekeepers in Sierra Leone and an impending war between Ethiopia and Eritrea. In May 2000 the Security Council was so preoccupied with events elsewhere in the world that it did not have time to authorize a routine renewal of the Oil for Food program, and it was precisely in May 2000 that the peak underpricing period began. Platt's Oilgram News (6/1/2000) reports “With the UN Security Council overloaded with work on the world's crisis spots this week, members have decided to delay a vote on the rollover of the Iraq oil-for-food program until next week.”

14. In the top panel in Figure IV, a regression of underpricing on number of articles yields $-.001$ (.001). In the middle panel this regression yields $-.030$ (.023). In the bottom panel it yields $-.005$ (.003). Results based on the number of first-page articles or the total number of words in articles about Iraq are similar.

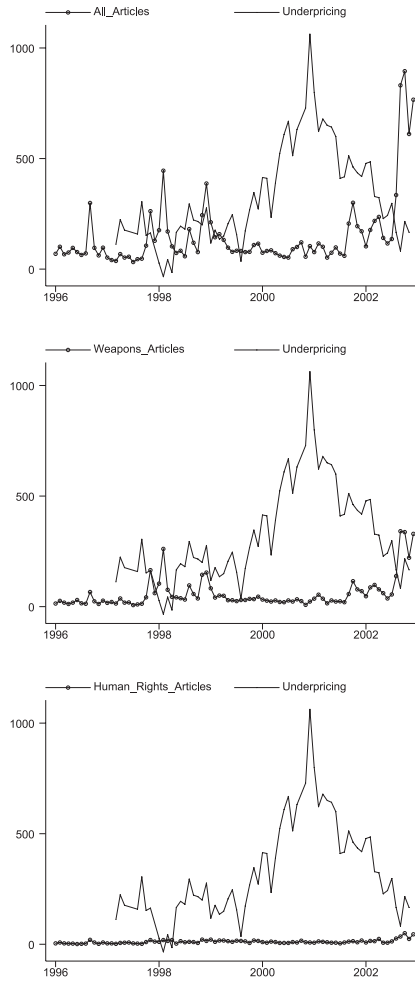


FIGURE IV

The Relationship between Newspaper Coverage of Iraq and Underpricing

In the top panel, the line with circles shows the number of articles that contain the word “Iraq” published in the *New York Times* and the *Wall Street Journal*, by month. The solid line shows the estimated underpricing, by month (based on the difference between market price and OSP of Basrah). The scale on the y-axis refers to the number of articles. A regression of underpricing on number of articles yields $-.001$ (.001). In the middle panel the line with circles shows the number of articles that contain the words “Iraq” and “Weapons” published in the *New York Times* and the *Wall Street Journal*, by month. A regression of underpricing on number of articles yields $-.030$ (.023). In the bottom panel, the line with circles shows the number of articles that contain the words “Iraq” and “Human Rights” published in the *New York Times* and the *Wall Street Journal*, by month. A regression of underpricing on number of articles yields $-.005$ (.003).

in Iraq shipping facilities. Years of embargo are likely to have curtailed Iraqi access to new technologies and to the spare parts necessary to keep its oil terminals up to date. As a consequence, it is possible that the decline in the price of Iraqi oil simply reflects longer delays experienced by oil tankers loading oil at Iraqi port, or higher variance of these delays. Both an increase in mean or an increase in variance in loading time would make Iraqi oil relatively less attractive and could therefore cause a decline in its relative price.

Can a decline in shipping facilities explain the underpricing? It seems unlikely. First, there is no particular reason for why the ports should have improved substantially after September 2001. More importantly, to measure delays experienced by oil tankers, we collected data on the arrival and departure date of every ship that loaded oil at the Basrah oil terminal, which is the main port facility in Iraq.¹⁵ If shipping facilities deteriorated, we should see an upward trend in the number of days that it takes to load an oil tanker, holding constant the number of ships. If port deterioration can explain underpricing, we should also see that periods when the delays are longest or when the variance in delay is highest are also periods when underpricing is largest.

Figure V shows that there does not seem to be any particular trend in the loading time in the five years under consideration. It also indicates that variation in underpricing is not positively correlated with changes in the length of time it takes to load tankers, or change in its variance. A regression of underpricing on average waiting time yields a negative but statistically insignificant coefficient (coeff: $-.40$; std error: $.26$). A regression of underpricing on the standard deviation of waiting time calculated across all shipments in the relevant month yields an insignificant coefficient (coeff: $-.05$; std error: $.26$). These results do not change if we control for number of ships to account for possible crowding at the port or if we use as dependent variable delays divided by tonnage to account for heterogeneity in the sizes of the tankers.

(3) *Decline in the Quality of Iraqi Oil.* A decline in the price of Iraqi oil could in theory be explained by a decline in the relative quality of Iraqi oil. For example, we have suggested

15. We purchased these data from Lloyds of London.

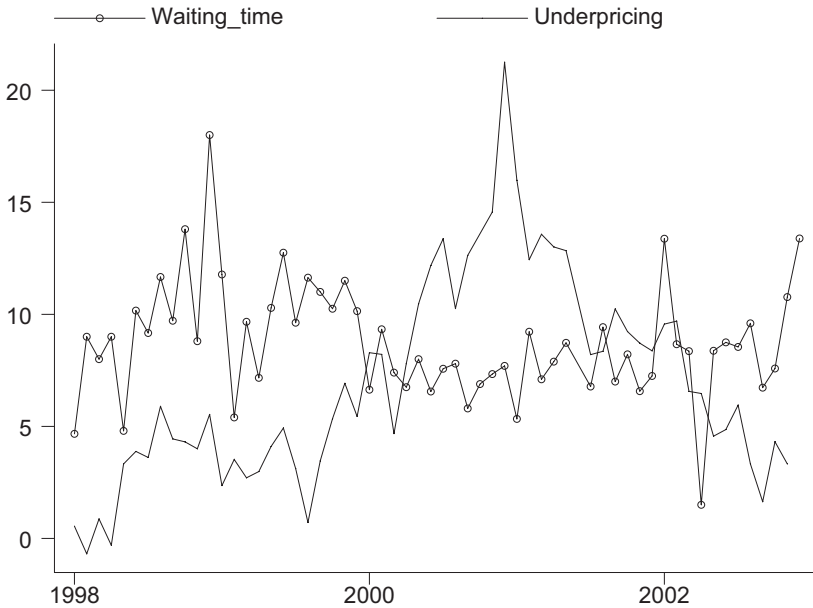


FIGURE V

The Relationship between Delays at Iraqi Ports and Underpricing

The line with circles shows the average waiting time (in days) experienced by tankers loading oil at the Basrah oil terminal, by month. The solid line shows the estimated underpricing, by month (based on the difference between market price and OSP of Basrah). The scale on the y-axis refer to days of waiting time. A regression of underpricing on waiting time yields $-.40 (.23)$. A graph based on the ratio of waiting time divided by the number of ships loading oil in the relevant month generates very similar picture.

that although Basrah Light has historically been virtually identical to Arab Light, there is some evidence that the quality of Basrah Light fell relative to Arab Light over the course of the Program. Therefore, the price gap between Arab Light and the official selling price of Basrah is likely to overstate the true degree of underpricing.

However, our best estimate of underpricing is based on the comparison of the official price of Iraqi oil and the market price of the same oil. This is our best estimate of underpricing because the two prices are for *exactly the same oil*, albeit at different stages of the distribution process. Therefore, the gap between Iraq's official selling price and the market price of Iraq oil are clearly are not contaminated by possible changes in oil quality.

Finally, we note again that the price gap between Iraqi oil and other oils drops to zero after the introduction of retroactive pricing, and there is no reason to believe that quality differences abruptly ended in September 2001.

(4) *Increased Supply*. It is in theory possible that the fall in the official price of Iraqi oil was a consequence of a deliberate effort by the Iraqi regime to sell more oil. Specifically, if Iraq valued the humanitarian supplies purchased by the Oil for Food Program, it might have chosen to sell more oil to increase the resources flowing into the program at a time when the marginal utility of consumption was particularly high. If the demand curve for Iraqi oil is not perfectly elastic, the increased supply of oil will result in a decline in the price of Iraqi oil.

Two pieces of evidence argue against this hypothesis. First, although it is probably less than infinite, the demand for Iraqi oil is likely to be very elastic.¹⁶ If the demand curve for Iraqi oil is very elastic, the large price differences documented in the previous section (more than \$1 on average, with peaks of up to \$4) would imply enormous increases in quantity sold. However, we simply do not observe such large variations in quantity sold. Figure VI presents a scatter-plot of the price gap and the quantity sold of Basrah (top panel) and Kirkuk (bottom panel). If the episodes of underpricing were triggered by an outward shift in the relative supply of Iraqi oil, we should see a *positive* relationship between the price gap and the quantity. However, the relationship between the price difference and quantity is either not statistically different from zero or negative.¹⁷ Second, the supply hypothesis cannot explain the evidence of underpricing even when we measure the gap between the market price of Iraqi oil and its official selling price.

(5) *Market Volatility*. Next, we address the impact of oil market volatility on our measures of underpricing. Until September 2001, the price of Iraqi oil was set a month in advance, and this price was stipulated as a fixed discount to the *future* price of

16. How elastic is the price elasticity of demand for Iraqi oil? In Hsieh and Moretti [2005], we use an exogenous change in the supply of Iraqi oil to show that the elasticity of demand of Iraqi oil is quite high.

17. A regression of quantity sold on the price gap yields a coefficient (std. error) of -2.17 (4.00) for Basrah and -8.61 (4.34) for Kirkuk. We note, however, that it is of course possible that these coefficients are biased, if unobserved shocks to the demand of Iraqi oil are systematically correlated with changes in its supply.

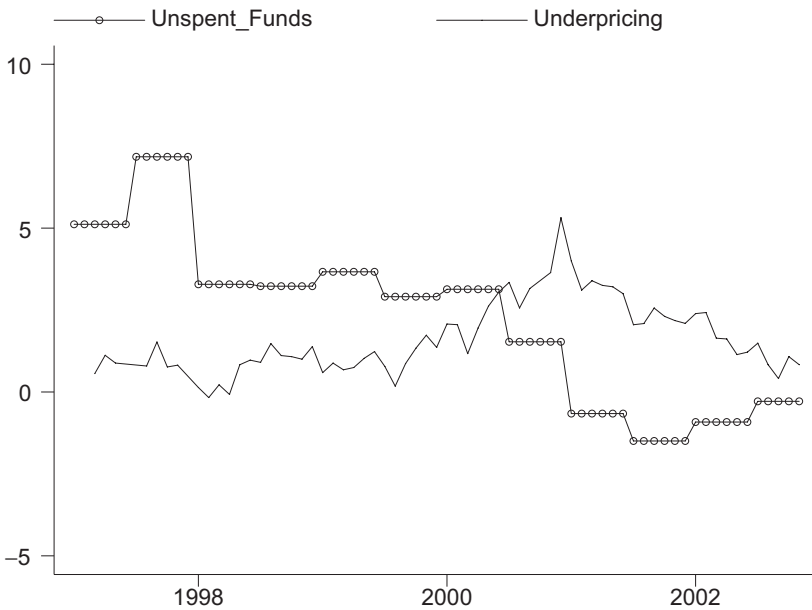


FIGURE VI

Quantity of Iraqi Oil Sold and Underpricing of Iraqi Oil

The top figure shows the quantity of Basrah oil sold and the difference between the market price of its closest competitor, Arabian Light, and the official selling price of Basrah. A larger price difference means a larger underpricing of Basrah. A regression of quantity sold on the price gap yields a coefficient (std error) of -2.17 (4.00). The bottom figure shows the quantity of Kirkuk oil sold and the difference between the market price of its closest competitor, Urals, and the official selling price of Kirkuk. A larger price difference means a larger underpricing of Kirkuk. A regression of quantity sold on the price gap yields a coefficient (std error) of -8.61 (4.34).

benchmark oils. Since the future price of the benchmark oils is not known at the time the buyer enters into a contract to purchase Iraqi oil, the buyer of Iraqi oil bears the risk of any price volatility. This volatility presumably lowers the value of an Iraqi oil contract.

Could this risk explain the lower price of Iraqi oil relative to its competitors that we have documented in the years before September 2001? It seems unlikely. The reason is that the competitors' oils are priced in a way similar to Iraqi oil. Specifically, Arab Light, Arab Medium, and Urals are priced at a discount to the future price of world benchmark oils. The price of Iraqi oil sold by oil traders is also quoted at a discount to the

relevant benchmark oils.¹⁸ Therefore, prior to September 2001, volatility in oil markets should have exactly the same effect on the comparison oils as on the official price of Iraqi oil.

It is the introduction of retroactive pricing in September 2001 that should have lowered the value of an official Iraqi oil contract relative to its competitors. The reason is that after the introduction of retroactive pricing, buyers of Iraqi oil learn the pricing formula only after the oil had been loaded onto the ships. Therefore, the observed price gap after September 2001 is likely to *overstate* the “risk-adjusted” price gap between Iraqi oil and the comparison oils. This could in part explain why the observed decline in underpricing is not sharper after September 2001. To get a sense of how much this may matter, we use options to value this risk. We purchased daily options data from NYMEX for the period 1998–2002 and run a regression of the call options price on the strike price and the strike date to estimate the price of a one-month call option for a strike price equal to the actual realized value of the oil price one month in the future.¹⁹ Our estimate suggests that the price gap after September 2001 is likely to be overstated by 30 cents. In other words, in the risk-adjusted version of Figure I, the underpricing line should be 30 cents lower after September 2001.

Finally, we note that our underpricing estimates are largely unchanged when we introduce controls for oil market volatility.²⁰ In the next section we show that volatility does affect underpricing, but only during the Oil for Food period.

(6) *Reduced Effort*. The official revenues of the Oil for Food Program were largely spent on humanitarian supplies. However, there were periods in which these funds were not spent due to administrative delays in the UN. Therefore, if Iraq did value these humanitarian supplies, but was frustrated by the fact that

18. The pricing of Saudi oils is identical to that of Iraqi oils. That is, the prices are set at a discount to world benchmark oils, where the relevant discount and benchmark used varied depending on the destination of the oil. The discounts obviously differed from that used by Iraqi oils, but the benchmarks oils used are exactly the same. In turn, since Urals is sold on a delivered basis in Italy, it is priced at a fixed discount to the benchmark oil on the day the oil is delivered in Italy. However, this price is obviously not known when the buyer of Urals enters into a contract to purchase the Russian oil.

19. Specifically, we ran a regression of the log of the option price on a quadratic in the log of the strike price and in the length of the strike date.

20. For example, estimates in column (1)–(3) in Table I became 2.45 (.25), 1.19 (.25), and .61 (.12). Estimates in column (1) and (2) of Table II became 1.49 (.21) and .32 (.13).

some of these funds were not spent, Iraq may have chosen to reduce its effort to obtain the best price for its oil. According to this “reduced effort” interpretation, Iraq did not underprice in an effort to extract bribes, but because the UN was not spending the official revenues in a timely manner. To investigate this possibility, we collected data on the amount of undisbursed funds in the UN escrow account. We focus on the fraction of revenues that is not spent for humanitarian aid in Iraq in each six-month period over the total oil revenues in that period as a measure of the implicit tax-rate faced by Iraq.²¹ If this argument is correct, we should observe that higher tax rate are associated with lower effort and therefore higher underpricing. Figure VII shows that this does not appear to be the case. Overall, the fraction of undisbursed funds is negatively, not positively, correlated with underpricing. The correlation is $-.46$.

V. COMPARATIVE STATICS: MARKET VOLATILITY AND TYPE OF BUYERS

In this section we investigate two implications of our hypothesis. First, we test whether underpricing was higher in periods of high price volatility, and lower in periods of price stability. Second, we test whether episodes of underpricing are associated with a decline in the share of major oil multinationals among the oil buyers.

V.A. *The Relationship between Underpricing and Price Volatility*

In this subsection we propose a specification check based on the notion that monitoring is easier when volatility is low. Until September 2001, Iraq could set the price of its oil, but the price had to be approved by the UN. In particular, each month Iraq set the price of its oil for transactions that would take place the *following* month. This price had to be approved by the UN, who had to ascertain that the price set was as close as possible to the *expected* market price for the following month. Since there are no future prices for Iraq oil, this was not an easy task for the UN overseers, especially in periods when world oil markets were volatile. When the price of oil on the world markets was stable, predicting the future price of oil was relatively easy. However, in periods when the price of oil was very volatile, predicting the

21. These data are available only every six months.

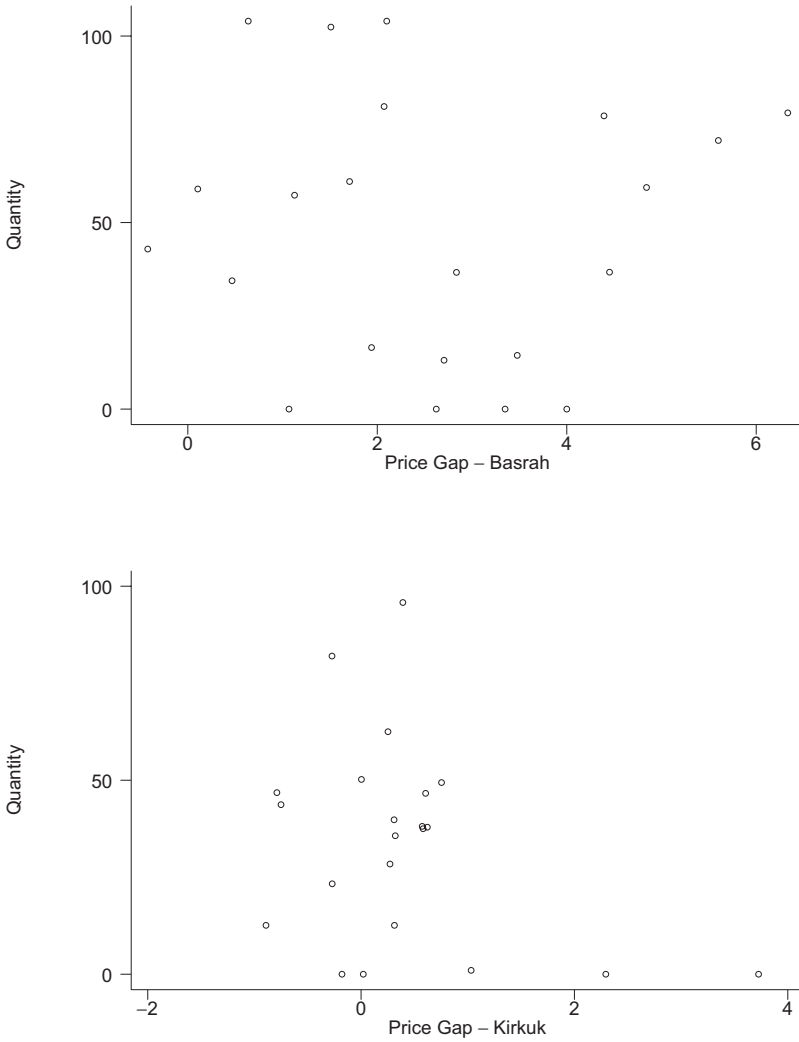


FIGURE VII

The Relationship between Unspent Funds and Underpricing

The line for "Unspent Funds" represents the fraction of revenues that is not spent for humanitarian aid in Iraq in each six-month period over the total oil revenues in that period.

future price of oil became much harder. This suggests that it should have been easier for Iraq to get away with a low price in periods characterized by high volatility in world oil prices. It was

TABLE III
THE RELATIONSHIP BETWEEN UNDERPRICING AND THE VOLATILITY
OF WORLD OIL PRICES

	Basrah (1)	Kirkuk (2)
Before retroactive pricing		
Coeff. on std. dev. of price of Brent	0.916** (0.094)	0.306** (0.084)
After retroactive pricing		
Coeff. on std. dev. of price of Brent	0.190 (0.176)	-0.067 (0.191)

Standard errors are in parentheses. Each entry is from a separate regression. The dependent variable is the estimated underpricing by week. The independent variable is the standard deviation of the price of Brent in the ten weeks preceding the relevant week. Estimates in row 1 are based on a sample that includes weeks before the introduction of retroactive pricing (September 2001). Estimates in row 2 are based on a sample that includes weeks after the introduction of retroactive pricing. The error structure is assumed to be autocorrelated up to five lags and heteroskedastic. The unit of measurement for both the dependent and independent variable is dollar per barrel. **Significant at the 5 percent level.

harder for the UN overseers to object to a low price in periods where the future price of oil was very uncertain.

After September 2001, the UN changed the way Iraqi oil was priced. Specifically, after this date, the price of Iraqi oil was not known until *after* the oil had already been loaded on the oil tankers. That is, the price of Iraqi oil was set by the UN only after the UN could observe the actual price in the market. Therefore, the relationship between underpricing and volatility presumably changed after the introduction of retroactive pricing, since retroactive pricing effectively forced Iraq to set prices equal to the actual market price, irrespective of volatility.

If this is true, we should observe two features in the data. First, in the period before the introduction of retroactive pricing, we should see that our estimates of underpricing are higher in weeks of high oil price volatility. Second, in the period after the introduction of retroactive pricing, we should see that our estimates of underpricing are not correlated with volatility. Finding a correlation between underpricing and volatility after the adoption of retroactive pricing would cast doubt on the validity of the interpretation of this specification test.

Table III presents estimates from four regressions of underpricing on market volatility. Because the UN overseers used the price of Brent (a widely used international benchmark price) as a reference price for shipments going to Europe, we measure volatility on the world oil market as the standard deviation of the price of Brent in the ten weeks preceding the observed transac-

tion.²² (Note that the dependent variable is of course a price difference, while the dependent variable is volatility in the price level.)

The table shows a pattern that matches remarkably well our predictions. In the period before retroactive pricing (September 2001), there is a positive association between our estimates of underpricing and volatility, both for Basrah and Kirkuk.²³ To have a sense of the magnitude of the coefficient, consider that a move from the twenty-fifth percentile of the standard deviation distribution to the seventy-fifth percentile (i.e., from a standard deviation of 0.89 to one of 1.77) would be associated with an increase in underpricing equal to \$0.80 for Basrah and \$0.26 for Kirkuk. Such an effect seems quantitatively large, especially for Basrah. Notably, after retroactive pricing is introduced, the correlation between underpricing and volatility disappears. The coefficients in the second row for both Basrah and Kirkuk are not statistically different from zero.²⁴

Overall, it appears that it was easier for Iraq to underprice its oil in periods of high price volatility than in periods of price stability. This finding is interesting for three reasons. First, it lends further credibility to our interpretation of underpricing. Second, it is useful in explaining the *timing* of the documented underpricing. Why was Iraq able to dramatically expand its underpricing in the spring of 2000? The evidence on the link between volatility and underpricing suggests that part of the answer may lie with the increased volatility of the price of oil in world markets. The months of peak underpricing (2000 and the first half of 2001) are months characterized by high volatility in the price of oil markets. By contrast, volatility was much lower in 1997, 1998, and 1999. Third, one might in theory be concerned

22. Using five weeks does not significantly change the results.

23. This association is stronger for Basrah than for Kirkuk. One possible explanation for this pattern is that Basrah is typically shipped to the United States, while Kirkuk is exported to European markets. Since it takes four to six weeks for Basrah to reach the United States from the Persian Gulf and only a few days for Kirkuk to reach its final destination in Europe (Kirkuk is shipped from a Turkish port on the Mediterranean), price volatility is likely to increase uncertainty about the final market price at the destination port by more for Basrah than for Kirkuk.

24. We also estimated models that include the price of Brent among the controls. The coefficients for Basrah become .445 (.174) and .122 (.131), while the coefficient for Kirkuk become .368 (.113) and -.0000 (.151). When we include year fixed effects, the coefficients for Basrah drop to .248 (.094) and .027 (.135), while the coefficient for Kirkuk become .346 (.093) and -.007 (.138).

that after the September 11 terrorist attacks, the Iraqis were worried that the United States could use the attacks as an excuse for heightened anti-Iraq actions and decided that continuing to solicit bribes from the Oil for Food program was too risky. We think that a more direct and likely explanation is that the end of underpricing was due to the adoption of the retroactive pricing scheme in September 2001. The finding that after retroactive pricing, the correlation between underpricing and volatility disappears is more consistent with our interpretation.

V.B. Underpricing and the Composition of Oil Buyers

In this subsection we investigate the relationship between underpricing and the importance of major oil companies as a fraction of Iraqi oil buyers. We stress that this evidence is very indirect and only suggestive, although it is consistent with the evidence presented so far. There were three types of buyers of Iraqi oil. First, there were major multinational oil companies that operate in Iraq as well as in many other oil markets. Examples include Shell, BP, Total (France), Agip (Italy), and Gasprom (Russia). On the other extreme, there were many individual dealers with no connection to reputable companies. These dealers were typically from countries such as Russia, Ukraine, or Switzerland, which are not major oil importers. Third, there were small but legitimate oil companies (not major multinationals). We expect that while a major multinational oil company might find it difficult to pay bribes, a small oil trader operating in Russia or Switzerland might be more willing to do so. Therefore, if greater underpricing was in fact associated with greater bribes, we should see a larger fraction of oil sold to obscure oil traders during periods when underpricing was more prevalent.²⁵

The top panel in Figure VIII shows the estimated total amount of illegal revenues, by period.²⁶ Most of the illegal revenues were obtained between 2000 and 2001. The middle panel presents the fraction of buyers of Iraqi oil that are individual traders (as opposed to corporations, both major and minor). This

25. The likelihood that the third type of buyers was willing to pay bribes is probably in between the likelihood for the first and the second type.

26. The evidence on buyers is based on data provided by the CIA's Iraq Survey Group [CIA 2004], which lists the buyers of Iraqi oil in each three- to six-month period. The largest number of contracts went to Russian traders. See Hsieh and Moretti [2005] for descriptive statistics.




FIGURE VIII

The Relationship between Illegal Revenues and Type of Buyers

In the top panel each bar shows the estimated total amount of illegal revenues in each period. In the middle panel each bar shows what fraction of the buyers of Iraqi oil are individuals (as opposed to corporations). In the bottom panel each bar shows what fraction of buyers of Iraqi oil are one of the 200 major oil companies (as defined by Forbes [2004]). The correlation between the top and the middle panel is .48. The correlation between the top and the bottom panel is $-.42$.

fraction is low in the years before 2000, and grows with time. It reaches a peak in the first trimester of 2001, and declines after that. A comparison with the top panel indicates that the first trimester of 2001 is the time when underpricing of Basrah reached its peak. The correlation between the degree of underpricing and the fraction of buyers that are individuals is .48. The bottom panel shows what fraction of buyers of Iraqi oil that are *major* oil companies, defined as the ones listed among the top 200 oil companies in the world. (Not all the companies are major oil companies, and so the fraction of major oil companies is not 1 minus the fraction of individual traders.) A large fraction of buyers are major oil companies in the earlier years of the program when underpricing was less prevalent. The share of major oil companies declines over time. Notably, there are no major oil companies among the oil buyers in the first trimester of 2001, which is exactly when underpricing reached its peak. The correlation between underpricing and the fraction of buyers that are major oil corporations is $-.42$.

VI. ESTIMATING THE TOTAL AMOUNT OF BRIBES TO IRAQ

In this final section we aim to quantify the amount of bribes that the Iraqi government may have received from this underpricing mechanism. Our best estimate of underpricing (based on the gap between the official price of Iraqi oil and the market price of the same oil) suggests that the rents generated through underpricing totaled \$3.5 billion.²⁷ Of course, not all this revenue ends up in Iraqi coffers. Presumably, the buyers (at least some of them) demand a share of the illegal revenues.

We have no direct information on the intermediary's share. However, we can use a simple monopolistic competition framework to estimate the split of the rents between the intermediary and the Iraqi regime. Specifically, if the demand for oil provided by each trader is isoelastic, the market price will simply be a constant markup over the trader's marginal cost:

$$(5) \quad P^{\text{market}} = \mu \times (P^{\text{official}} + \text{Bribe}).$$

Here, μ is the markup over marginal cost and is equal to $(\sigma - 1)/\sigma$ where σ is the elasticity of substitution between traders. The

27. That is, the total rent is the sum of the two entries in the bottom row in Table II.

important thing to notice about this expression is that the trader's marginal cost is the sum of the official price of Iraqi oil and the bribe paid to Iraq to obtain the underpriced oil. Therefore, if one knows the markup, the total amount of bribes obtained by Iraq could be estimated as

$$(6) \quad \sum_t \text{Bribe}_t = \sum_t \left[\frac{(\Delta P_t - \hat{\alpha}) + (1 - \mu)P^{\text{official}}}{\mu} \right] \times (Q_t \times \gamma_t),$$

where γ_t is the fraction of the relevant oil in week t that was bought by buyers willing to pay bribes. Here, we use the fact that we estimate $P^{\text{market}} - P^{\text{official}}$ as $(\Delta P_t - \hat{\alpha})$ (3).

We have no direct information on the markup charged by traders of Iraqi crude oil. However, we can get an approximate estimate by examining evidence on the standard markup charged by crude oil traders. The 1997 U. S. Service Census indicates that the markup of crude oil traders is 1.6 percent.²⁸ We therefore use $\mu = 1.016$ as our central estimate. This implies a very high elasticity of substitution between traders ($\sigma = 62$), which is consistent with the fact that oil is commodity. However, it is possible that the elasticity of substitution is lower than that implied by a markup of 1.6 percent. For this reason, we also provide alternative estimates of the total bribes under different assumptions on the size of μ .

In Table IV we present estimates of (6) under different assumptions on μ and γ . Our central estimate of the split of the \$3.5 billion rent between trader profits and bribes paid to Iraq (first entry in column (1) assumes that the dealers' markup is 1.6 percent and that $\gamma = 1$. In this case, the total amount of bribes obtained by Iraq would be \$2.60 billion, or 75 percent of the total rent. If the dealers' markup was instead 3 or 5 percent, the total amount of bribes obtained by Iraq would be lower.

However, not all buyers may have been willing to pay bribes. As we argue above, major oil companies may be concerned about their reputation and may find it difficult to pay bribes. Entries in column (2) show estimates of the amount of bribes obtained by Iraq under the assumption that only individual traders not affiliated with oil companies are willing to

28. U. S. Census Bureau [2001]. We define crude oil traders as "agents, brokers, and commission merchants for petroleum wholesale" (NAICS code 42272).

TABLE IV
ESTIMATED AMOUNT OF BRIBES OBTAINED BY IRAQ FROM UNDERPRICING

	If all buyers pay bribes ($\gamma = 1$)		If only individual traders pay bribes (γ is estimated)	
	Iraqi bribes (billion)	Iraqi share of total rent (%)	Iraqi bribes (billion)	Iraqi share of total rent (%)
	(1)	(2)	(3)	(4)
Dealer's markup is 1.6%	2.64	75	1.30	76
Dealer's markup is 3%	1.82	52	0.95	51
Dealer's markup is 5%	0.78	22	0.47	28

Entries in column (1) and (3) are estimates of (6). Underpricing is estimated using the difference between the market price of Iraqi oil and the official selling price of Iraqi oil. Entries in column (3) are obtained by estimating γ as the fraction of buyers that are individual traders in the relevant phase of the Oil for Food Program.

pay bribes. Those entries are obtained by applying (6), where γ_t is estimated as the fraction of buyers who are individual traders in the relevant period of the Oil for Food Program. Note that under this scenario, the major oil companies that were allocated contracts for underpriced Iraqi oil get to keep the entire rent. Individual traders, on the other hand, would buy Iraqi oil below market price, resell it at market price, but would have to split the rent with the Iraqi regime. Obviously, the Iraqi regime would prefer to deal only with individual traders. But it is possible that, in some periods, the supply of Iraqi oil exceeded the demand from traders that were willing to pay bribes to purchase underpriced oil. Alternatively, it is possible that Iraq felt compelled to deal at least in part with reputable oil companies to minimize international scrutiny.

On the basis of the estimates in column (2), the total amount of bribes obtained by Iraq would be between \$0.47 billion and \$1.3 billion, depending on the dealer's markup. Because 1.6 percent is our best guess of the magnitude of the markup, our estimate in row 1 provides our best estimate of the bribes obtained by Iraq: \$1.3 billion. The estimated bribes amount to 2 percent of the total value of oil sales during the Oil for Food Program. Compared with the amount of corruption typically observed in projects in some developing countries, a 2 percent of funds lost to corruption does

not appear to be particularly large.²⁹ However, it is possible that there were other irregularities in the Oil for Food Program that allowed Iraq to siphon funds from the program. For example, it is possible that overbilling took place in the purchase of humanitarian goods. We do not measure this in this paper, as we do not have a tractable way of measuring its magnitude. For this reason, the 2 percent figure should be interpreted as a lower bound.

How does our estimate of the total amount of bribes compare with the existing estimates? The CIA [2004] estimates that Iraq obtained \$230 million in bribes. This estimate is based on documentation provided by Iraq's state-owned oil company on the cash bribes paid by the oil buyers.³⁰ However, there are several reasons to believe that the CIA's data might understate the amount of the illegal cash payments. First, bribes may have also been paid to other branches of the Iraqi government (other than the oil company) and to Iraqi officials on a personal basis. The CIA estimate would miss these bribes. Second, the source of the CIA's data—Iraq's oil company—has an obvious incentive to minimize the extent of the bribes. This is a common problem with all self-reported data on illegal activities.

More importantly, it is possible that some recipients of the oil contracts (or their associates) compensated Iraq with political favors instead of monetary bribes. If Iraq was compensated by some buyers in the form of political favors, the CIA's estimate would understate the degree by which Iraqi oil was underpriced and, by extension, the illegal resources that Iraq might have been able to obtain through this scheme. For example, the head of the Oil for Food program is alleged to have obtained oil allocations for a shell company called AMEP. In exchange, the Iraqis wanted assistance in getting permission to use the Oil for Food program funds to purchase oil-industry equipment [IIC 2005a]. In this specific example, documentary evidence compiled by the Volcker commission indicates that a cash bribe was paid in only one of the five AMEP transactions and that

29. For example, Olken [2004a] finds that 29 percent of funds allocated to a road building project in Indonesia were stolen. In another program in Indonesia, Olken [2004b] finds that 18 percent of subsidized rice in a large antipoverty program in Indonesia is stolen. A study of hospital procurement in Argentina finds a 15 percent overcharging for inputs [Di Tella and Schargrotsky 2003]. A study of Uganda schools finds that 80 percent of the funds allocated by the Ministry of Education for schools never reached the schools [Reinikka and Svensson 2004].

30. The Volcker Commission [IIC 2005b] provides similar estimates based on the same data. The GAO [2004] estimates that Iraq obtained almost \$900 million in bribes. However, the GAO's estimate is not based on any evidence, but simply assumes that Iraq collected a 50-cent bribe on every barrel of oil.

the cash bribe equaled only a tenth of the rent generated by underpricing. In contrast, we remind the reader that our estimates are remarkably consistent with the Volcker commission's documentary evidence of underpricing in the AMEP oil sales. Therefore, our estimates based on the price gap are arguably a more accurate measure of the *sum* of the monetary bribes and the monetary value of the political favors that Iraq obtained through the deliberate underpricing of its oil.

VII. CONCLUSION

In this paper we find modest evidence that the official selling price of Iraqi oil was below its market value from 1997 through 1999, but clear evidence of underpricing from May 2000 through September 2001. In addition, there is little evidence of underpricing after September 2001, when the UN introduced a retroactive pricing scheme that made it more difficult for Iraq to underprice. Moreover, we find that it was easier for Iraq to underprice its oil in periods of high price volatility in world oil markets. This is expected, as detection of underpricing by the UN is more difficult in periods of high volatility. As expected, the relationship between underpricing and volatility disappears after the adoption of retroactive pricing. We also find a suggestive relationship between the degree of underpricing and the composition of the buyers of Iraqi oil. Periods of underpricing were periods in which a larger share of the oil was purchased by obscure oil traders rather than by multinational oil companies. The peak of the underpricing occurred at a time when no major multinational companies appear among the buyers.

Why did Iraq sell its oil below market price? We hypothesize that underpricing was a way for the Iraq regime to obtain illegal kickbacks or political favors from the oil buyers. Because all the legal oil revenues were controlled by the UN, the Iraq regime had an incentive to sell its oil below its market price in exchange for a bribe from the buyer. Our estimates suggest that Iraq created 3.5 billion dollars in rents by underpricing the oil. Using a simple monopolistic competition framework and an estimate of the elasticity of substitution between traders, we calculate that Iraq illegally received 1.3 billion dollars in bribes, or 2 percent of the oil sales under the Oil for Food program. This estimate should be interpreted as the sum of the direct cash payments and the monetary value of the political favors provided by the oil buyers.

Compared with the amount of corruption observed in projects in developing countries, this estimate appears to be small.

There are a number of important questions we do not address in this paper. First, we only focus on the bribes that Iraq might have been able to obtain through underpricing its oil, and not on other mechanisms by which Iraq could have obtained illegal resources (for example, overinvoicing the purchases of humanitarian goods). We focus on illicit income from oil sales, rather than on illegal revenue from other sources, because we have a tractable way of measuring the bribes Iraq might have been able to extract from the oil buyers.

Second, we clearly cannot say much about the benefit of the program. Although we find evidence of underpricing, the program did provide valuable aid that significantly alleviated the humanitarian crisis in Iraq. Furthermore, it appears that the UN were successful in reducing the scope of underpricing after September 2001. Indeed, it is remarkable how much a small change in the rules of the program—specifically, the introduction of retroactive pricing—was able to reduce the scope of underpricing. Therefore, one obvious lesson from the Oil for Food program is that incentives do matter, and careful institutional design on the part of international development agencies is crucial to minimize the amount of waste in humanitarian programs.

APPENDIX: DATA

Prior to September 2001, Iraq (with the approval of the UN's oil overseers) set the official selling price of Basrah Light and Kirkuk at the end of the month. The quoted prices are "free on board" (fob) prices for loadings in Ceyhan in Turkey (Kirkuk) or at Mina-ak-Bahr (now renamed as Basrah) in the Persian Gulf (Basrah Light) and apply to oil liftings over the subsequent month. The price of Iraqi oil is set as fixed discount to dated Brent for oil shipped to European markets, as a discount to second-month West Texas Intermediate for shipments to the United States, and as a discount to an average of Oman and Dubai for exports to Asia. There are two implications from the use of this pricing formula. First, the fob price clearly can differ depending on the destination of the oil. Second, although the pricing formula is known at the beginning of the month, the exact price is not. Specifically, because second-month WTI refers to the price of WTI in the first day of the second month after the oil is

loaded, the final price of oil shipped to the United States is only known one month after the oil has been lifted from an Iraqi port. In turn, dated Brent refers to the price of Brent no more than seven days in the future, the final price of oil shipped to Europe is known seven days after the oil has been shipped from Turkey. After September 2001, the exact discount to the price of the benchmark oils was only set at the end of the month, and so the buyers of Iraqi oil would not even know the pricing formula at the time the oil is lifted.

Starting in September 1998, the UN reports the average official selling price of Iraqi crude oil each week.³¹ The UN does not provide information on the individual prices of the two Iraqi crude oils, nor does it provide the fob prices for the different markets. For this reason, we obtained the official selling price of the two Iraqi crude oils from Platts.³² Because Basrah Light is typically exported to the United States, we take the official selling price for Basrah Light destined for the United States. In turn, Kirkuk is largely exported to Europe, and so we take the official selling price of Kirkuk destined for European markets.

We compare the official selling price of Iraqi oil with the market price of Arabian Light and Arabian Medium (Saudi Arabia) and Urals (Russia). The oil industry press frequently cites "Arab Light," a crude oil produced by Saudi Arabia, as the closest substitute for Basrah Light. In turn, a Russian crude oil known as "Urals" is widely viewed as a close substitute for Kirkuk. We define the price of the two Saudi crude oils as the fob price in the Persian Gulf (from ICIS-LOR, obtained from Datastream's database). In turn, the market price of Urals is typically quoted as a delivered (or cif) price in Italy (also from ICIS-LOR, provided through Datastream). We convert the delivered price of Urals in Italy to a fob price by subtracting oil tanker rates from the Black Sea to Italy.³³ The average transportation cost used for Urals between 1989 and 2002 is \$.87. During the program years, trans-

31. Available from <http://www.un.org/Depts/oip>.

32. We purchased these data from Platts. There are no price data in 1995 and 1996, because no trading took place. The data from the UN match reasonably well the properly weighted data from Platts for the years when both sources are available (mid-1998–2002).

33. For the years 1991–2002, we use weekly shipping data from Bloomberg. Unfortunately the Bloomberg series is not available prior to 1991. For the years 1987–1990, we use the monthly series from the annual issues of OPEC's Statistical Bulletin (OPEC). We could not find monthly level data for 1985 and 1986. We impute transportation costs in 1985 and 1986 based on the mean transportation cost observed between 1987 and 1990.

portation costs fluctuate between .58 in 1999 and 1.22 in 2000. Transportation costs are slightly higher during the Oil for Food program years. A regression of transportation cost on a dummy equal to one for the Oil for Food years yields .10 (.04). Transportation costs are also positively correlated with the price of oil.

Turning to the market price of the Iraqi oils, we combine data from three sources to obtain the spot market price of Basrah Light and Kirkuk. These three sources are ICIS-LOR, Platts, and Petroleum Argus.³⁴ For Kirkuk, oil traders will sell the oil contracts to the end users on a fob basis in Turkey. The end user will then be responsible for dispatching a tanker to Turkey to load the oil. The market price for Kirkuk provided by all three sources are therefore spot fob prices for loadings at the port of Ceyhan (in Turkey). For Basrah Light, oil-traders typically sell the oil on a delivered basis in the U. S. Gulf Coast. Therefore, ICIS-LOR and Petroleum Argus reports the cif price at the U. S. Gulf Coast. We convert the cif price to a fob price by subtracting the shipping rates for crude oil from the Persian Gulf to the U. S. Gulf Coast.³⁵ Platts also reports the cif price of Basrah Light in the U. S. Gulf Coast, but our sources that report Platts' data already adjust the cif price for shipping costs from the Persian Gulf. Unfortunately, no individual source for spot prices covers the entire period. To construct a series for the spot market price of the two Iraqi crude oils over the entire time period, we take the average of the price gaps from the data sources available in each week. We choose a simple average as the most transparent way of combining the three alternative series.

Finally, the web site of the UN Oil for Food program reports the weekly quantity of oil sold. Starting in mid-1998, these data are reported separately for exports through Ceyhan (Turkey) and Mina-ak-Bahr (on the Persian Gulf) for most weeks, but there are many weeks in which this breakdown is not provided publicly.

UNIVERSITY OF CALIFORNIA, BERKELEY AND NBER

34. The data from Petroleum Argus are obtained from weekly issues of "Petroleum Argus Weekly Global Markets." We obtained Platts data on the spot price of Iraqi oil from OPEC's annual statistical bulletin and weekly issues of the Energy Information Administration's "Weekly Petroleum Status Report." Finally, we purchased data on the spot prices of the two Iraqi oils from ICIS-LOR.

35. We obtained weekly series on shipping costs from the Persian Gulf to the U. S. Gulf Coast from Bloomberg.

REFERENCES

- Baltagi, Badi H., and Ping X. Wu, "Unequally Spaced Panel Data Regressions with AR (1) Disturbances," 1999.
- Central Intelligence Agency (CIA), Iraq Survey Group, Comprehensive Report of the Special Advisor of the DCI on Iraq's WMD (I), Washington, DC: Central Intelligence Agency, 2004. http://www.cia.gov/cia/reports/iraq_wmd_2004/.
- Di Tella, Rafael, and Ernesto Schargrotsky, "The Role of Wages and Auditing during a Crackdown on Corruption in the City of Buenos Aires," *Journal of Law and Economics*, XLVI (2003), 269–292.
- FAO, and WFP, Special Report: FAO/WFP Crop. Food Supply and Nutrition Assessment Mission to Iraq, 2003.
- Fisman, Ray, and Shang-Jin Wei, "Tax Rates and Tax Evasion: Evidence from 'Missing Imports' in China," *Journal of Political Economy*, CXII (2004), 471–496.
- General Accounting Office (GAO), Observations of the Oil for Food Program and Iraq's Food Security, Washington, DC: GAO, 2004. <http://www.gao.gov/cgi-bin/getrpt?GAO-04-880T>.
- Glaeser, Edward, "Public Ownership in the American City, in Urban Issues and Public Finance," in *Essays in Honor of Dick Netzer*, A. E. Schwartz, ed. (Northampton, MA: Edward Elga Publishing, 2004), pp. 130–162.
- Hsieh, Chang-Tai, and Enrico Moretti, "Did Iraq Cheat the United Nations? Underpricing, Bribes, and the Oil for Food Program," NBER Working Paper No. 11202, 2005.
- Independent Inquiry Committee into the United Nations Oil for Food Program (IIC), Interim Report, New York, NY: IIC, 2005a.
- , Manipulation of the Oil-For-Food Program by the Iraqi Regime, New York, NY: IIC, 2005b.
- Newey, Whitney K., and Kenneth D. West, "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, LV (1987), 703–708.
- Olken, Ben, "Monitoring Corruption: Evidence from a Field Experiment in Indonesia," Harvard University, 2004a.
- , "Corruption and the Costs of Redistribution: Micro Evidence from Indonesia," Harvard University, 2004b.
- Reinikka, Ritva, and Jacob Svensson, "Local Capture: Evidence from a Central Government Transfer Program in Uganda," *Quarterly Journal of Economics*, CXIX (2004), 679–705.
- Shleifer, Andrei, and Robert Vishny, "Pervasive Shortages Under Socialism," *RAND Journal of Economics*, XXIII (1992), 237–246.
- U. S. Census Bureau, Miscellaneous Subjects: 1997 Economic Census, Wholesale Trade, Subject Series (Washington, DC: U. S. Census Bureau, 2001).