Hedging Input Commodity Price Risk: An Equilibrium View

Hamed Ghoddusi, Sheridan Titman, Stathis Tompaidis [1,2]

Abstract

We present a model of a commodity processing chain with endogenously determined input and output prices and characterize the effectiveness of hedging policies that employ forward contracts on the price of the input. The model illustrates that the variance minimizing hedge ratio depends on operational characteristics, such as the convexity of the production function, and on economic variables – the elasticity of supply and demand and the relative magnitude of input and output risks; i.e., the size of supply and demand shocks. We find that the optimal hedge ratio can change over time as capacity utilization in the industry changes. To gauge the quantitative importance of our implications we estimate a parametrized version of our model for the crude oil to refined products supply chain using the simulated method of moments.

Keywords: Supply Chain, Spread Dynamics, Optimal Hedging, Simulated Method of Moments (SMM)

[1] Ghoddusi is with the School of Business, Stevens Institute of Technology, `hghoddus@stevens.edu`. Titman is with the McCombs School of Business, University of Texas at Austin, Department of Finance, `Sheridan.Titman@mccombs.utexas.edu`. Tompaidis is with the McCombs School of Business, University of Texas at Austin, Department of Information, Risk and Operations Management and Department of Finance, `Stathis.Tompaidis@mccombs.utexas.edu`.

[2] We thank participants at seminars at HEC Montreal, Oxford Institute for Energy Studies, UT Austin, UC Berkeley, Humboldt University, Rutgers, Alberta Finance Institute Conference, BI Norwegian Business School, Athens University of Economics and Business, Sharif University of Technology, UC Davis, HEC Paris, University of Vienna, UQAM Montreal, the conference on Commodities and Financialization at the Institute for Pure and Applied Mathematics, the 2012 INFORMS meeting, the Commodity Futures Trading Commission, and the Supply Chain Finance and Risk Management Workshop for comments. All remaining errors are ours.

I. Introduction

A typical commodity processing chain consists of three components: an upstream primary input – such as crude oil – a downstream output product – for example refined products like gasoline and heating oil – and a capital asset that turns the raw commodity into the processed commodity – for example a refinery. Commodity supply chains are ubiquitous – smelters transform bauxite into aluminum; power plants convert natural gas to electricity; tomato processing plants transform tomatoes into ketchup, airlines transform jet fuel into air travel; etc.

Hedging input (and output) commodity prices is a major issue for processors. Power plants are concerned with the rising cost of natural gas; airlines use variety of methods to mitigate jet fuel price risk; and food processors try to protect themselves against agriculture product price shocks. Among various risk-management strategies, financial hedging (e.g., the use of futures, options, and swaps) is the most popular one in practice. The existing literature typically treats the price of input and output commodities as exogenous variables for the commodity processing industry. However, the equilibrium prices of input and output commodities are formed as a result of endogenous decisions of both processors and consumers.

This paper develops an equilibrium model of prices and quantities in a commodity processing chain with uncertainties regarding supply and demand conditions. The model describes how characteristics of the input and output markets, such as supply and demand elasticities, and the degree of competition, properties of the production function, e.g., its convexity, and the relative magnitude of supply and demand shocks, influence input and output prices, the spread between the two, and optimal hedging policies.

We use the model to explore the endogenous dynamics of profitability and the effectiveness of hedging in a market with financial contracts only on inputs or output but not on both. For example, airlines can hedge the cost of jet fuel but not their revenues; food processors can hedge commodity inputs (e.g. coffee beans, corn, wheat, even eggs) but typically cannot hedge their outputs (e.g. instant coffee, bread, corn syrup, sauces). The opposite is also true for certain firms: futures contracts exist for base metals such as copper, zinc, and aluminum (the output of a smelting unit), but not for metal ores and chemicals used as inputs and producers of major crops can hedge their output price risks but not their inputs (e.g. water, labor, fertilizers or pesticides). For the application that we study, the oil refinery business, there are in fact forward contracts for both the input (oil) and some of the outputs, e.g., heating oil and gasoline. However, although the market for oil is quite liquid, the derivative markets for refined products is very illiquid for contracts exceeding two years.

The conventional wisdom is that the price risk associated with inputs can be hedged by buying the inputs forward. For example, a number of airlines make forward purchases of either fuel oil or oil to hedge their fuel price exposure. As our model illustrates, this view is largely based on the case where input price changes arise because of exogenous supply shocks; e.g., a rise in oil prices due to a conflict in the Middle East. When this is the case, an increase in the cost of production caused by a shock to the supply of the input decreases the spread between input and output prices and hence, the profitability of the producer. Our model illustrates that if supply shocks are the only source of uncertainty, the conventional view of hedging holds and the variance-minimizing hedge ratio increases when the convexity of the production function increases, and decreases when the elasticity of the demand for the output increases.

An analysis of the oil to refined petroleum products supply chain reveals that this conventional wisdom does not always hold. Figure 1 plots the correlation between the spread between refined petroleum products and oil prices (the crack spread) from 1985 to 2017. The correlations plotted at each point in time are calculated from the following 60 months. As the figure illustrates, the correlations start out negative, which is consistent with the conventional wisdom, and then become positive. In recent years the correlations are close to zero, indicating that producers could not have benefited from hedging input risk.

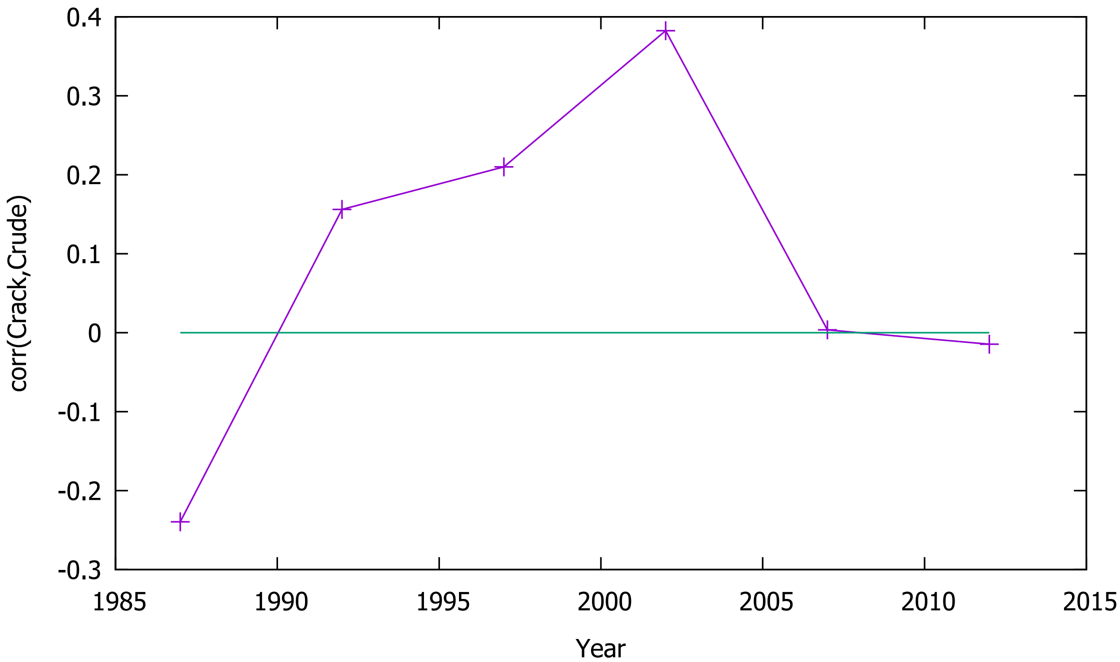

Figure 1: Correlation between Crack Spread and Crude Oil. The correlations are calculated using monthly data. Each correlation is calculated over a 60-month, non-overlapping, window. Correlations calculated between January 1987 and December 1991 are shown under the label for 1987; correlations calculated between January 1992 and December 1996 under the label for 1992; etc.

As our model illustrates, the conventional wisdom may have failed in this example because uncertainty comes from demand shocks as well as supply shocks. When the input cost increases because of a demand shock, e.g., because of a booming economy, the output price can move more than the input price, which means that spreads and input prices move in the same direction. As a result, input prices and the processing margin move together, which means that hedging demand shocks entails selling rather than buying the input forward. The optimal hedge ratio again depends on the convexity of the production cost function, but it is now the elasticity of the supply curve rather than the elasticity of the demand curve that determines the hedge ratio.

In the general case with both demand and supply shocks the correlation between profits, proxied by the spread, and input prices may be low, as illustrated in Figure 1, and it is generally not possible to perfectly hedge with just one financial instrument. However, the intuition developed in the case with a single source of uncertainty still holds. We show that increased supply uncertainty increases the amount of the input that a variance-minimizing hedger will purchase forward, while increased demand uncertainty reduces the magnitude of the forward purchase. The correlation between the spreads and the input prices, and thus the variance-minimizing hedge ratio, is still determined by the interaction between the convexity of the production function and the elasticities of the supply and demand functions.

We derive these results in a simple model where the relevant functions; i.e., the marginal production cost and the supply and demand functions, are all linear. This model can be solved in closed form and conveys our basic intuition. However, to a large extent, the questions that we ask are quantitative: what are the variance-minimizing hedge ratios, and do we expect them to change over time? How effective is hedging in plausible scenarios, i.e., to what extent can hedging reduce the variance of the producer's cash flows? Getting even rough answers to these questions requires a more realistic model that must be solved numerically.

Our more realistic model assumes demand and supply functions that are tailored for the crude oil to refined products supply chain. The main difference between our simple linear model and the model that we take to the data is that the estimated model allows for a more flexible production function that becomes more convex as utilization rates increase and we account for the possibility that production is subject to random production costs, caused for example by hurricanes or fluctuations in the price of natural gas. Using the simulated method of moments, we estimate the parameters of the model, i.e., the convexity of the production function, the volatility of the exogenous demand and supply factors, the demand elasticity of refined products, and the supply elasticity of crude oil.

Based on the estimates of these parameters, we run simulations that allow us to address issues that relate to hedging and its effectiveness. First, we consider the average percentage of the input quantity that is purchased forward in the variance minimizing hedge, and how this quantity changes as capacity utilization changes. We then compare the effectiveness of a dynamic hedging strategy that accounts for the changing levels of capacity with a strategy that simply holds the hedge ratio constant. Finally, we provide a more general analysis of hedging effectiveness, i.e., to what extent can firms reduce the variation of their cash flows by hedging input costs.

We find that in most of the scenarios we consider, hedging reduces cash flow variation by only a modest amount. For our base case, the optimal hedge ratio involves selling – rather than buying – oil forward, and hedging reduces the variance of profits by only 10%. There are two reasons why hedging looks to be ineffective. The first reason is that the correlation between oil prices and the crack spread is quite week because of the confounding effects of supply and demand shocks. Indeed, the correlation between oil prices and the crack spread is weakly positive for most of our sample, indicating that demand shocks are more important than supply shocks. The second reason is that matching the observed volatility of crack spreads requires relatively volatile noise in the portion of processing costs that cannot be hedged.

In the estimated base case, most of the systematic uncertainty comes from demand shocks. Our simulations show that in this case, if the exogenous sources of uncertainty can be eliminated, hedging can reduce cash flow variability by over 80%. However, in our comparative statics we show that when the magnitudes of supply and demand shocks have similar magnitudes, hedging effectiveness is quite low, even when the exogenous sources of uncertainty are eliminated.

The model also highlights the potential advantage of dynamic versus static hedging. A major difference between conventional statistical approaches to optimal hedging versus our structural approach is that the former is backward-looking: hedge ratios are estimated based on past values of realized prices. In contrast, our structural model illustrates how one can incorporate forward-looking information to dynamically adjust the optimal hedge. Example of forward-looking information that can be used include the implied volatility of options on input and output commodity prices, changes in the production function (due to new investments, technology improvements, and expected physical shocks), and the level of future supply and demand.

Using our calibrated model with exogenous shocks to the production function, we show in a simulation that improvement is modest in the base case. Both static and dynamic annual hedging deliver the same level of effectiveness – approximately 10%. The improvement is significantly better when the magnitude of exogenous shocks to the production function is small. The hedging effectiveness of a dynamically adjusted hedging strategy improves by up to 84%, while that of a static strategy by up to 76%. If structural parameters were to change; e.g., the variance of supply or demand shocks, then the benefit of dynamically adjusting the hedging strategy can be potentially even larger.

In summary, we offer the following novel contributions to the literature. First, we characterize the behavior of optimal hedging policies when there are shocks to both the supply of inputs and the demand for the downstream product. As we show, the conventional intuition that is based on shocks that only effect input prices often fail to hold in this equilibrium setting. Second, we offer a framework to highlight the drivers of time variation in dynamic hedge ratios and to incorporate forward-looking information in optimal hedging decisions. Finally, we use a novel calibration approach in the context of hedging and apply it to the case of the refinery industry to provide quantitative insights regarding the value of static versus dynamic hedging.

The rest of the paper is organized as the following. Section II provides a brief review of the relevant literature. We introduce the basic model in Section III. Section IV describes the theoretical results of the model. The quantitative exercise is introduced in Section VI.

II. Literature Review

Our model is most closely related to models developed by Hirshleifer (1988a), Hirshleifer (1988b), and Hirshleifer (1989), which explore the determinants of equilibrium spot and futures prices of a processed commodity. Hirshleifer (1988a) considers supply shocks and shows that the optimal hedge is long the input, while Hirshleifer (1988b) considers the case of demand shocks and shows that the optimal hedge is long the output. The papers focus on the difference between spot and futures prices, and the role played by transaction costs. Hirshleifer (1989) considers how the risk premium between the spot and the futures price varies with stock market variability, limited participation in the futures market, and the elasticity of supply and demand. Beyond the difference in focus, our framework differs in that we consider supply and demand shocks jointly, explore properties of the production function, and consider both competitive and monopolistic markets. We also characterize how the relative size of demand and supply shocks determine the optimal hedge ratio and hedging effectiveness, and derive results regarding the bounds on the ratio of the volatility of the input to the volatility of the output prices.

Carter et al. (2017) provide a recent review of the literature on commodity risk management. Kamara (1993) studies the optimal production and hedging decisions of the owner of a capital asset that can adjust its production plans. The paper shows that production decisions can be independent of risk preferences in the presence of futures contracts. Bessembinder and Lemmon (2002) consider the case of the electricity market and build an equilibrium model for electricity spot and forward prices. Unlike our paper, which focuses on the relationship between input and output prices, the focus of Bessembinder and Lemmon (2002) is on the spot and forward output price and how they are influenced by expected demand and uncertainty about shocks to demand. Another difference is that Bessembinder and Lemmon (2002) do not consider the feedback between a shock to the electricity demand and the price of the input; i.e., they assume that an increased demand for electricity does not affect the price of natural gas. Bessembinder and Lemmon (2002) show that the skewness in output prices is driven from the convexity of the production function but do not study the correlation between input and output prices and their spread, or hedging and its effectiveness.

Casassus et al. (2012) consider the possibility that the inputs and outputs can be stored and show that the ability to store commodities can influence correlations of prices and spreads. While storage is very important for understanding daily fluctuations of input and output prices, especially when demand or supply are seasonal, we expect it to be less important for understanding long term dynamics. For the case of the crude oil, storage accounts for approximately 3 months of consumption. We abstract away from storage and estimate our model using annual data.

Several papers study the behavior of the refinery industry. Wu and Chen (2010) apply a dynamic model of inventory and production to the refined products market. They also consider a market facing shocks to input and output. Dong et al. (2014) model the value of flexibility in the production function of the refinery industry. The focus of this literature is on the optimal level of endogenous production and inventory and its effect on price dynamics.

In contrast, our focus is on hedging and how it depends on supply and demand shocks and on the characteristics of the cost of production.

Our paper is also related to the literature on the dynamics of spreads and the valuation of spread options on a pair of commodities. Carmona and Durrleman (2003) introduce reduced-form pricing methods for spread options. Secomandi (2010) considers the effect of capacity constraints in natural gas pipelines. Our paper is related to this literature but, rather than a reduced-form approach, we build a micro-founded model to understand the dynamics of spreads.

III. Model

In the baseline model we consider a competitive industry which uses a capital asset to convert one unit of input to a unit of output. We consider the optimal production and profit dynamics of a *representative* owner of the capital asset; i.e., an owner of a continuum of small production units with different levels of efficiency; i.e., processing costs. We assume that the production function of the representative owner is proportional to the production function of the entire industry.

In addition to the production function, the primitives of our model are supply and demand functions of input and output. In the baseline model, we assume that the supply and demand functions are linear with respect to quantities, and that the production cost is quadratic with respect to quantity; i.e., the marginal cost of production is linear. We assume that there are no adjustment costs to production; that production is chosen once the levels of supply and demand have been observed and it is instantaneous; and that storage is not possible. These assumptions allow us to derive closed-form solutions – given supply and demand shocks the prices of the input and the output are determined endogenously from the elasticities of supply and demand and the convexity of the production function.

Table 1 summarizes the variables of our baseline model.

Symbol	Definition	Remark
P_C	Price of input	Endogenous
P_g	Price of output	Endogenous
Q^*	Optimal production quantity	Endogenous
X_S	Supply factor	Exogenous
X_D	Demand factor	Exogenous
γ_d	Demand elasticity	Constant parameter
γ_s	Supply elasticity	Constant parameter
$\mathrm{TC}(Q)$	Total cost of producing Q units	Convex in the production quantity
$\phi(Q)$	Capacity-related costs	Convex in capacity utilization
λ	Intensity of capacity-related costs	Capturing heterogeneity in production cost
$P_{I,g}$	Unit cost of other inputs	Constant parameter
F_C	Fixed cost of processing	Constant parameter

Table 1: Notations Used in the Model

A. Input Market

The price of a unit of the input; e.g. crude oil, is determined from a linear inverse supply function.

$$P_c(Q) = X_s + \gamma_s^{-1} Q \tag{1}$$

where $X_s(t)$ is the value of an exogenous, stochastic, supply factor at time t, Q is the supply of the input, and γ_s the elasticity of supply; i.e., the change in quantity supplied for a unit change in price. We assume that one unit of input is processed into one unit of output. Based on this assumption, the variable Q represents both the supply of the input and the quantity of the output. A positive supply shock, which corresponds to a decrease in the supply factor, X_s, decreases the price of the input and increases the quantity supplied.

B. Output Market

The price of a unit of output is determined by a demand factor X_d and the quantity of the output commodity Q. The inverse demand function is linear in quantity:

$$P_g(Q) = X_d - \gamma_d^{-1} Q \tag{2}$$

The demand factor X_d is exogenous and depends on long-term variables like the efficiency of the stock of existing appliances in use, and short-term shocks to income, taste, and seasonal factors. An increase in the demand factor X_d, increases the price as well as the quantity produced. We assume that demand and supply shocks are independent of each other.[3]

C. Output Production

We assume that the competitive industry is populated with many small representative firms which produce according to an increasing marginal cost schedule. This increasing marginal cost can be understood as the outcome of the deployment of production units on a *merit-order*, where the most efficient units of production are utilized first – as the production increases, higher marginal cost units are activated.

We assume that the marginal cost of production increases linearly with the amount produced; i.e., the total cost to produce Q units, $\mathrm{TC}(Q)$, is given by

$$\mathrm{TC}(Q) = F_c + QP_c + P_{I,g}Q + \phi(Q) = F_c + QP_c + P_{I,g}Q + \frac{\lambda}{2}Q^2 \tag{3}$$

where the total cost includes a fixed cost component, F_C, and a variable component that includes the cost of purchasing Q units of the input, QP_c, the cost of inputs other than the main input commodity, such as energy (electricity, natural gas, etc.), labor, materials, maintenance, etc., $P_{I,g}Q$,[4] and an increasing, quadratic, component $\phi(Q) = \lambda Q^2/2$, which

[3] In the case of the crude oil to refined products supply chain, supply shocks correspond to the discovery of additional supply, or to disruptions due to wars. Demand shocks correspond to unexpected economic growth, increased efficiency through new technologies, or the use of refined products in ways that were previously unanticipated. Given this intuition, it is reasonable to expect that supply and demand shocks are uncorrelated, or that, at least, their covariance is small enough that it can be safely ignored.

[4] We assume that the price of these inputs is not stochastic; i.e., it does not depend on the price of the input or the output. Assuming that their price is given by a random factor, creates an additional source of

represents the activation of higher marginal cost units.

D. Equilibrium: The competitive case

We assume that the output market is competitive, which implies that in equilibrium the market price of output equals the marginal cost. The equilibrium price and production satisfy

$$P_g = X_d - \gamma_d Q^* = \left.\frac{\partial \mathrm{TC}(.)}{\partial Q}\right|_{Q=Q^*} \tag{4}$$

The equilibrium quantity produced is given by:

$$Q^* = \frac{X_d - P_{I,g} - X_s}{\gamma_d + \gamma_s + \lambda} \tag{5}$$

The equilibrium price of the input and output, and their spread, are given by

$$\begin{aligned}
P_c &= X_s + \gamma_s^{-1} Q^* = \frac{(\gamma_d^{-1} + \lambda) X_s + \gamma_s^{-1} X_d - \gamma_s^{-1} P_{I,g}}{\gamma_d^{-1} + \gamma_s^{-1} + \lambda} \\
P_g &= X_d - \gamma_d^{-1} Q^* = \frac{(\gamma_s^{-1} + \lambda) X_d + \gamma_d^{-1} X_s + \gamma_d^{-1} P_{I,g}}{\gamma_d^{-1} + \gamma_s^{-1} + \lambda} \\
\text{Spread} = P_g - P_c &= X_d - X_s - \left(\gamma_s^{-1} + \gamma_d^{-1}\right) Q^* = \frac{\lambda (X_d - X_s) + (\gamma_d^{-1} + \gamma_s^{-1}) P_{I,G}}{\gamma_d^{-1} + \gamma_s^{-1} + \lambda}
\end{aligned} \tag{6}$$

The variances and covariances of input prices, output prices, the spread between output price and input price, and the variance of the quantity produced, are given by

volatility in spreads.

$$
\begin{aligned}
\mathrm{var}(P_c) &= \left(\frac{\gamma_d^{-1}+\lambda}{\gamma_d^{-1}+\gamma_s^{-1}+\lambda}\right)^2 \mathrm{var}(X_s) + \left(\frac{\gamma_s^{-1}}{\gamma_d^{-1}+\gamma_s^{-1}+\lambda}\right)^2 \mathrm{var}(X_d) \\
\mathrm{var}(P_g) &= \left(\frac{\gamma_d^{-1}}{\gamma_d^{-1}+\gamma_s^{-1}+\lambda}\right)^2 \mathrm{var}(X_s) + \left(\frac{\gamma_s^{-1}+\lambda}{\gamma_d^{-1}+\gamma_s^{-1}+\lambda}\right)^2 \mathrm{var}(X_d) \\
\mathrm{var}(P_g - P_c) &= \frac{\lambda^2}{(\gamma_d^{-1}+\gamma_s^{-1}+\lambda)^2}\left(\mathrm{var}(X_d)+\mathrm{var}(X_s)\right) \\
\mathrm{var}(Q^*) &= \frac{1}{(\gamma_d^{-1}+\gamma_s^{-1}+\lambda)^2}\left(\mathrm{var}(X_d)+\mathrm{var}(X_s)\right) \\
\mathrm{covar}(P_g - P_c, P_c) &= \frac{-\lambda(\lambda+\gamma_d^{-1})\mathrm{var}(X_s)+\lambda\gamma_s^{-1}\mathrm{var}(X_d)}{(\gamma_d^{-1}+\gamma_s^{-1}+\lambda)^2} \\
\mathrm{covar}(P_g - P_c, P_g) &= \frac{\lambda(\lambda+\gamma_s^{-1})\mathrm{var}(X_d)-\lambda\gamma_d^{-1}\mathrm{var}(X_s)}{(\gamma_d^{-1}+\gamma_s^{-1}+\lambda)^2}
\end{aligned}
\tag{7}
$$

IV. Properties of Prices and Spreads

Our model allows the study of the transmission of shocks across the supply chain. It helps us determine how properties of prices, such as the variance and the correlations of the price of input, the price of output, and their spread, depend on the convexity of the production function and the elasticities of supply and demand.

Proposition 1. *In a market described by Equations (1)-(5), the variance of the spread between the output price and the input price increases when the supply elasticity,* γ_s*, the demand elasticity,* γ_d*, and the coefficient of convexity,* λ*, increase; i.e.,*

$$
\begin{aligned}
\frac{\partial\, Var(P_g - P_c)}{\partial \gamma_s} &> 0 \\
\frac{\partial\, Var(P_g - P_c)}{\partial \gamma_d} &> 0 \\
\frac{\partial\, Var(P_g - P_c)}{\partial \lambda} &> 0.
\end{aligned}
$$

The proof of Proposition 1 follows from Equation (7).

Our next proposition illustrates the mechanism for risk transmission along the supply chain.

Proposition 2. *In a market described by Equations (1)-(5), the ratio of the variance of the price of the input to the variance of the price of the output is greater than 1 when only supply shocks are present, and smaller than 1 when only demand shocks are present. Moreover, with both supply and demand shocks, the value of the ratio is bounded by*

$$\frac{\gamma_s^{-2}}{(\gamma_s^{-1}+\lambda)^2} \leq \frac{var(P_c)}{var(P_g)} \leq \frac{(\gamma_d^{-1}+\lambda)^2}{\gamma_d^{-2}}$$

Proposition 2 follows from Equation (7), which shows that

$$\frac{\mathrm{var}(P_c)}{\mathrm{var}(P_g)} = \frac{(\gamma_d^{-1}+\lambda)^2\mathrm{var}(X_s)+\gamma_s^{-2}\mathrm{var}(X_d)}{\gamma_d^{-2}\mathrm{var}(X_s)+(\gamma_s^{-1}+\lambda)^2\mathrm{var}(X_d)},$$

which implies that the ratio increases (decreases) when the variance of the supply (demand) shocks increases.

Proposition 2 shows that the value of the ratio of the variance of the price of the input to the variance of the price of the output depends on the relative variance of supply shocks and demand shocks. If implied values for the variance of input and output prices are available, the proposition provides a mechanism to identify whether expected shocks are likely to be in the supply of the input or the demand of the output. Assuming that the elasticities of demand and supply, as well as the convexity of the production function remain constant, then an increase (decrease) in the value of the ratio can be attributed to an increase in the variance of supply (demand) shocks.

Proposition 3. *In a market described by Equations (1)-(5), an increase in the variance of the supply: a) decreases the covariance between the input price and the spread between the output price and the input price; b) decreases the covariance between the output price and the spread between the output price and the input price; and, c) increases the ratio of the* *variance of the input price to the variance of the output price. Similarly, an increase in the variance of the demand: d) increases the covariance between the input price and the spread between the output price and the input price; e) increases the covariance between the output price and the spread between the output price and the input price; and, f) decreases the ratio of the variance of the input price to the variance of the output price.*

The proof of the proposition follows from Equation (7).

Proposition 3 implies that when uncertainty is concentrated on the supply side, the input and output prices covary negatively with the spread, and when uncertainty is concentrated on the demand side, input and output prices covary positively. For the producer of the output, this behavior has an impact on whether to hedge by buying or selling forward contracts. The proposition also shows the relative importance of supply and demand shocks on prices. Parts c) and f) of the proposition show that supply shocks have a bigger impact on the input price, while demand shocks have a bigger impact in the output price.

A. The monopoly case

Up to this point we have assumed that producers are perfectly competitive price takers. In this section, we consider the case of a single producer that accounts for the effect of its production on output prices.

The monopolist maximizes the total profit $\pi = P_g(Q)Q - TC(Q)$ by considering the effect of its production decisions on the output market:

$$\max_Q Q(X_d - \gamma_d Q) - Q(P_{I,g} + X_s + \gamma_s^{-1}Q + \lambda Q) \tag{8}$$

The optimal amount produced by a monopolist is:

$$Q_{\mathrm{M}}^* = \frac{X_d - P_{I,g} - X_s}{2\gamma_d^{-1} + \gamma_s^{-1} + \lambda} \tag{9}$$

The equilibrium price of the input and output, and their spread, are given by

$$
\begin{aligned}
P_c = X_s + \gamma_s^{-1} Q_M^* &= \frac{(2\gamma_d^{-1} + \lambda) X_s + \gamma_s^{-1} X_d - \gamma_s^{-1} P_{I,g}}{2\gamma_d^{-1} + \gamma_s^{-1} + \lambda} \\
P_g = X_d - \gamma_d^{-1} Q_M^* &= \frac{(\gamma_s^{-1} + \gamma_d^{-1} + \lambda) X_d + \gamma_d^{-1} X_s + \gamma_d^{-1} P_{I,g}}{2\gamma_d^{-1} + \gamma_s^{-1} + \lambda} \\
\text{Spread} = P_g - P_c = X_d - X_s - \left(\gamma_s^{-1} + \gamma_d^{-1}\right) Q_M^* &= \frac{(\lambda + \gamma_d^{-1})(X_d - X_s) + (\gamma_d^{-1} + \gamma_s^{-1}) P_{I,G}}{2\gamma_d^{-1} + \gamma_s^{-1} + \lambda}
\end{aligned}
\tag{10}
$$

The following propositions summarize the difference between a competitive market and a monopoly.[5]

Proposition 4. *In a market described by Equations (1)-(3), the variance of the quantity produced by a monopolist is smaller than the variance of the quantity produced in a competitive market.*

Proof. Comparing Q^*, and Q^*_{M} from Equations (5) and (9), we observe that the numerators of the three equations are the same but the denominator is greater in the monopoly case. Thus, the level and the variance of quantity are both smaller in the case of a monopoly.

□

Proposition 5. *In a market described by Equations (1)-(3), the variance of the spread between the output price and the input price is smaller in a competitive market compared to the variance of the spread in a monopoly.*

Proof. We note that the variance of the spread between the output and the input price in the case of a monopoly is given by

$$
\mathrm{Var}(P_g - P_c)^{\mathrm{M}} = \left(\frac{\lambda + \gamma_d^{-1}}{2\gamma_d^{-1} + \gamma_s^{-1} + \lambda}\right)^2 (\mathrm{Var}(X_d) + \mathrm{Var}(X_s)) \tag{11}
$$

[5] Similar results can be derived for a firm that has monopsony power in the input market.

The result follows from the fact that

$$\frac{d}{dx}\frac{(\lambda+x)^2}{(\gamma_s^{-1}+\gamma_d^{-1}+\lambda+x)^2} = 2(\gamma_s^{-1}+\gamma_d^{-1})\frac{\lambda+x}{(\gamma_s^{-1}+\gamma_d^{-1}+\lambda+x)^3}$$

which is positive for $x \geq 0$.

$\square$

The intuition underlying Propositions 4 and 5 is that the production choices of the monopolist, compared to those of a producer in a competitive market, are less sensitive to both demand and supply shocks.

V. Minimizing variance of profits through hedging

Our results have implications for the optimal hedging program of the owner of the capital asset. We do not model the incentives of the firm to hedge. Instead, we assume that the objective of the commodity processing firm is to minimize the variance of its profit.

A. Profitability and Hedging

To determine the optimal hedge, we assume the firm has access to a single financial contract, a forward contract on the price of the input.

The firm's profit depends on two variables: 1) the profit margin $(P_g - P_c)$; and 2) the production quantity(Q).

$$\pi(Q) \approx Q(P_g - P_c) - QP_{I,g}$$

where $\pi(Q)$ is the profit for producing Q units of output – we have ignored costs that do not depend on the production quantity.

For a representative producer of the competitive industry, with a marginal cost that mirrors the marginal cost of the industry, the amount produced is proportional to the aggregate optimal amount produced, Q^*. As long as the price elasticity of demand is sufficiently small

– the gasoline market is one such case – the changes in quantity are negligible compared to changes in spreads. In such markets, the production quantity of a larger producer changes very little over time; however, its profit margin can vary significantly.[6]

When the changes in quantity are negligible compared to changes in the profit margin, the producer minimizes the variance of its profit by selling β units of the input to minimize the residuals in the relation

$$(P_g - P_c)_t = \alpha + \beta(P_{c,t} - F_{c,0}) + \epsilon \tag{12}$$

where $F_{c,0}$ is the futures price at time 0 for a unit of input delivered at time t. The optimal β, the hedge ratio, is given by the standard formula:

$$\text{hedge ratio} = -\frac{\text{covar}(P_g - P_c, P_c)}{\text{var}(P_c)} \tag{13}$$

B. Properties of the Variance-Minimizing Hedge Ratio

Equation (13) suggests that the variance-minimizing hedge ratio depends on the variance of the spread, the variance of the price of the input, and the correlation between the spread and the price of the input. Proposition 3 shows that, if uncertainty is driven by supply shocks, then the optimal hedge for the owner of the capital asset is to buy the input commodity in the forward market. On the other hand, if uncertainty is driven from demand shocks, the owner of the capital asset should sell the input commodity in the forward market.

In addition to determining the direction of the hedge; i.e., whether to buy or sell forward contracts on the input, our model offers guidance regarding the hedging amount, and the effectiveness of hedging.

The following proposition describes the optimal hedge ratio:

[6]For oil, the capacity utilization rate of the refinery industry only changes by a few percent. On the other hand, the profit margins can vary from \$6 per barrel to \$25 per barrel.

Proposition 6. *In a competitive market described by Equations (1)-(5), when demand is certain and supply is uncertain, an increase in the convexity coefficient, λ, increases the optimal amount bought; when only demand is uncertain, an increase in the convexity coefficient, λ, increases the optimal amount sold. The hedge ratio that minimizes the variance of profits is positive; i.e., long futures contracts, when demand is certain, but supply is uncertain, and negative when supply is certain and demand is uncertain. When both supply and demand are uncertain, the hedge ratio is bounded between*

$$-\frac{\lambda}{\gamma_s^{-1}} \leq \textit{hedge ratio} \leq \frac{\lambda}{(\lambda + \gamma_d^{-1})}$$

Proof. The results in Proposition 6 follow from the optimal hedge ratio

$$-\frac{\mathrm{covar}(P_g - P_c, P_c)}{\mathrm{var}(P_c)} = -\frac{-\lambda(\lambda + \gamma_d^{-1})\mathrm{var}(X_s) + \lambda\gamma_s^{-1}\mathrm{var}(X_d)}{(\lambda + \gamma_d^{-1})^2\mathrm{var}(X_s) + \gamma_s^{-2}\mathrm{var}(X_d)} \tag{14}$$

When demand is deterministic; i.e., $\mathrm{var}(X_d) = 0$, and supply is uncertain, the optimal hedge is to buy $\lambda/(\lambda + \gamma_d^{-1})$ forward contracts, which implies that the hedge ratio is less than one, but tends to one when the convexity coefficient approaches infinity.

It is easy to see, from Equation (14), that the hedge ratio decreases as the variance of demand increases relative to the variance of supply. In the limit when supply is deterministic; i.e., $\mathrm{var}(X_s) = 0$, and demand is uncertain, the optimal hedge is to sell λ/γ_s forward contracts. The bounds on the hedge ratio follow.

□

We note that, with only supply uncertain, the hedge ratio is between 0 and 1, while, when only demand is uncertain the hedge ratio can be very large in magnitude, especially when the convexity of the production function is large.

Two corollaries follow from the hedge ratio in Equation (14).

Corollary 1. *In a competitive market described by Equations (1)-(5), when demand is certain and supply is uncertain, the hedge ratio that minimizes the variance of profits decreases as the variance of demand shocks increases relative to the variance of supply shocks, and increases as the variance of supply shocks increases relative to the variance of demand shocks.*

Corollary 2. *In a competitive market described by Equations (1)-(5), the hedge ratio is positive when*

$$\frac{var(X_d)}{var(X_s)} > \frac{\lambda + \gamma_d^{-1}}{\gamma_s^{-1}}$$

and is negative otherwise.

In addition to changes in the optimal hedge ratio as uncertainty shifts from the input to the output, we can quantify the effectiveness of hedging using futures contracts for the input. We define hedging effectiveness as the reduction in the variance of the profits by hedging; i.e., the coefficient of determination, R^2, in the regression of the spread between profits and the value of futures contracts on the input.

Proposition 7. *In a competitive market described by Equations (1)-(3), the variance of hedged profits is zero, and hedging effectiveness is 100%, when there is a single source of uncertainty; either supply or demand shocks. If both supply and demand are uncertain, hedging effectiveness is zero when the hedge ratio that minimizes the variance of profits is zero.*

Proof. The coefficient R^2 is given by the square of the correlation between the profits of the firm and the value of the futures contract in the input.

$$R^2 = \frac{\left(-\lambda(\lambda + \gamma_d^{-1})\mathrm{var}(X_s) + \lambda\gamma_s^{-1}\mathrm{var}(X_d)\right)^2}{\left((\lambda + \gamma_d^{-1})^2\mathrm{var}(X_s) + \gamma_s^{-2}\mathrm{var}(X_d)\right)\lambda^2(\mathrm{var}(X_s) + \mathrm{var}(X_d))} \tag{15}$$

From Equation (15), it is obvious that for fixed variance of demand (supply), the numerator initially decreases (increases) and subsequently increases (decreases) as the variance of supply (demand) increases. When the variance of supply and demand are such that the hedge ratio is zero, the effectiveness of the hedge is also zero; i.e., there is no benefit in hedging the profits of the refiner with a futures contract on the input. Hedging effectiveness reaches 100% when either the variance of the supply shocks, $\text{var}(X_s)$, or the variance of the demand shocks, $\text{var}(X_d)$, is equal to zero. $\square$

Proposition 7 shows that hedging effectiveness is highest when there is a single source of uncertainty, either supply or demand, and that, when both demand and supply are uncertain, it is a non-monotonic function of the variance of supply or demand.

Another implication of Proposition 7 is that higher input or output volatility does not necessarily result in larger hedging positions; firms may indeed optimally reduce their hedging in response to higher volatility.

Our next proposition compares the hedge ratios of a monopolist to that of an operator in a competitive market.

Proposition 8. *In a market described by Equations (1)-(3), when there are only supply shocks, the hedge ratio that minimizes the variance of profits for a monopolist is greater than for an operator in a competitive market; i.e., the monopolist buys more of the input forward. On the other hand, when there are only demand shocks, the monopolist sells more of the input forward than an operator in a competitive market.*

Proof. The hedge ratio for a monopolist is given by

$$\frac{\lambda\gamma_s^{-1}\text{var}(X_d) - \lambda(\lambda+\gamma_d^{-1})\text{var}(X_s)}{\gamma_s^{-2}\text{var}(X_d) + (\lambda+\gamma_d^{-1})^2\text{var}(X_s)}$$

while the hedge ratio for an operator in a competitive market is given by

$$\frac{(\lambda+\gamma_d^{-1})\gamma_s^{-1}\text{var}(X_d) - (\lambda+\gamma_d^{-1})(\lambda+2\gamma_d^{-1})\text{var}(X_s)}{\gamma_s^{-2}\text{var}(X_d) + (\lambda+2\gamma_d^{-1})^2\text{var}(X_s)}$$

The proof follows from comparing the hedge ratios when the variance of the demand shocks, $\mathrm{var}(X_d)$, and the variance of the supply shocks, $\mathrm{var}(X_s)$ are equal to zero. With both supply and demand shocks, the hedge ratio for the monopolist is equal to zero when

$$\frac{\mathrm{var}(X_d)}{\mathrm{var}(X_s)} = \frac{\lambda + 2\gamma_d^{-1}}{\gamma_s^{-1}}$$

while, for an operator in a competitive market, it is equal to zero when

$$\frac{\mathrm{var}(X_d)}{\mathrm{var}(X_s)} = \frac{\lambda + \gamma_d^{-1}}{\gamma_s^{-1}}$$

□

VI. Quantitative Analysis of Hedging in the Refinery Industry

We apply our model to the specific case of the refinery industry: the supply chain that transforms crude oil to refined products through a refinery. This is arguably one of the most important supply chains in the world, influencing almost every industry. There are major refining hubs located in North America, Western Europe, the Middle East, South/East Asia, and significant recent growth in refining capacity in South America. The refined products market in the United States is the largest in the world, consuming almost 20% of the global crude oil production.

In Section III we examined a model with linear demand and supply functions and linear marginal production cost and were able to derive closed-form solutions. The setting that we explore in this section is more realistic, but must be solved numerically. There are multiple objectives for such a numerical analysis. First, we want to test whether the theoretical results, derived in a simple model, hold in a real-world application. Second, we are interested in determining the magnitude of hedging effectiveness in a particular context. For example, one of the theoretical lessons of the model is that in cases with both supply and demand shocks one may not be able to reduce variance very much by hedging. However, this is a quantitative question, and we need a calibrated model to illustrate this.

A. Data

Crude oil can be refined to an array of products, including gasoline, distillates – i.e. diesel fuel, jet fuel, and heating oil – and heavy, or residual, products, such as gas oil, lubricants, and asphalt. Gasoline, represents close to 60% of refinery revenue and is the most important product for the majority of refineries. Gasoline yield; i.e., the ratio of gasoline output to total refinery output ranges between 42% and 48%.

We use a weighted basket of gasoline and heating oil to proxy for the output of a refinery. The difference in the price of this basket, made of two parts gasoline and one part heating oil, and the price of the crude oil input, is called the crack spread, and is a proxy for the gross margins of the refining industry.[7] The same 3-2-1 ratio is used in the crack spread contracts traded in NYMEX and other commodity exchanges.

We use annual data on prices, quantities of crude oil refined, and refining capacity.[8,9,10] Our data spans the period from 1987 to 2017. We obtain prices for crude oil, gasoline, and heating oil from the Energy Information Administration (EIA).[11] We obtain prices for the total amount of crude oil refined and the global refining capacity from the statistical review report issued by British Petroleum.[12]

[7]The profit of individual refineries can vary. For example, refineries that can process heavy or sour crude benefit significantly from the price difference between these types and light/sweet crude oil. We do not model the heterogeneity in crude oil prices.

[8]Gasoline prices are typically higher during the summer in the northern hemisphere, while heating oil prices are typically higher during the winter – using annual prices avoids the potential seasonality.

[9]Annual prices are the average of monthly prices within the year.

[10]Our choice to use annual data somewhat mitigates the need to model storage. Storage is particularly relevant for explaining monthly data – for example it is common to transfer excess production of gasoline from the low demand months of winter to the high demand months of summer. However, storage from year to year is less common, potentially due to the limited amount of storage available – total, worldwide, available storage is approximately equal to 3 months of consumption.

[11]`https://www.eia.gov/dnav/pet/pet_pri_spt_s1_d.htm`

[12]`https://www.bp.com/en/global/corporate/energy-economics/statistical-review-of-world-energy.`

Statistic	Crude Oil	Gasoline	Heating Oil	Spread
Mean	54.12	64.27	64.35	10.12
Median	38.93	49.05	48.35	9.78
Maximum	117.56	128.17	136.13	15.51
Minimum	18.62	26.18	24.04	5.99
Standard Deviation	31.30	31.76	34.87	2.57
Skewness	0.90	0.79	0.90	0.27
Kurtosis	2.38	2.21	2.42	2.09

Table 2: Descriptive Statistics for prices of crude oil, gasoline, heating oil, and crack spread. Prices are real, deflated to 2012 dollars.

Figure 2 presents the time-series of the price of crude oil and the crack spread for the period 1987-2017. The prices are deflated using the consumer price index – the base year is 2012.

Table 2 provides descriptive statistics for annual prices and crack spreads. We use New York harbor gasoline, NYMEX heating oil, and Brent crude oil prices to calculate the crack spread. The values for the skewness and kurtosis suggest that the distributions, especially for the crack spread, are not symmetric, and that there is a large right tail.

Figure 3 presents a time-series of the ratio of the volatilities of the prices of crude oil and refined product. Similar to Figure 1, that illustrates that correlations between the price of the input and the price of crack spread vary over time, the relative volatility also varies over time. Both Figure 1 and Figure 3 suggest that the relative magnitude of supply and demand shocks varies over time – supply shocks are relatively larger in the late '80s and early '90s, while demand shocks dominate in the mid '90s and mid '00s.

B. Supply and demand functions

While the model in Section III allows for closed form solutions, to estimate a structural model for the crude oil to refined products supply chain we need a more general model.

Based on the crude oil production of low-cost members of the Organization of Petroleum

`html`

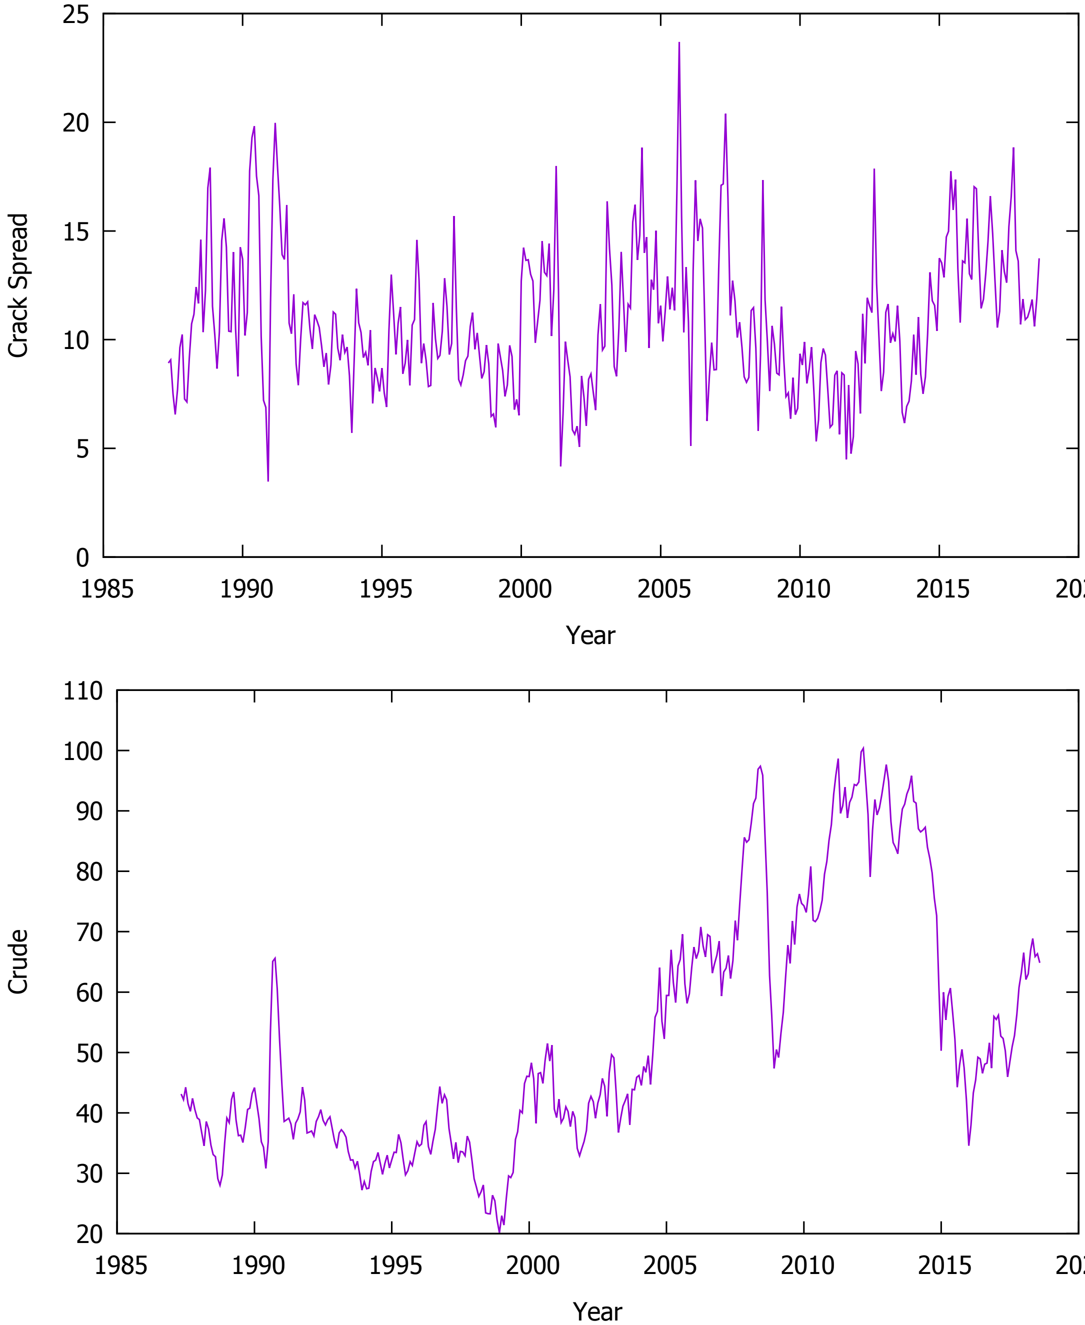

Figure 2: The figure presents the crack spread – top panel, measured in real terms, based in 2012 values and deflated using the consumer price index – and the price of crude oil – specifically the price of Brent, bottom panel. Prices and spreads are observed monthly.

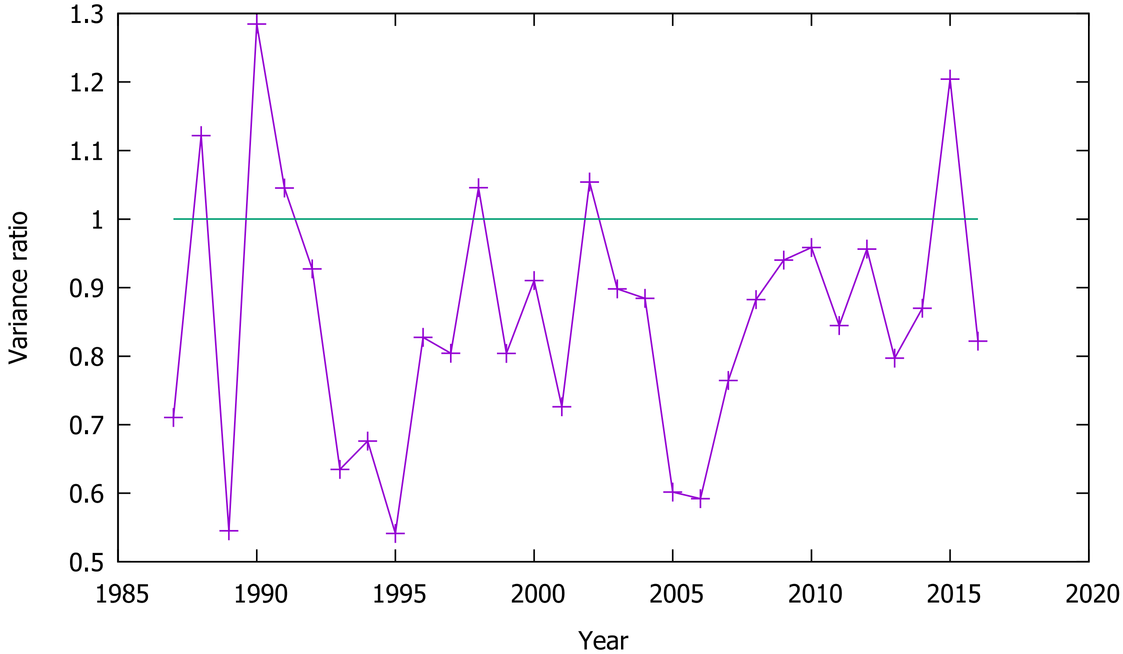

Figure 3: The figure presents the ratio of the volatility of the price of crude oil to the volatility of the price of refined products. The ratio is reported annually, with the volatilities calculated from the standard deviation of corresponding monthly returns within each year.

Exporting Countries (OPEC), we assume that the supply function is flat for low levels of production. The inverse supply function becomes steeper when the marginal supply moves to non-OPEC producers and unconventional sources. According to the literature; e.g. Dale (2016), the slope of the supply curve starts to increase at a quantity equal to 80% of global refining capacity. The inverse supply function is given by:

$$P_C = e^{X_s} + \left((Q - \underline{Q})\mathbb{1}_{Q>\underline{Q}}\right)^{\gamma_s} \tag{16}$$

where the factor, X_s, captures exogenous shocks to the global supply of oil, and shifts the inverse supply function; $\underline{Q}$ is the threshold where the steep part of the inverse supply function begins; and γ_s is the sensitivity of the marginal cost of crude oil to the level of production of crude oil at levels of production above the threshold.

We model the dynamics of the stochastic factor X_s as a mean-reverting process.

$$dX_s = \mu_s(\overline{X}_s - X_s)dt + \sigma_s dW_s \tag{17}$$

where the mean reversion rate μ_s, the long term level $\overline{X}_s$, and the volatility σ_s are assumed constant.

Consistent with the demand elasticity literature; e.g., Liu (2014) we use a constant-elasticity function for the demand of refined products[13]

$$P_G = e^{X_d} Q^{\gamma_d} \tag{18}$$

We assume that the stochastic demand factor, X_d, follows a mean-reverting process with shocks that are correlated with the shocks to the net supply of crude oil

$$dX_d = \mu_d(\overline{X}_d - X_d)dt + \sigma_d dW_d \tag{19}$$

where μ_d is the speed of mean reversion, X_d is the long term level of the demand factor, and σ_d is the volatility of the shocks to the demand of refined products.

Variable	Functional Form
Crude oil supply	$P_C = e^{X_s} + \left((Q - \underline{Q})\mathbb{1}_{Q>\underline{Q}}\right)^{\gamma_s}$
Supply shift factor	$dX_s = \mu_s(\overline{X}_s - X_s)dt + \sigma_s dW_s$
Refined products demand	$P_G = e^{X_d} Q^{\gamma_d}$
Demand shift factor	$dX_d = \mu_d(\overline{X}_d - X_d)dt + \sigma_d dW_d$
Marginal cost of production	$\mathrm{MC}_t = P_{I,g} + P_C + \phi_t(Q)$
Capacity-related costs	$\phi_t(Q) = \lambda_t Q^{\eta}$
Shocks to capacity	$\lambda_t = \overline{\lambda} e^{\sigma_\lambda \epsilon_t}$

Table 3: Supply, demand, and production functions for the crude oil to refined products supply chain.

[13]This functional form is equivalent to the log-log specification typically used to estimate gasoline demand elasticity in the literature.

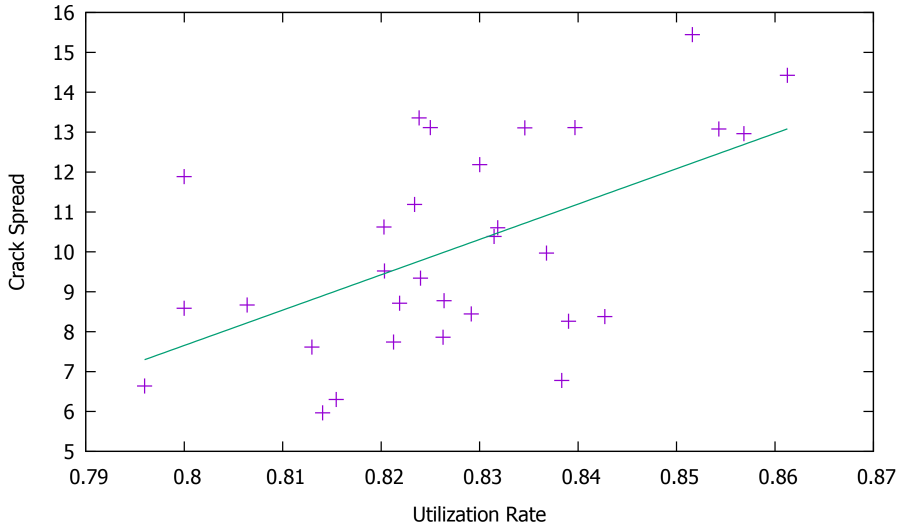

Figure 4: Crack Spread vs. Global Capacity Utilization. The crack spread values are expressed in real terms – deflated to 2012 prices using the consumer price index. The data are reported annually between 1987 and 2017.

C. Production Function

Figure 4 illustrates the relationship between global capacity utilization and the level of crack spreads. In a competitive market the crack spread equals the marginal production cost – the figure suggests that the marginal production cost is increasing as capacity utilization increases. The figure includes the linear least squares fit – a linear function is a close match to the data. The coefficient of determination, R^2; i.e., the percentage of the variance in the crack spread explained by capacity utilization is 32%.

To accommodate the residuals in Figure 4, potentially due to changes in production technologies, the delay between planning and building a refinery which results in shortages and oversupply, and exogenous shocks; e.g., hurricanes, that affect production, we assume that the production function is influenced by random fluctuations.

We note that, during the period we study, the amount of crude oil refined has increased significantly. However, Figure 4, suggests that the cost to refine a certain percentage of the total refining capacity remains relatively stable, possibly due to improvements in refining technology. To account for this feature, rather than expressing the quantity of crude oil that is refined each year in absolute terms; e.g., number of barrels of oil per day, we instead express it as a percentage of refining capacity.

Given the information in Figure 4, we assume that the marginal cost of refining Q units of crude oil in period t is given by

$$\mathrm{MC}_t = P_{I,g} + P_C + \phi_t(Q) \tag{20}$$

where $P_{I,g}$ is the cost of other inputs, P_C is the price of crude oil, and $\phi_t(Q)$ are costs that are capacity-related. The capacity-related costs, $\phi_t(Q)$, are assumed to be given by a power function, with two unknown coefficients λ_t and η.

$$\phi_t(Q) = \lambda_t Q^{\eta} \tag{21}$$

We assume that the convexity of the production function, described by the power η, is constant. However, to account for the residuals in Figure 4, we assume that the coefficient λ_t is a random variable drawn from the following distribution

$$\lambda_t = \overline{\lambda} e^{\sigma_\lambda \epsilon_t} \tag{22}$$

where $\overline{\lambda}$ is the baseline value for the parameter, σ_λ is the standard deviation of shocks to the parameter, and ϵ_t are standard, normally distributed, i.i.d. random variables.

The fluctuations in the production function have significant empirical consequences: for example, they influence the volatility of the crack spread. Given the i.i.d. nature of the random variables, ϵ_t, we expect that the marginal production costs will be autocorrelated.

D. Caveats

Our data consists of a time-series of annual prices and quantities between 1987-2017 – a total of 31 observations. It is clear that the small number of observations limits our choices.

There are several, potential, improvements that we could attempt with a longer time-series.

Supply Function of Crude Oil. We use a two-piece power function to model the increasing marginal cost of crude oil. In reality, the crude oil cost curve is much more complicated. In addition, several papers in the literature argue that the oil sector contains dynamic and delayed responses, see Wirl (2008).[14]

Time-Varying Variables. Given the small number of observations, we assume that the parameters of our model are constant. In particular, we assume that the parameters for the production function for the refinery sector, the demand for refined products, and the production function of crude oil are constant. However, refining technology has significantly improved, and our adjustment of using quantity produced as a percentage of refining capacity may not be accurate. In addition, there is evidence that consumer demand elasticity for gasoline has changed over time, see Hughes et al. (2006). Finally, the supply curve for crude oil has shifted significantly due to several factors, including the introduction of unconventional energy sources.

Storage. Our model is an instantaneous production model, with no inter-temporal storage of refined products or crude oil. In reality, both crude oil and refined products are stored. In the presence of storage, the optimal production will be less sensitive to demand shocks and more sensitive to input shocks because the refiners can produce in low-demand periods and store the refined products in order to sell them in periods with higher demand, and possibly with more expensive input. Since we do not model storage, our estimate of the elasticity of demand may be biased.

Shut-Downs. Refiners need to stop operation for necessary maintenance. The planned shut-downs generate patterns in the production rate that we do not capture. A maintenance

[14] For example, the presence of price adjusting mechanisms, such as the Organization of Petroleum Exporting Countries, prevent the oil market from following a static cost curve.

schedule chosen strategically; e.g., plan shut-downs when demand is expected to be low, would generate endogenous capacity adjustments. For example, in years in which oil prices are high, refiners may decide to postpone maintenance activities and keep the average operable capacity at a higher level. Since most of the planned shut-downs happen within the year, using annual data helps mitigate this potential bias.

E. Estimation

We use the Simulated Method of Moments (SMM) to estimate the structural parameters of the model. Starting with an initial guess of the parameters, we simulate several scenarios for the model. We calculate the model-generated moments for each scenario, and then average over all the scenarios. We compare the model-generated moments to the empirical ones and modify the initial guess of the model parameters until the percentage difference between the model-generated moments and the empirical moments is smaller than a cut-off. Additional details on SMM can be found in Strebulaev and Whited (2012).

The empirical moments we use in the SMM procedure, are the mean and standard deviation of the price of refined products; the mean and standard deviation of crude oil prices; the annual autocorrelation of the price of refined products; the correlations between the price of refined products and the crack spread, the price of refined products and crude oil, and the price of crude oil and the crack spread; the mean and standard deviation of capacity utilization; the mean of the crack spread; and the explained variance of the crack spread. Overall, we use 12 empirical moments.

The model parameters are: the demand and supply elasticities; the mean-reversion rates for supply and demand shocks; the standard deviation of supply and demand shocks; the long-run level of the supply and demand factors; the degree of convexity and the coefficient of the convexity term in the production function; and the standard deviation of fluctuations to the output from the production function. The total number of parameters is 11, implying that our model is over-identified.

F. Starting Values

A subset of parameters can potentially be estimated directly. This is possible by either using estimates provided in the literature, or through direct identification. In order to improve the performance of the SMM estimation, we use starting values for every parameter.

Elasticity of Demand. The literature on demand for refined products; especially gasoline, reports a wide range of values for short-term price elasticity of demand between 0.00 to -0.15 for different countries – see Cooper (2003), and Hughes et al. (2008).[15] We choose a starting value of γ equal to -20.

Supply Elasticity. In anticipation of, and during the disruption in 2011 to the supply of Libyan crude oil to the global markets, prices increased by close to $10 per barrel. Given that Libyan crude production is close to one million barrels per day, we use a starting value for supply elasticity, that corresponds to the assumption that an additional production of a million barrels of crude oil per day increases the price of crude oil by $10 per barrel. This assumption corresponds to a starting value for supply elasticity of $\gamma_s = 1.45$.

Stochastic Processes for Supply and Demand Factors. We use statistical properties of the historical prices of crude oil and refined products to determine the starting values for the long-term levels, volatility parameters, and mean-reversion rates of the supply and demand processes.

[15] Note that our elasticity parameter γ is the inverse of typically reported elasticity parameters. The estimates in the literature are based on price elasticity for a particular country or region. Since our model reflects global demand, it is possible that the value of the estimated demand elasticity may be outside the range given in the literature.

Cost of Processing. Based on industry estimates, we assume that the processing cost, net of the cost of crude oil and the nonlinear, convex, term, is \$3 per barrel.[16] Going forward, we report the crack spread net of this processing cost.

Convexity of the Production Function. Figure 4 suggests that the relationship between the deflated crack spread and the global capacity utilization is close to linear. Since, in a competitive market, the crack spread corresponds to the marginal cost of the production function, we choose a starting value for the power of the convex part of the marginal cost to be $\eta = 1$, corresponding to quadratic production costs with respect to the quantity of oil refined. To estimate a starting value for the coefficient λ we use the intercept of the linear relationship between crack spreads and capacity utilization, which provides a starting value equal to 3.2. To approximately match the explained variance of the crack spread, we set the initial value of the standard deviation of the random variables in the production function to 10%.

Table 4 lists the 11 parameters we estimate, their starting values, and the values that result in model-generated moments that best match the empirically observed moments.

Simulation Parameters. For the SMM procedure, we use 100 iterations. For each iteration, we set the number of time periods of the simulated vector T to 200 – each period corresponds to one year. We discard the first 100 periods in each iteration and use the next 100 periods to estimate the various moments.

We report the performance of the model in matching empirically observed moments in Table 5.

We observe that the model provides a reasonable match for the level and standard deviation of the price of crude oil, the price of refined products, and the crack spread. The autocorrelation between the prices of crude oil, refined products, and crack spread, are also

[16]See `https://www.iea.org/media/omrreports/Refining_Margin_Supplement_OMRAUG_12SEP2012.pdf`.

Notation	Parameter	Starting Value	Source of Starting Value	Estimated Value
γ_d	Demand elasticity	-20	Literature	-22.4
γ_s	Quantity sensitivity of crude oil price	1.45	Change of $10/barrel per extra 1M barrels	2.07
μ_d	Mean-reversion rate for demand (year^{-1})	0.3	Statistics of prices	0.11
μ_s	Mean-reversion rate for supply (year^{-1})	0.3	Statistics of prices	0.05
σ_d	Standard deviation of demand	0.3	Statistics of prices	0.53
σ_s	Standard deviation of supply	0.3	Statistics of oil prices	0.14
$\overline{X}_s$	Long-run supply factor	3	Statistics of crude oil prices	2.90
$\overline{X}_d$	Long-run demand factor	100	Statistics of refined products prices	101.24
η	Convexity of marginal cost function	1.00	Regression of crack spreads on capacity utilization	1.85
λ	Coefficient of marginal cost function	3.42	Regression of crack spreads on capacity utilization	0.0022
σ_λ	Shocks to Capacity	0.1	Regression of crack spreads on capacity utilization	0.07

Table 4: List of Parameters

Moment	Empirical Value	Model-Generated Value	Weight
Mean Refined Prices	64.27	64.85	20
S.D of Refined Prices	31.76	31.23	10
Mean Crude Oil Prices	54.12	54.62	20
S.D of Crude Oil Prices	31.30	30.93	10
Mean of Crack Spreads	7.12	7.23	20
Explained Volatility of Crack Spreads	0.10	0.10	20
Average Capacity Utilization	82%	80%	10
S.D of Capacity Utilization	1.7%	2.4%	5
Autocorrelation of Annual Refined Prices	0.86	0.86	5
Correlation of Crack Spread and Refined Prices	0.39	0.27	5
Correlation of Crack Spread and Crude Oil Price	0.26	0.23	5
Correlation of Crude Oil and Refined Prices	0.99	0.99	5

Table 5: Moments Used in the Simulated Method of Moments Estimation

accurate in terms of the sign and reasonably close in terms of size.

The worst match occurs for the standard deviation of capacity utilization. The model-generated standard deviation of capacity utilization is 50% larger than the empirical value.[17]

Our results suggest that demand shocks are the major source of shocks during the period we study. This finding is in line with recent empirical results; e.g., Kilian (2009), which decompose the volatility of the crude oil prices to supply and demand driven shocks and conclude that the share of demand shocks in explaining the volatility of crude oil prices has increased over time.

[17]This mismatch is largely due to the method trying to match the model volatility of crack spreads with the empirical value. Since we do not account for certain shocks; e.g., shocks to the prices of energy and materials used by in the refining process, the method increases the volatility of capacity-utilization to match the volatility of crack spreads with the empirical value.

VII. Hedging Effectiveness

Our calibrated model can provide quantitative guidance regarding the optimal hedging policies for a refiner that can only access the forward market on input (i.e. crude oil).[18]

The optimal number of forward contracts required to hedge the refiners profits is given by

$$\frac{Q_F}{Q} = -\frac{\text{covar}(P_g - P_c, P_c)}{\text{var}(P_c)}$$

The correlation of profits with input or output prices can change because of changes in the level of supply and demand, as well as changes in their variances.[19]

To evaluate the effectiveness of hedging, we generate 1000 random scenarios for demand and supply shocks for the next year. The variance of realized refinery profits, defined as the crack spread plus the pay-off from hedging, is then calculated under different hedging policies.

A. Baseline

We first consider the effectiveness of hedging using our model for the calibrated set of parameters. We find that, for the basecase, since demand shocks are relatively larger than supply shocks, to hedge profits, oil refiners with a production cost that reflects the production function of the entire industry, should sell forward a small number of contracts of crude oil. Without hedging, we find that the variance of next-period profits, in dollars per barrel per unit of production is 26.0%, while the variance of hedged profits is 23.5%, an improvement

[18]The effectiveness of optimal hedging strategies for refiners using both input and output forward contracts has been discussed by several recent papers; e.g., Ji and Fan (2011), Alexander et al. (2013), and Liu et al. (2017). Rather than use a structural approach, these papers use statistical estimation to compare various strategies using both forward contracts for both crude oil and refined products to hedge the crack spread.

[19]We note that our results assume that the decision-maker knows the true parameters of the model and that the model is correctly specified. Otherwise, the calculation is subject to both sampling and specification errors. The literature on optimal decision making under parameter uncertainty; e.g., Smith and Winkler (2006), shows that a rational decision-maker, in the face of sampling and specification errors, may refrain from taking the suggested optimal decision.

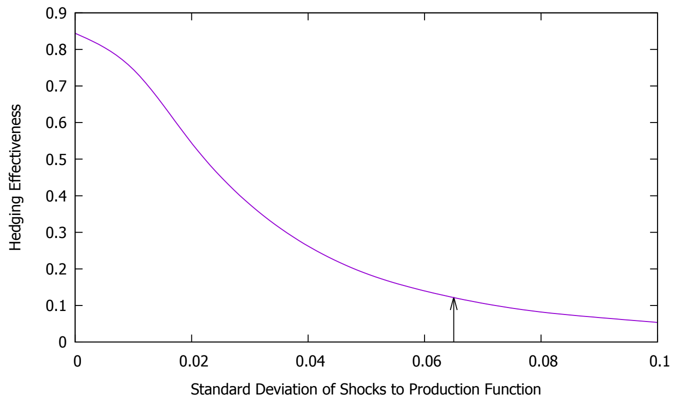

Figure 5: Impact of Exogenous Shocks on Production Function on Hedging Effectiveness. The vertical line with the arrow corresponds to the base case of the calibrated model.

of approximately 10%.

Hedging effectiveness is low for two reasons: the hedge ratio is small, due to both supply and demand shocks; and, there are fluctuations to the production cost function that further diminish hedging effectiveness.

We plot the impact of the volatility of shocks to the production function on the effectiveness of the hedging strategy in Figure 5. We note that hedging effectiveness can be above 80% if the fluctuations in the production function are small. Hedging effectiveness quickly deteriorates as the size of the fluctuations increases.

B. Dynamic versus Static Hedging

The optimal hedge ratio in our structural model varies over time for two reasons: endogenously – following the realization of supply and demand shocks, the optimal quantity produced, the crack spread, the correlation between the crack spread and the input price, and the hedge ratio change; and exogenously – if, for example, the variance of supply or demand shocks changes, so does the hedge ratio.

We consider the variation due to endogenous changes and quantify the impact on hedging effectiveness in two ways. First, we consider the cross-sectional variation of the hedge ratio and hedging effectiveness as the variance of supply and demand change. Figure 6 shows

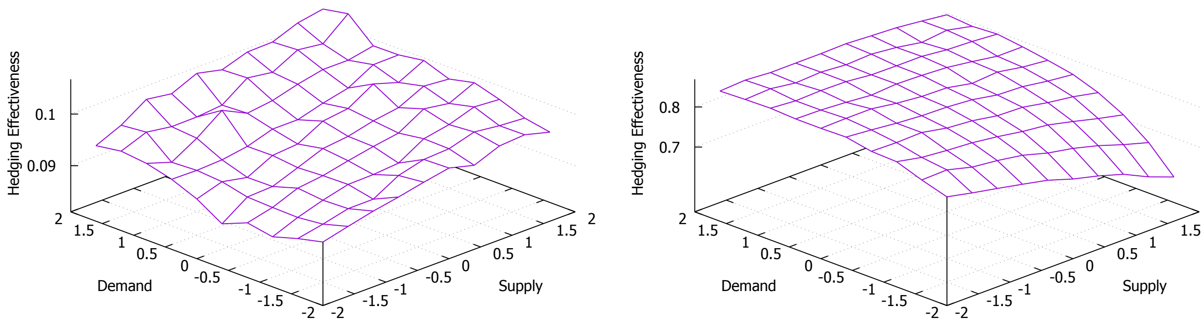

Figure 6: Hedging effectiveness vs. supply and demand shocks over a single year. The left panel shows hedging effectiveness for the baseline case. The right panel shows hedging effectiveness in the case without exogenous shocks to the production function. The units for the supply and demand shocks are in terms of standard deviations away from their long term mean.

the optimal hedge ratio and hedging effectiveness, over a single year, for different levels of the initial realization of demand and supply shocks, keeping all structural parameters of the model constant. The value of the supply and demand shocks varies within two standard deviations of their long-term level. The figure illustrates that the presence of fluctuations to the production function reduces hedging effectiveness significantly across all values of supply and demand shocks. We note that the fluctuations, evident in the left panel, are due to the large variance of the results – the values of hedging effectiveness are not statistically different across the values of the supply and demand shocks displayed. On the other hand, when random fluctuations to the production function are eliminated, the right panel illustrates that hedging effectiveness is high

Without fluctuations to the production function, hedging effectiveness depends on capacity utilization. Since, in the calibrated model, the production function is convex, when capacity utilization is low convexity is low, and the correlation of spreads and input prices is close to zero. On the other hand, when the capacity utilization is high, the correlation and hedging effectiveness can be much higher.

The second way we quantify hedging effectiveness is across time. Starting with the calibrated values for the parameters and the levels of supply and demand, we simulate 500 paths, each for 100 years, and evaluate hedging effectiveness by comparing the cumulative variance of profits for a dynamic strategy, where the hedge ratio is updated each year, and a static strategy, where the hedge ratio is set once, based on the long-term level of the supply and demand shocks. The results suggest that, relative to static hedging, dynamic hedging decreases the variance of profits only when fluctuations to the production function are small. In particular, without fluctuations to the production function, the cumulative variance of profit is reduced by 84% using dynamic hedging, and 76% using static hedging. With the calibrated level of fluctuations to the production function on the other hand, both strategies reduce the cumulative variance of profit by 10%.[20]

C. Comparative Statics

Our results suggest that hedging effectiveness is high when there is a single source of uncertainty; either supply or demand shocks; when fluctuations to the production function are small; and when the convexity of the production function is large.

Figure 7 explores hedging effectivess as the volatilities of supply and demand change. The figure shows that hedging effectiveness is non-monotonic. For example, as the volatility of supply increases, the hedging effectiveness initially drops and then starts increasing. The intuition behind this result is that the correlation between crude oil prices and crack spreads, when the volatility of supply is very low, is large and positive, and hedging effectiveness high. As the volatility of supply increases, the correlation drops. At some point – which is a function of demand and supply volatility as well as other, structural, parameters – the correlation of the crack spread and the price of crude oil, the hedge ratio, and the hedging effectiveness

[20] Our simulation assumes that the model parameters are constant. The benefit of dynamic hedging can hedging can potentially be significantly higher when volatitilies of supply and demand change, and can be determined from the price of forward looking financial instruments; e.g., the implied volatility of option prices.

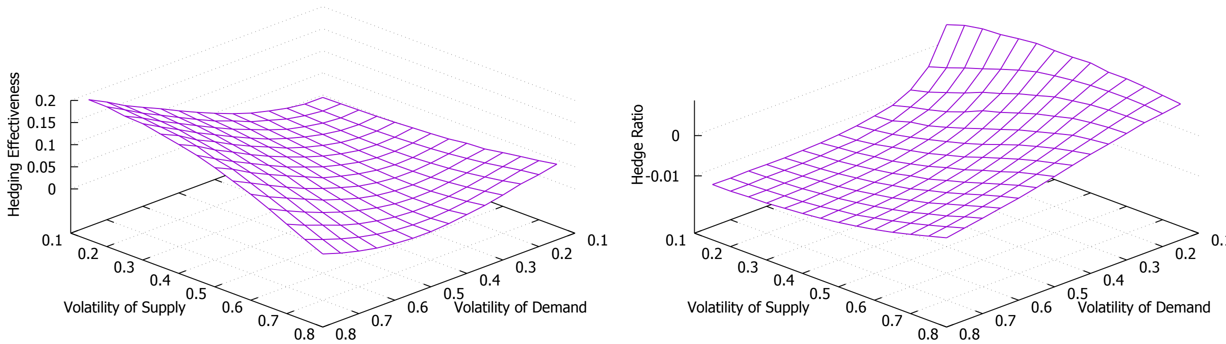

Figure 7: Hedging effectiveness (left panel) and hedge ratio (right panel) against the volatility of supply and demand shocks.

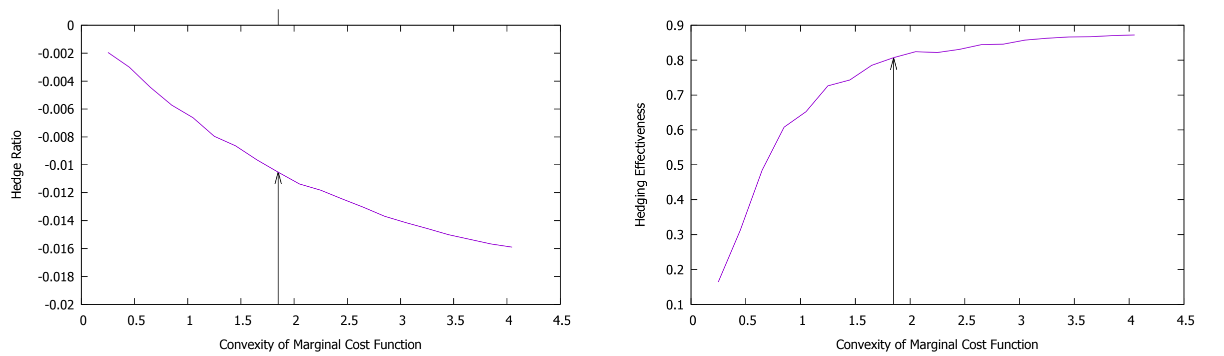

Figure 8: The left panel shows the hedge ratio that minimizes the variance of profits against the convexity of the marginal cost function. The right panel shows the corresponding hedging effectiveness. The vertical lines with the arrow correspond to the base case of the calibrated model

become zero. As the volatility of supply increases further, the correlation between crude oil price and crack spreads becomes negative, and hedging effectiveness increases.

Figure 8 shows the optimal hedge ratio (left panel) and hedging effectiveness (right panel) as the convexity of the production function varies. When the convexity of the production function is low, the crack spread is less volatile, and its correlation with input and output prices, as well as the hedge ratio drop. When the production function becomes more convex, the hedge ratio increases in absolute terms. The right panel of Figure 8 shows that hedging effectiveness increases as the convexity of the production function increases.

VIII. Conclusions

There is a substantial literature that examines the motivations for firms to reduce the variance of their cash flows by hedging.[21] This literature implicitly assumes that firms can in fact effectively hedge. In this paper we have taken a careful look at the economic factors that affect the effectiveness of hedging in an equilibrium model with endogenous input and output prices that are determined by stochastic supply and demand shocks as well as by the characteristics of the production process that transforms the input into the output. We estimated our model with data on the oil refinery industry to get a calibrated base case and provided numerical comparative statics on the determinants of hedging effectiveness. As we show, hedging is not expected to be particularly effective in most plausible cases.

To summarize, our analysis generates the following insights: 1) the source of shock (demand versus supply) is a key determinant of the direction of the co-movement between profits and input/output costs and consequently optimal hedging position of firms. In industries where most shocks come from the supply side, the firm should take a short position on the futures contracts for the input; when most come from the demand side, the optimal hedging requires a long position in input commodity; 2) there are natural bounds on the effectiveness of hedging, which is determined by the convexity of the cost function and the elasticity of demand; 3) market competition reduces the volatility of spreads in commodity processing industries; 4) one-sided financial hedging, where a producer can hedge the costs of his input, but not the price of his output, may not be efficient in plausible scenarios, when supply shocks and demand shocks are of similar magnitudes, or in the presence of significant exogenous noise in the production function.

Although our message is relatively negative about the efficacy of one-sided hedging, it

[21] This theoretical literature started in the 1980s and 1990s with papers by Shapiro and Titman (1986), Smith and Stulz (1985) and Froot et al. (1993). Empirical studies include Tufano (1996) and Geczy et al. (1997).

does point to a number of ways that an informed dynamic hedging policy can increase its effectiveness. For example, to account for the observed variation in the crack spread, our model includes a substantial amount of exogenous noise in the cost of refining that we assume cannot be observed in advance by the hedger. In reality, some of these costs can be observed in advance, and the producer can hedge more effectively by accounting for them. In addition, we have assumed that producers cannot anticipate changes in the variance of future demand and supply shocks. In reality, changes in these variances can be directly observed, e.g., an increase in tensions in the Middle East implies an increase in the variance of future supply shocks, or they can be indirectly observed from information in the financial markets. Indeed, an important implication of our model is that the variance of the input price increases more than the variance of the output price when the supply shock variance increases, while the variance of the output price increases more when the variance of the demand shock increases. In theory, one can glean useful information about these variances from the commodity options markets.

We focused our attention on the petroleum refining industry because of the availability of data. However, our analysis applies to any industry that converts an input commodity into an output. In some sense, the refining industry is not ideal for our analysis of effectiveness of hedging inputs, since this is an industry where it is possible to at least partially hedge the outputs as well as the inputs. For this reason, in future research, it makes sense to study markets with outputs like broiler chickens and airline travel, which cannot be directly hedged.

The tight relation between its inputs and outputs also makes the petroleum refining industry somewhat unique. It might also be interesting to consider how our model applies to industries where these linkages are not as tight. For example, electricity is an important input for aluminum smelters, however, the link between electricity prices and aluminum prices do not co-move to the same extent as the prices of oil and gasoline for a couple of reasons. First, since aluminum production accounts for a very small part of the overall use of electricity, shocks to aluminum demand is likely to have a very small effect on electricity prices. Second, the price of aluminum is determined on world markets, since it is inexpensive to ship, while electricity is a local commodity, whose price is determined by regional supply and demand conditions that do not directly influence the price of aluminum. For these reasons, shocks to the supply of electricity have very little influence on aluminum prices, which is why aluminum producers tend to buy electricity with long term forward contracts.

A final extension of our model that we will leave to future research has to do with the production function. In our current model producers convert one unit of the input good into one unit of the output good at a cost that increases as capacity utilization increases. Alternatively, we can assume that the process converts one unit of the input to X units of the output, where X can vary from plant to plant (some are more or less efficient) and it can vary as capacity utilization increases and decreases. For example, older gas fired power plants have higher heat rates than newer plants, which means that they use more natural gas to produce the same amount of electricity. Similarly, big airplanes use less fuel per passenger than smaller airplanes.

For these industries, the optimal hedge ratio for the more or less efficient producer can be very different. We leave these extensions for future research.

Alexander, C., Prokopczuk, M., Sumawong, A., 2013. The (de) merits of minimum-variance hedging: Application to the crack spread. Energy Economics 36, 698–707.

Bessembinder, H., Lemmon, M.L., 2002. Equilibrium pricing and optimal hedging in electricity forward markets. Journal of Finance 57, 1347–1382.

Carmona, R., Durrleman, V., 2003. Pricing and hedging spread options. SIAM Review 45, 627–685.

Carter, D.A., Rogers, D.A., Simkins, B.J., Treanor, S.D., 2017. A review of the literature on commodity risk management. Journal of Commodity Markets 8, 1–17.

Casassus, J., Liu, P., Tang, K., 2012. Economic linkages, relative scarcity, and commodity futures returns. Review of Financial Studies 26, 1324–1362.

Cooper, J.C., 2003. Price elasticity of demand for crude oil: estimates for 23 countries. OPEC Energy Review 27, 1–8.

Dale, S., 2016. New economics of oil. Oil & Gas, Nat. Resources & Energy J. 1, 365.

Dong, L., Kouvelis, P., Wu, X., 2014. The value of operational flexibility in the presence of input and output price uncertainties with oil refining applications. Management Science 60, 2908–2926.

Froot, K.A., Scharfstein, D.S., Stein, J.C., 1993. Risk management: Coordinating corporate investment and financing policies. Journal of Finance 48, 1629–1658.

Geczy, C., Minton, B.A., Schrand, C., 1997. Why firms use currency derivatives. Journal of Finance 52, 1323–1354.

Hirshleifer, D., 1988a. Residual risk, trading costs, and commodity futures risk premia. Review of Financial Studies 1, 173–193.

Hirshleifer, D., 1988b. Risk, futures pricing, and the organization of production in commodity markets. Journal of Political Economy 96, 1206–1220.

Hirshleifer, D., 1989. Determinants of hedging and risk premia in commodity futures markets. Journal of Financial and Quantitative Analysis 24, 313–331.

Hughes, J., Knittel, C., Sperling, D., et al., 2008. Evidence of a shift in the short-run price elasticity of gasoline demand. Energy Journal 29, 113–134.

Hughes, J.E., Knittel, C.R., Sperling, D., 2006. Evidence of a shift in the short-run price elasticity of gasoline demand. Technical Report. National Bureau of Economic Research.

Ji, Q., Fan, Y., 2011. A dynamic hedging approach for refineries in multiproduct oil markets. Energy 36, 881–887.

Kamara, A., 1993. Production flexibility, stochastic separation, hedging, and futures prices. Review of Financial Studies 6, 935–957.

Kilian, L., 2009. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. American Economic Review 99, 1053–69.

Liu, P., Vedenov, D., Power, G.J., 2017. Is hedging the crack spread no longer all it's cracked up to be? Energy Economics 63, 31–40.

Liu, W., 2014. Modeling gasoline demand in the united states: A flexible semiparametric approach. Energy Economics 45, 244–253.

Secomandi, N., 2010. On the pricing of natural gas pipeline capacity. Manufacturing & Service Operations Management 12, 393–408.

Shapiro, A.C., Titman, S., 1986. An integrated approach to corporate risk management. The revolution in corporate finance 3, 251–265.

Smith, C.W., Stulz, R.M., 1985. The determinants of firms' hedging policies. Journal of Financial and Quantitative Analysis 20, 391–405.

Smith, J.E., Winkler, R.L., 2006. The optimizers curse: Skepticism and postdecision surprise in decision analysis. Management Science 52, 311–322.

Strebulaev, I.A., Whited, T.M., 2012. Dynamic models and structural estimation in corporate finance. Foundations and Trends in Finance 6, 1–163.

Tufano, P., 1996. Who manages risk? an empirical examination of risk management practices in the gold mining industry. Journal of Finance 51, 1097–1137.

Wirl, F., 2008. Why do oil prices jump (or fall)? Energy Policy 36, 1029–1043.

Wu, O.Q., Chen, H., 2010. Optimal control and equilibrium behavior of production-inventory systems. Management Science 56, 1362–1379.