JOURNAL OF FINANCIAL AND QUANTITATIVE ANALYSIS Vol. 58, No. 3, May 2023, pp. 1202–1229

doi:10.1017/S0022109022001478

Hedging Commodity Price Risk

Hamed Ghoddusi
California Polytechnic State University Orfalea College of Business
hghoddus@calpoly.edu

Sheridan Titman
University of Texas at Austin McCombs School of Business
Sheridan.Titman@mccombs.utexas.edu (corresponding author)

Stathis Tompaidis
University of Texas at Austin McCombs School of Business
Stathis.Tompaidis@mccombs.utexas.edu

Abstract

We present an equilibrium model of hedging for commodity processing firms. We show the optimal hedge ratio depends on the convexity of the firm's cost function and the elasticity of the supply of the input and the demand for the output. Our calibrated model suggests that hedging tends to be ineffective. When uncertainty comes exclusively from either the supply or from the demand side, updating the hedge dynamically, and using nonlinear contracts improves hedging effectiveness. However, with both supply and demand uncertainty, hedging effectiveness can be low even with option-based and dynamic hedging strategies.

I. Introduction

Commodity processing industries (e.g., refineries, power plants, metal smelters, food processors, fish farms, and even airlines) are ubiquitous and play a critical role in the modern economy. Commodity processors convert one form of a commodity (e.g., crude oil, jet fuel, or fish feedstock) into other products (e.g., gasoline, air travel, or salmon) and generate profits that are closely related to the spread between the price of the output and input goods. Given the high level of price uncertainty in these markets, it is natural for processors to use derivative markets to hedge their cash-flow risk. The evidence, however, on both hedging activities and hedging effectiveness is mixed. While some papers document the active use and positive effects of hedging (Bartram (2017), Gilje and Taillard (2017)),

We also thank participants and discussants for their input at seminars at HEC Montreal, Oxford Institute for Energy Studies, UTAustin, 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 Commodity Futures Trading Commission (CFTC), the 2011 Applied Financial Economics conference, the 2011 International Conference on Mathematical and Financial Economics, the 2011 CFTC Conference on Commodities Markets, the 2012 Workshop on Probability and Statistics in Finance, the 2012 Alberta Finance Institute Conference on Speculation, Risk Premiums, and Financing Conditions in Commodity Markets, "the 2015 workshop on Commodities and Financialization at the Institute for Pure and Applied Mathematics, the 2019 Oklahoma University Energy and Commodities Finance Research Conference, and the 2019 Supply Chain Finance and Risk Management Workshop."

others question the intensity and effectiveness of hedging (e.g., Mello and Parsons (2000), Guay and Kothari (2003), Jin and Jorion (2006), and Lievenbrück and Schmid (2014)).

We contribute to this debate by offering an equilibrium model of input and output prices that allows us to gauge the extent to which cash-flow risk can be reduced by hedging. Our model focuses on markets with two sources of uncertainty – shocks to the supply of the upstream commodity (e.g., jet fuel) and shocks to the demand for the downstream product (e.g., air travel), but in markets where firms only have the ability to hedge the input price. The model describes how characteristics of the input and output markets (e.g., supply and demand elasticities), the degree of competition, the properties of the production cost function (e.g., its convexity), and finally the relative magnitude of supply and demand shocks, influence input and output prices, the spread between the two, and optimal hedging policies.

The conventional view is that the firm's exposure to input cost risk is reduced if it locks in its input costs in advance. For example, a producer may hedge by buying inputs in forward markets. One key result of our model is that this conventional view holds only when supply shocks are the main source of uncertainty. When this is the case, input prices are more volatile than output prices, and changes in input prices are negatively correlated with changes in the spread between output and input prices. In contrast, if the main source of uncertainty comes from demand shocks, output prices are more volatile than input prices and changes in input prices are *positively* correlated with the changes in the spread. As we show, producers in this situation hedge perversely relative to the conventional view (i.e., they hedge by *selling* rather than buying the input forward).

In the general case with both demand and supply shocks, the correlation between the processor's profits, that is, the spread between output and input prices, and input prices can be low, or even 0, implying that hedging the price of the input may not materially reduce the variance of the producer's profits. As we show, depending on the magnitude of expected supply and demand shocks, the variance-minimizing hedge ratio can be either positive or negative. In addition to the variance of these shocks, the hedge ratio depends on the convexity of the cost function and the elasticities of the demand for the output and the supply of the input.

Our model also illustrates how the nature of competition affects the relationship between input and output prices, and as a result, the variance-minimizing hedge ratio. We show that while competition does not change the direction of the hedge, it does change its magnitude. In the case with only supply shocks, a monopolist that minimizes the variance of its profits purchases more units forward than a competitive firm. Similarly, when there are only demand shocks, the monopolist sells more of the input forward than a producer in a competitive market.

We illustrate these results analytically in a simple model where the relevant functions, that is, the marginal production cost and the supply and demand functions, are all linear. Despite its simple structure, the model delivers powerful insights on the direction of hedging policies, bounds on hedge ratios, and the overall effectiveness of hedging. However, since we are interested in the effectiveness of hedging, which is a quantitative issue, we also develop a more realistic model that allows for nonlinear demand and supply functions, and an arbitrary degree of production cost convexity.

The more realistic model is solved numerically, using parameters that roughly match quantities observed in the oil refinery industry. We consider a variety of scenarios and conclude that, under the most plausible scenarios, hedging reduces the variance of profits very little. We also consider whether nonlinear (e.g., option-based) and dynamic hedging strategies can improve hedging effectiveness, and find that option contracts improve hedging in some cases (i.e., when uncertainty is limited to either the supply of the input or the demand for the output). However, for the most part, hedging still tends to be ineffective in plausible situations. Similarly, we find that dynamic hedging (i.e., the frequent rebalancing of the hedge based on changing supply and demand) has the potential to improve hedging effectiveness, especially when the hedge is rebalanced frequently. But dynamic hedging remains ineffective with two sources of uncertainty.

Our analysis is most closely related to models developed by Hirshleifer (1988a), (1988b), and (1991) which explore the determinants of equilibrium spot and futures prices of a processed commodity. For a broad review of the commodity risk management literature, see Carter, Rogers, Simkins, and Treanor (2017). Hirshleifer (1988a) considers a case with only supply shocks and shows that the optimal hedge is long the input, and Hirshleifer (1988b) considers a case with only demand shocks and shows that the optimal hedge is long the output. Hirshleifer (1991) considers the case of different growers of an agricultural commodity, where depending on the time of year, uncertainty comes from either demand or supply shocks. Early in the season, there is a lot of output uncertainty, suggesting little hedging owing to the partial offsetting of price risk and quantity risk. Later in the season, after most of the quantity risk has been resolved, remaining demand uncertainty translates to price risk and an optimal hedge that is large in absolute value. These papers differ from ours in that they consider output risk as well as price risk. However, we extend this literature by jointly considering supply and demand shocks, and by exploring the role of the production cost function, the impact of market structure, and the effect of demand and supply elasticity.

Another related paper by Dybvig, Liang, and Marshall (2013) shows that hedging inputs with forward contracts can be improved with nonlinear hedging instruments such as options. We show that options improve the effectiveness of hedging in cases with just supply shocks, but when there is uncertain demand as well as uncertain supply, even option-based hedging tends to be ineffective. Kamara (1993) studies the optimal production and hedging decisions of the owner of a capital asset that can adjust its production plans and shows that production decisions can be independent of risk preferences in the presence of futures contracts. Papers such as Ho (1984) and Brown and Toft (2002) are concerned with hedging decisions of firms that face both price and quantity risks, and derive optimal hedging policies using various forms of derivative contracts. Fehle and Tsyplakov (2005) consider a dynamic hedging framework in continuous time to study the optimal use of short-term contracts for hedging long-term risks. These models assume that input and output prices are exogenous, so hedging effectiveness is also effectively exogenous.

Bessembinder and Lemmon (2002) develop 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 price of the output good and on how these prices are influenced by expected demand and uncertainty about shocks to demand. Bessembinder and Lemmon (2002) do not consider the feedback between a shock to electricity demand and the price of the input (i.e., they assume that increased demand for electricity does not affect the price of natural gas, which is a key feature of our model). Bessembinder and Lemmon (2002) show that the skewness in output prices is driven by the convexity of the cost function but do not study the correlation between input and output prices and their spread, or hedging and its effectiveness. Finally, Routledge, Spatt, and Seppi (1998) and Casassus, Liu, and Tang (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. To keep our model simple, we ignore storage, which is very important for understanding daily fluctuations of input and output prices, especially when demand or supply is seasonal, but may be less important for understanding the longer-term dynamics that is our focus.

II. Motivation

The conventional wisdom on hedging commodity price risk can be illustrated with an anecdote from the first book of Politics, Part XI, written by Aristotle in the fourth century BCE. Aristotle describes how Thales of Miletus secured the use of all the local olive presses based on his forecast of plentiful olive production on the island of Chios. The subsequent realization of a good harvest resulted in a reduction in olive prices and an increased demand for the limited number of olive presses, allowing Thales to rent out his olive presses for a large profit. This example illustrates how a positive shock to the supply of the input commodity decreases the price of the input and increases the spread between the price of the output and the price of the input when there are constraints on the capacity of the processors.

Moving to the present, the oil refinery industry provides a similar setting for evaluating the relationship between the price of commodity inputs and the profits of firms that use the inputs to produce consumer goods. Since crude oil and refined products are actively bought and sold on fairly liquid markets, input and output prices as well as their spreads are observable, which allows us to calculate correlations between crude oil prices and what is known as the crack spread (i.e., the spread between the price of the refined products that are produced by the refinery and the price of crude oil).

As illustrated in Figure 1, a scatter plot of changes in crude oil prices versus changes in crack spreads, there is, at best, a weak relationship between the crack spread and oil prices. Indeed, the correlation calculated using monthly data over the entire 1989–2019 sample period is close to 0. However, Table 1 and Figure 2 show that the zero correlation over the entire sample period masks the fact that the correlation in the initial 5 year period is negative (and statistically significant), and in each of the subsequent 5 year periods the correlations are relatively low and are not statistically significant. We will show that this pattern is consistent with a setting where price movements are mainly generated by supply shocks in the initial period, with the volatility of demand shocks becoming more important in later periods.

FIGURE 1

Crude Oil Price Versus Crack Spread

Figure 1 shows the scatter plot of monthly changes in the prices of crude oil and crack spreads between 1989 and 2019.

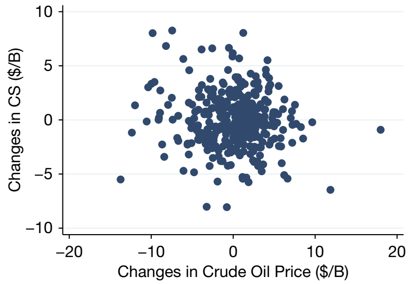

TABLE 1

Crude Oil Price Versus Crack Spread

Table 1 shows the correlation of crude oil prices and crack spreads over the 1989–2019 period. ** indicates that the correlation is significant at the 5% level.

Period	No. of Obs.	Correlation
1989/05–1994/04	60	–0.26**
1994/05–1999/04	60	0.11
1999/05–2004/04	60	0.04
2004/05–2009/04	60	0.20
2009/05–2014/04	60	0.00
2014/05–2019/04	60	–0.03

FIGURE 2

Correlation Between Crack Spread and Crude Oil Price

Figure 2 shows the correlation between crack spread and crude oil prices. The correlations are calculated using monthly data over 60-month, nonoverlapping, windows. Correlations calculated between Jan. 1989 and Dec. 1993 are shown under the label for 1989; correlations calculated between Jan. 1994 and Dec. 1998 under the label for 1994; and so forth.

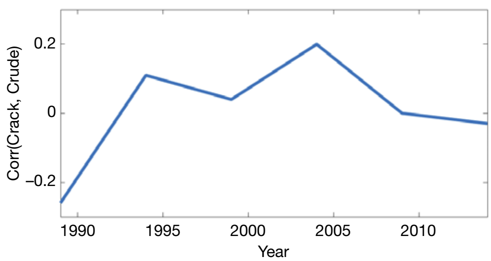

Figure 3 shows the volatility of crude oil prices relative to the volatility of the price of refined products for various subperiods. We will show that this ratio provides information about the extent to which price changes are due to supply versus demand shocks. In particular, consistent with the results in Figures 1 and 3, we will show that when crude oil prices are more volatile than the prices of refined products, the correlation between the prices of crude oil and refined products is negative. The ratio flips when the correlation is positive.

FIGURE 3

Ratio of Volatility of Crude Oil Price over Volatility of Gasoline Price

Figure 3 shows the relative volatility of crude oil prices to gasoline prices. We use a 60-month window and report the ratio of the standard deviation of crude oil and gasoline monthly returns within each window.

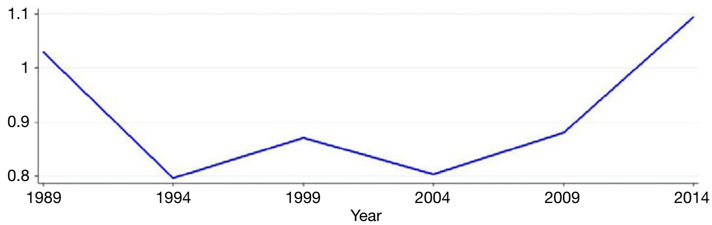

III. Model

We consider a competitive industry that uses a capital asset to convert one unit of input to one unit of output. We assume that the firms in this industry are identical price takers for both the input and output commodities.

The price of a unit of the input is determined by a linear inverse supply function

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

where X_s is the value of an exogenous, stochastic, supply factor; Q is the supply of the input; and γ_s is the inverse elasticity of supply (i.e., the inverse change in quantity supplied for a unit change in price). Since we assume that one unit of input is converted into one unit of output, the variable Q represents both the quantity of the input and the quantity of the output.

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 the quantity produced:

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

The demand factor X_d is exogenous, while γ_d is the inverse elasticity of supply. We assume that demand and supply shocks are independent of each other.[1]

The representative firm's production function exhibits an increasing marginal cost schedule. The increasing marginal cost can be viewed 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

[1]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 (e.g., improvements in the fuel efficiency of the stock of existing automobiles). They can also correspond to short-term shocks to income and consumer tastes. Given this intuition, it is reasonable to expect that supply and demand shocks are uncorrelated, or that their covariance is small enough that it can be safely ignored.

TABLE 2

Model Notations

Table 2 presents the list of variables, parameters, and functions used in the 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	Inverse of demand elasticity	Constant
γ_s	Inverse of supply elasticity	Constant
$\mathrm{TC}(Q)$	Total cost of producing Q units	Convex in the production quantity
$\phi(Q)$	Capacity-related costs	Convex in capacity utilization
$P_{I,g}$	Unit cost of other inputs	Constant
F_C	Fixed cost of processing	Constant

linearly with the amount produced, that is, the total cost to produce Q units, $\mathrm{TC}(Q)$, is given by

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

The total cost function includes a fixed cost component, F_C, and a variable component, QP_c, that represents the cost of purchasing Q units of the input. The cost of variable inputs other than the main input commodity, such as energy, labor, maintenance, and so forth, is given by $P_{I,g}Q$. While, in reality, the prices for these inputs may be random, for simplicity we assume that they are certain.[2] Finally, the cost function contains an increasing, quadratic, component $\phi(Q) = \lambda Q^2/2$, which represents the activation of higher marginal cost units.

Table 2 summarizes our notation.

A. Equilibrium

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

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

In Appendix A, we solve for the equilibrium quantity produced, as well as the price of the input and output, their spread, and their variances and covariances. Each of these quantities depends on demand and supply elasticities, the coefficient of convexity of the production cost function, and the magnitudes of supply and demand shocks.

[2]Assuming that the price of other inputs, such as natural gas, is random creates an additional source of volatility in spreads and further complicates the hedging decision.

B. Properties of Prices and Spreads

We present two propositions that characterize the relationships between input prices, output prices, and their spread (proofs are provided in Appendix A). The first proposition provides properties of the input and output prices relative to the magnitudes of supply and demand shocks.

Proposition 1. In a competitive market described by equations (1)–(3), the covariance of input and output prices with their spread depends on the extent to which uncertainty is generated by supply versus demand shocks.

- If supply uncertainty dominates, then the covariance of the price of either the input or the output and their spread is negative.
- If demand uncertainty dominates, then both covariances are positive.
- 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.
- With both supply and demand shocks, the value of the ratio is bounded by

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

Proposition 1 shows that the source of uncertainty has an impact on the magnitude of the variances of the input and output prices. This result can help identify the source of uncertainty when market prices for option contracts are available. If the implied values for the variance of input and output prices are available, we can identify whether expected shocks are likely to come from the supply of the input or the demand for the output: An increase (decrease) in the value of the ratio can be attributed to an increase in the variance of supply (demand) shocks.

The second proposition describes the sensitivity of the spread between the price of the output and the price of the input versus the elasticities of demand and supply, as well as the coefficient of convexity of the cost function.

Proposition 2. In a competitive market described by equations (1)–(3), the variance of the spread between the output price and the input price decreases with the inverse supply elasticity, γ_s, the inverse demand elasticity, γ_d, and increases with the coefficient of convexity, λ, that is,

$$\frac{\partial var(P_g - P_c)}{\partial \gamma_s} < 0$$
$$\frac{\partial var(P_g - P_c)}{\partial \gamma_d} < 0$$
$$\frac{\partial var(P_g - P_c)}{\partial \lambda} > 0.$$

IV. Hedging

The literature provides several explanations for why firms hedge their profits. For example, Smith and Stulz (1985) and Shapiro and Titman (1986) consider bankruptcy costs and financial distress; Froot, Scharfstein, and Stein (1993) consider costly external financing; Graham and Smith (1999) consider a convex tax structure; and DeMarzo and Duffie (1995) and Breeden and Viswanathan (2015) consider information asymmetry. Rather than focus on the underlying motivation for hedging, we instead focus on the effectiveness of hedging, measured by the reduction in the variance of the firm's profits. We assume that, for hedging purposes, the firm has access to a single financial contract, a forward contract on the price of the input.

We assume that the representative firm seeks to minimize the variance of profit per unit produced (i.e., the variance of the spread between the price of the output and the price of the input, $P_g - P_c$, rather than the variance of total profit). Hedging profit-per-unit is a reasonable approximation to hedging total profit in cases where the *changes* in the quantity of the output are small relative to the level of production. Two examples where this happens are when the price elasticity of demand is sufficiently small and when the production costs are sufficiently convex. The gasoline market and refineries fit these examples: the price elasticity of demand is small, while the level of production for refineries ranges between 79% and 87% of the total capacity. The range in production, relative to average capacity utilization is approximately 10% while the inflation-adjusted spreads range between \$5/barrel to \$16/barrel; corresponding to a range of up to 100% relative to the average spread, an order of magnitude larger than the range in the quantity of production.

Given this approximation, the producer minimizes the variance of its profit per unit produced by selling β units of the input forward, that is, minimizes the residuals in the relation

$$(5) \qquad \left(P_{g,t} - P_{c,t}\right) = \alpha + \beta(P_{c,t} - F_{c,0}) + \tilde{\varepsilon}_t,$$

where $F_{c,0}$ is the forward price at time 0 for a unit of input delivered at time t.[3]

The optimal hedge ratio, β, is given by the coefficient from the regression of P_c on $P_g - P_c$[4]:

$$(6) \qquad \text{Hedge Ratio} = -\frac{\text{cov}(P_g - P_c, P_c)}{\text{var}(P_c)}.$$

A. Properties of the Variance-Minimizing Hedge Ratio

Equation (6) indicates that the variance-minimizing hedge ratio depends on the variance of the spread, the variance of the price of the input, and the correlation

[3]In our analysis, we do not consider the, potential, cost of hedging. In particular, since our objective is to minimize the variance of profits, risk premiums in forward markets are not relevant in our analysis. An objective that takes into account both variance of profits, as well as expected profits, would need to consider the existence of risk premiums, and whether hedging would give the risk premium up.

[4]We note that since the forward price, $F_{c,0}$ is known, it does not influence the estimate of the regression coefficient, β.

between the spread and the price of the input. Proposition 1 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 forward. On the other hand, if uncertainty is driven by demand shocks, the producer should sell the input commodity forward.

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 optimal hedge ratio and the effectiveness of hedging.

Proposition 3. In a competitive market described by equations (1)–(3), the optimal hedge ratio is given by

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

- The optimal hedge ratio is bounded between the hedge ratio when there is no uncertainty in demand, $\lambda/(\lambda+\gamma_d)$, and the hedge ratio when there is no uncertainty in supply,$-\lambda/\gamma_s$

$$-\frac{\lambda}{\gamma_s} \leq \text{Hedge Ratio} \leq \frac{\lambda}{(\lambda+\gamma_d)}$$

- When demand is certain, but supply is uncertain, the hedge ratio that minimizes the variance of profits is positive (i.e., long futures contracts). The hedge ratio is negative when supply is certain and demand is uncertain.
- 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.
- When demand is certain and supply is uncertain, an increase in the inverse elasticity of the demand coefficient, γ_d, decreases the optimal amount bought; when only demand is uncertain, an increase in the inverse elasticity of supply coefficient, γ_s, decreases the optimal amount sold.

Proposition 3 has several implications. First, with only supply uncertainty, the hedge ratio is positive and between 0 and 1 – notably, it is always less than 100% when demand is not perfectly elastic since shocks to supply lead to a change in the input price that is greater than the change in the output price. In this case, the hedge ratio is an increasing function of both the convexity coefficient, λ, and the elasticity of demand, $1/\gamma_s$. Intuitively, when the cost function is more convex and demand is more elastic a smaller portion of an input price increase can be passed along to the consumer of the output good.

The importance of the convexity coefficient is easy to understand if one considers the extreme case where the processing cost is constant (i.e., $\lambda=0$). In this case, the spread between the input and output prices is constant regardless of the demand elasticity, so the hedge ratio is 0. When the convexity coefficient is positive, $\lambda>0$, the elasticity of demand is also important. In this case, an increase in the price of the input is not fully passed on to the consumer of the output, because the quantity demanded decreases when the output price increases, and, because of this decrease in quantity demanded, the marginal cost of each firm's production declines. In this case, the spread between the output and input prices declines by an amount determined by the drop in the quantity demanded, which is, in turn, determined by the elasticity of the demand function.

When only demand is uncertain, the hedge ratio is negative, and is an increasing function of the elasticity of supply, $1/\gamma_s$. The intuition is that an increase in the demand for the output good increases the profit of the producer and at the same time increases the price of the input good. The hedge ratio in this case is determined by how the gains associated with an increase in demand are shared by the output producers and the input suppliers. In the extreme case where the supply is almost perfectly elastic, almost all of the gains are captured by the output producers, so a small increase in the input price is associated with a very large increase in the spread between the input and output prices. In contrast, when the supply of the input is inelastic, the input price increases substantially when demand for the output good increases, so the producers of the output capture less of the gain. Regardless of the supply elasticity, the hedge ratio is again 0 when the coefficient of convexity of the production function is 0, $\lambda = 0$, since competition forces the producers to sell the output at a constant spread over the input price.

With both supply and demand shocks, the hedge ratio is 0 when the variance of supply and demand shocks balance, that is, when

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

B. Hedging Effectiveness

Using our framework, we can quantify the effectiveness of hedging profits using forward contracts on the input. We define hedging effectiveness as the reduction in the variance of the profits by hedging, which is the coefficient of determination, R^2, in the regression of the spread and the profit on the forward contract.

Proposition 4. In a competitive market described by equations (1)–(3), hedging effectiveness depends on the magnitude of supply and demand shocks in a non-monotone way.

- With certain demand (supply) and uncertain supply (demand), hedging effectiveness is 100%.
- Hedging effectiveness decreases as uncertainty in demand (supply) increases, until the hedge ratio becomes 0 (in that case, hedging effectiveness is 0).
- As the uncertainty in demand (supply) increases further, hedging effectiveness increases – in the limit where supply (demand) is certain and demand (supply) uncertain, it reaches 100% again.

Proposition 4 shows that hedging effectiveness is highest when there is single source of uncertainty, either supply or demand. Another implication of Proposition 4 is that, with two sources of uncertainty, higher supply or demand uncertainty do not necessarily result in larger hedging positions; firms may indeed optimally *reduce* their hedge in response to greater uncertainty.

V. Market Power

Up to this point, we have assumed that firms are price takers. In this section, we consider firms that account for the effect of their production on output prices – the firm may be a monopolist or an oligopolist. We formulate the problem for the case of a single firm – the case with several firms is similar.

Rather than take prices as given, a firm with market power in the output market maximizes its total profit $\pi = P_g(Q)Q - \mathrm{TC}(Q)$ by considering the effect of its production decisions on prices:

$$\max_Q Q(X_d - \gamma_d Q) - Q\left(P_{I,g} + X_s + \gamma_s Q + \lambda Q\right). \tag{7}$$

The optimal amount produced, as well as the prices of the input, output, and their spread are provided in Appendix B.

The following propositions summarize the difference between a competitive market and a market where firms have market power.[5]

Proposition 5. In a market described by equations (1)–(3), the variance of the quantity produced by a firm with market power, as well as the variance of the spread between the output price and the input price, is smaller than the variance of the quantity produced in a competitive market.

Proposition 5 indicates that the production choices of a firm with market power, compared to firms in a competitive market, are less sensitive to both demand and supply shocks. Our next proposition compares the hedging behavior of a firm with market power to that of a firm in a competitive market.

Proposition 6. In a market described by equations (1)–(3), when there are only supply shocks, the hedge ratio that minimizes the variance of profits is greater for a firm with market power than for a firm in a competitive market (i.e., the firm with market power buys more of the input forward). On the other hand, when there are only demand shocks, it sells more of the input forward than a firm in a competitive market.

In the case with only supply shocks, the hedge ratio for the firm with market power is greater than the hedge ratio of the competitive firm (the firm with market power hedges by buying more forward contracts). In the case with only demand shocks the firm with market power hedges by selling more forward contracts than the competitive firm. However, because the optimal hedge ratios for the firm with market power and the competitive firm become 0 at different combinations of the variance of supply and demand shocks, one cannot make general statements about differences in these hedge ratios. Overall, our results are qualitatively the same under both market structures (the correlation between prices and spreads, as well as the hedge ratio, depend on the relative magnitude of supply and demand shocks).

[5]Similar results can be derived for firms that have monopsony power in the input market.

https://doi.org/10.1017/S0022109022001478 Published online by Cambridge University Press

VI. A Calibrated Model

Up until now, we have considered the relationship between input prices and firm profits within a stylized linear model that allows us to derive easy-to-interpret expressions. However, since one of our goals is to gauge the effectiveness of hedging, which is a quantitative exercise, we also present a model with more realistic supply, demand, and production functions. While we believe that our analysis applies to any firm that converts an input commodity into an output, for clarity, we roughly calibrate the model to match a refinery that converts crude oil into refined products like gasoline. In contrast to our linear model, this more realistic model must be solved numerically, that is, we must use numerical methods to find the equilibrium production level and corresponding input and output prices, and their spread.

A. Model Specification

Our stylized linear model focuses on hedging effectiveness over a single period, and as a result, dynamics are not important; we only care about the variances of supply and demand shocks over the next period. However, when we calibrate a more realistic model using actual data, we need to account for dynamic aspects of the supply and demand shocks. In the case of the crude oil to refined products supply chain, we expect that both supply and demand shocks exhibit mean-reversion due to, for example, increased exploration when crude oil prices are high for the case of supply, or the nature of the business cycle for the case of demand. We capture this behavior in our model with mean-reverting, stochastic, processes for supply and demand shocks.

1. Supply and Demand Functions

Based on the crude oil production of low-cost members of the Organization of Petroleum 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. Specifically, we approximate the inverse supply function by a function that does not depend on quantity below a production threshold, and that, beyond that threshold, exhibits constant elasticity. Supply shocks shift the entire function up or down. The inverse supply function is given by

$$P_C = e^{X_s} + ((Q - \underline{Q})\mathbf{1}_{Q>\underline{Q}})^{\gamma_s}. \tag{8}$$

Consistent with the demand elasticity literature (e.g., Liu (2014)), we use a constant-elasticity function for the demand of refined products[6]

$$P_G = e^{X_d} Q^{\gamma_d}. \tag{9}$$

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

The factor X_s captures exogenous shocks to the supply of crude oil, while X_d captures exogenous shocks to demand; $\underline{Q}$ is the threshold where the steep part of the inverse supply function begins; γ_s is the sensitivity of the marginal cost of crude oil to the supply of crude oil at levels of output above the threshold, while γ_d is the elasticity of demand.

2. Stochastic Processes for Supply and Demand Factors

To compare model quantities to observed quantities, we assume that the stochastic supply and demand factors follow discrete-time AR(1) mean-reverting dynamics.

$$\text{(10)}\qquad \begin{aligned}\Delta X_s &= \mu_s(\overline{X}_s - X_s(t))\Delta t + \sigma_s \varepsilon_s(t)\sqrt{\Delta t},\\ \Delta X_d &= \mu_d(\overline{X}_d - X_d(t))\Delta t + \sigma_d \varepsilon_d(t)\sqrt{\Delta t},\end{aligned}$$

where we assume that the mean reversion rates μ_i, the long-term levels $\overline{X}_i$, and the volatilities σ_i, $i \in \{s,d\}$, are constant. The values of the stochastic supply and demand factors at time t, are given by $X_s(t)$, and $X_d(t)$, respectively. The time step is Δt, and the shocks to supply and demand, $\varepsilon_s(t), \varepsilon_d(t)$, are assumed to be IID, normally distributed, with mean zero and standard deviation 1. We also assume that the correlation between random shocks, ε_s and ε_d, is 0.[7]

Observing the realization of state variables X_s and X_d, the firm tries to minimize the conditional variance of its profit in the next period.

3. Cost Function

The marginal cost of refining Q units of crude oil is determined by the following expression:

$$\text{(11)}\qquad \mathrm{MC}(Q) = P_{I,g} + P_C + \phi(Q) = P_{I,g} + P_C + \lambda Q^{\eta},$$

where $P_{I,g}$ is the cost of other inputs, P_C is the price of crude oil, and $\phi(Q)$ are costs that depend on the level of production.

4. Shocks to Processing Costs

We assume that marginal costs increase as refineries approach capacity, and these capacity-related costs, $\phi(Q) = \lambda Q^{\eta}$, are given by a power function, with two parameters λ and η. We assume that the convexity of the cost function, described by the coefficient η, is constant. To account for fluctuations in the production cost function over the period that we study due to, for example, randomness in the cost of other inputs to the refining process (e.g., natural gas), we allow the coefficient λ to be a random variable drawn from the following distribution:

$$\text{(12)}\qquad \lambda_t = \overline{\lambda} e^{\sigma_\lambda \varepsilon_t},$$

[7] We note that assuming independent shocks to demand and supply does not mean independent input and output prices.

https://doi.org/10.1017/S0022109022001478 Published online by Cambridge University Press

where $\bar{\lambda}$ is the baseline value, σ_λ is the standard deviation of shocks to the coefficient, and ε_t are standard, normally distributed, IID random variables.[8]

We note that the fluctuations in the cost function have significant empirical consequences: For example, they influence the volatility of the crack spread. Given the IID nature of the random variables, ε_t, we expect that changes in the marginal production costs will be negatively autocorrelated.

B. Calibration

To calibrate the model parameters we match the moments of prices and spreads for the crude oil to gasoline supply chain. The prices we examine include oil, the output of a refinery which we define as the weighted basket of two parts gasoline and one part heating oil (the two most important refinery outputs) and the crack spread (i.e., the difference in the price of this basket and the price of the crude oil input). These prices, as well as the quantities of crude oil, refined products, and refining capacity are observed annually from 1987 to 2018.[9] We obtain prices for crude oil, gasoline, and heating oil from the Energy Information Administration (EIA) (https://www.eia.gov/dnav/pet/pet_pri_spt_s1_d.htm), and data for the total amount of crude oil refined and the global refining capacity from the statistical review report issued by British Petroleum (https://www.bp.com/en/global/corporate/energy-economics/statistical-review-of-world-energy.html).

1. Parameter Estimates

We use simulated method of moments (SMM) to estimate the structural parameters of the model (Strebulaev and Whited (2012)). 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.[10] 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.

The empirical moments used in the SMM procedure are the following: i) the mean and standard deviation of the price of refined products; ii) the mean and standard deviation of crude oil prices; iii) the annual autocorrelation of the price of refined products; iv) the correlations between the price of refined products and the crack spread; v) the price of refined products and crude oil, and the price of crude oil and the crack spread; vi) the mean and standard deviation of capacity utilization; vii) the mean of the crack spread; and viii) the variance of the crack spread explained by forward crude oil prices. Overall, we use 12 empirical moments.

The model parameters we estimate are the following: the demand and supply elasticities; the mean-reversion rates for supply and demand shocks; the standard

[8]We note that ignoring this source of randomness makes matching moments generated from the model to moments in the data very difficult. For example, without this source of randomness, it is difficult to match the standard deviation of the prices of oil, refined products, and spreads, simultaneously.

[9]Our choice to use annual data somewhat mitigates the need to model storage because storage from year to year is less common, potentially due to the limited amount of storage available.

[10]We use 100 scenarios. For each scenario, we set the number of time periods of the simulated vector, T, to 200 (each period corresponds to 1 year). We discard the first 100 periods in each scenario and use the next 100 periods to estimate the various moments.

https://doi.org/10.1017/S0022109022001478 Published online by Cambridge University Press

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 cost function; and the standard deviation of fluctuations to the output from the cost function. The total number of estimated parameters is 11, implying that our model is over-identified.

2. Moment Response

In order to identify model parameters, the response of moments to changes in parameters should be strong, smooth, and well-behaved. In particular, at least some model-generated moments should be sufficiently sensitive to the value of the parameter; otherwise, the parameter cannot be identified. Moreover, it is desirable to have a clear local maximum for the distance between model-generated and empirical moments. We plot the behavior of all moments with respect to each of the parameters and make sure the conditions for the SMM identification are satisfied (e.g., there is a smooth response).[11]

3. Starting Values

We estimate plausible starting values (through the literature or direct approximate estimation) for every parameter.

Elasticity of Demand. The literature reports a range of values for short-term price elasticity of demand between 0.00 to -0.15 for different countries – see Cooper (2003), and Hughes, Knittel, and Sperling (2008). We choose a starting value of γ (the inverse elasticity) equal to -20.

Supply Elasticity. We use a starting value for supply elasticity of $\gamma_s = 1.45$, which is equivalent to assuming that a \$10 per barrel increase in price results in the production of an additional million barrels of crude oil per day.

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.

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 (see https://iea.blob.core.windows.net/assets/cbf37dfc-5fe1-4854-b248-95e6a2e5240a/Refining_Margin_Supplement_OMRAUG_12SEP2012.pdf). Going forward, we report the crack spread net of this processing cost.

Convexity of the Cost Function. Figure 4 suggests that the relationship between the deflated crack spread and global capacity utilization is close to linear. Thus, 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

[11]The plots are available from the authors.

FIGURE 4

Crack Spread Versus Capacity Utilization

Figure 4 shows the crack spread versus 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.

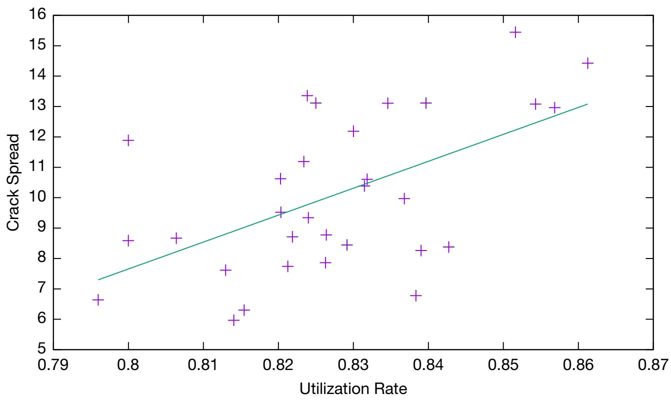

TABLE 3

Model Parameters

Table 3 presents the list of parameters and their values.

Notation	Parameter	Value
γ_d	Inverse of demand elasticity	−5.0
γ_s	Quantity sensitivity of crude oil price	11.2
μ_d	Mean-reversion rate for demand (year^{-1})	0.11
μ_s	Mean-reversion rate for supply (year^{-1})	0.05
$P_{I,g}$	Cost of other inputs	3 ($/b)
$\underline{Q}$	Capacity of crude oil supply	90 (mb/d)
η	Convexity of marginal cost function utilization	1.85
$\bar{\lambda}$	Coefficient of marginal cost function	0.0022
σ_λ	Shocks to capacity	0.07
$\overline{X}_s$	Mean of supply factor	3.82
$\overline{X}_d$	Mean of demand factor	25.47
σ_d	Standard deviation of demand factor	0.53
σ_s	Standard deviation of supply factor	0.14

linear relationship between crack spreads and capacity utilization, which provides a starting value equal to 3.2.

4. Results

Table 3 reports the parameter values from the SMM procedure. Table 4 reports the empirically observed values of the moments as well as the values generated by the calibrated model.

We find 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 correlation between the prices of crude oil, refined products, and the crack spread, are also accurate in terms of the sign and are reasonably close in terms of size.

https://doi.org/10.1017/S0022109022001478 Published online by Cambridge University Press

TABLE 4
SMM Moments

Table 4 shows the moments used in the simulated method of moments estimation.

Moment	Empirically Observed Value	Model-Generated Value	Weight
Mean refined price	64.27	64.85	20
Std. Dev. of refined price	31.76	31.23	10
Mean crude oil price	54.12	54.62	20
Std. Dev. of crude oil price	31.30	30.93	10
Mean of crack spread	7.12	7.23	20
Explained variance of crack spread	0.10	0.10	20
Average capacity utilization	82%	80%	10
Std. Dev. of capacity utilization	1.7%	2.4%	5
Autocorrelation of annual refined price	0.86	0.86	5
Correlation of crack spread and refined price	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

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. This mismatch could be due to a capacity mismeasurement problem. It can also be due to the fact that we do not incorporate shocks to nonoil input costs, like natural gas. Since we do not account for these shocks the method increases the volatility of capacity utilization to match the volatility of crack spreads with the empirical value.

VII. Numerical Results for Hedging Effectiveness

In this section, we use numerical simulations to evaluate the effectiveness of hedging, which we define as the reduction in the variance of profits per unit of production achieved using financial instruments. Given the calibrated parameter values, we randomly draw 5,000 realizations of the annual changes in the demand and supply factors (one value of the demand factor and one value of the supply factor per draw). For each draw, we numerically solve the model for the equilibrium quantity of gasoline produced, and the corresponding prices of crude oil and gasoline and their spread. We use the prices across the 5,000 draws to calculate the covariance of profits with the price of the input, as well as the variance of the price of the input. These quantities are then used to estimate the variance-minimizing hedge ratio, and calculate the ratio of the variance of the hedged profits to the variance of the unhedged profits (i.e., hedging effectiveness). For our base case we perform this process with initial values for the demand and supply factors set to their long-term means, and then repeat the process with alternative parameters which we vary to calculate comparative statics. Specifically, we are interested in how the exogenous parameters affect hedging effectiveness.

A. Comparative Statics

Our comparative statics examine how hedging effectiveness is influenced by the standard deviation of supply and demand shocks, and the convexity of the production function. We first evaluate the effectiveness of hedging with just forward contracts on the input commodity, and then consider a more complicated

TABLE 5
Hedging Effectiveness

Table 5 presents the hedging effectiveness as a function of the annualized standard deviation of supply and demand shocks for a forward-based hedging strategy.

Std. Dev. of Demand	Std. Dev. of Supply						
	0.00	0.10	0.20	0.30	0.40	0.50	0.60
0.00	–	0.99	0.99	0.98	0.96	0.94	0.91
0.10	0.27	0.00	0.39	0.67	0.79	0.84	0.85
0.20	0.23	0.08	0.01	0.20	0.43	0.58	0.67
0.30	0.21	0.14	0.02	0.01	0.13	0.30	0.44
0.40	0.20	0.16	0.07	0.00	0.01	0.10	0.22
0.50	0.18	0.16	0.11	0.04	0.00	0.01	0.08
0.60	0.17	0.16	0.12	0.07	0.02	0.00	0.01

TABLE 6
Hedge Ratio

Table 6 presents the hedge ratio, expressed as percentage of production.

Std. Dev. of Demand	Std. Dev. of Supply						
	0.00	0.10	0.20	0.30	0.40	0.50	0.60
0.00	–	1.6	1.6	1.6	1.5	1.4	1.3
0.10	−2.3	−1.6	−0.4	0.4	0.8	1.0	1.1
0.20	−2.3	−2.0	−1.5	−0.9	−0.2	0.3	0.6
0.30	−2.2	−2.1	−1.9	−1.5	−0.9	−0.4	0.1
0.40	−2.2	−2.2	−2.0	−1.7	−1.4	−0.9	−0.4
0.50	−2.2	−2.2	−2.1	−1.9	−1.6	−1.3	−0.8
0.60	−2.2	−2.2	−2.1	−2.0	−1.8	−1.5	−1.1

hedge that allows for both forward contracts and option contracts. The more complicated hedge includes one in-the-money call option, one out-of-the-money put option, and one forward contract on the input commodity. The strike prices of options are 1-standard-deviation above the forward price.[12]

1. Standard Deviation of Supply and Demand

For the hedging strategy using forward contracts, Tables 5 and 6 present the hedging effectiveness and hedge ratio for different values of the standard deviations of the supply and demand shocks. The results are in line with the intuition developed in the linear model and Proposition 4. The hedge ratio is positive when most uncertainty comes from input supply shocks, and is negative when most uncertainty comes from shocks to the demand for the output.

Table 5 illustrates that hedging effectiveness is not monotonic with respect to the standard deviation of either supply or demand shocks. Similar to the linear case, when the standard deviation of supply shocks is high, the correlation between input price and spreads is large and negative, and hedging effectiveness is high. As the standard deviation of supply decreases, this correlation drops. At some point –

[12]We did not use the forward price as the strike price for the options to avoid a collinearity problem. In results we do not report, we experimented with options with different strike prices and observed that they had little impact on the overall hedging effectiveness.

TABLE 7

Hedging Effectiveness with Options

Table 7 presents the hedging effectiveness of option-based strategies against the annualized standard deviation of supply and demand shocks.

Std. Dev. of Demand	Std. Dev. of Supply						
	0.00	0.10	0.20	0.30	0.40	0.50	0.60
0.00	–	0.99	0.99	0.99	0.99	0.99	0.98
0.10	0.49	0.00	0.46	0.72	0.83	0.88	0.91
0.20	0.46	0.08	0.01	0.23	0.47	0.61	0.70
0.30	0.45	0.17	0.03	0.01	0.15	0.32	0.46
0.40	0.46	0.25	0.08	0.01	0.01	0.10	0.24
0.50	0.46	0.29	0.13	0.04	0.01	0.01	0.08
0.60	0.46	0.31	0.19	0.08	0.03	0.00	0.01

which is a function of the standard deviation of supply and demand shocks as well as other structural parameters – the correlation of the spread and the price of the input, the hedge ratio, and hedging effectiveness becomes 0. As the standard deviation of supply decreases further, the correlation becomes negative, and hedging effectiveness increases again. The results in Table 5 also suggest that nonlinearities in the marginal cost function and the supply and demand functions, result in low hedging effectiveness even with a single source of uncertainty. The table also indicates that the relationship between hedging effectiveness and supply and demand standard deviations is complicated, with potentially multiple inflection points.

Table 5 also reveals that hedging with forward contracts on the input is more efficient when the refiner faces only supply shocks relative to the case when the refiner faces only demand shocks. This behavior can be understood by considering the limiting case of a production function with a capacity limit. Once production reaches the capacity limit, the price of the output decouples from the price of the input and is determined by the demand factor alone. In a market dominated by supply shocks, the output price is relatively stable, while the input price varies. In this situation hedging by buying the input forward is very efficient. Conversely, when demand shocks dominate, the price of the input is relatively stable, while the price of the output varies. While this result is no longer true without a capacity limit, the example illustrates that hedging supply shocks with contracts on the price of the input is generally more effective than hedging demand shocks. The results in Table 5 illustrate this intuition.

Table 6 reports optimal hedge ratios (i.e., positions in the forward contract) for different standard deviations of demand and supply shocks. The hedge ratios, reported in Table 6, are quite small – they vary between being long less than 2% to being short slightly more than 2% of the amount produced. This is a consequence of the assumed small convexity of the production function and the relatively inelastic demand for gasoline, which in combination imply that refineries pass almost all of the increase in oil prices to end users, resulting in hedge ratios that are very small.

Table 7 presents the hedging effectiveness of the options-based hedging strategy. We note that options improve performance across the board (an indication that the relationship between the price of the input and the spread between the price of the output and the price of the input is nonlinear). The improvement is greatest

FIGURE 5

Hedging Effectiveness

Figure 5 shows the Hedging effectiveness of forward-based and option-based hedging strategies versus the convexity of the cost function. Graph A corresponds to the case with only supply shocks and Graph B to the case with only demand shocks.

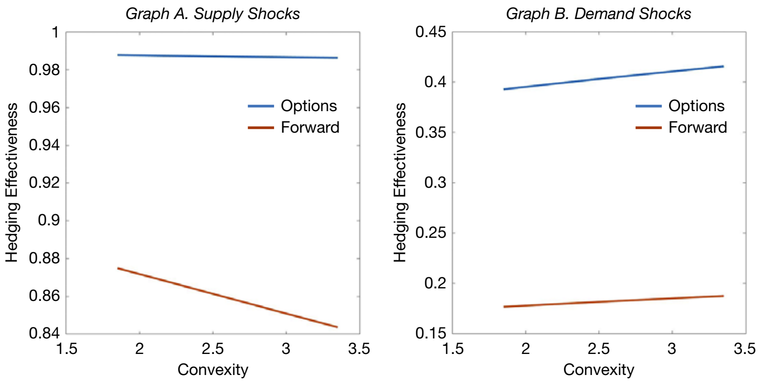

when uncertainty comes from demand shocks alone (this is the case where hedging with the forward-based strategy is particularly ineffective). When both supply and demand shocks are present, hedging, either with a forward-based strategy or an option-based strategy is not effective, and the difference between the two is small. This result confirms the intuition that, with two sources of uncertainty, the relationship between the spread and the input price can be very weak.

2. Effect of Production Cost Function Convexity

The convexity of the cost function can vary across industries and also over time for the same industry. For example, the introduction of new plants and the retirement of old ones may flatten the cost curve and reduce its convexity.

We examine this relationship in our model by varying the magnitude of the convexity coefficient of the cost function, λ. Figure 5 illustrates hedging effectiveness for two cases. Graph A corresponds to a market with only supply shocks, where $\sigma_C = 0.6$. Graph B corresponds to a market in which uncertainty is concentrated on the demand factor (the standard deviation of the supply factor $\sigma_C = 0.0$, i.e., the supply factor is constant, while the standard deviation of the demand factor is $\sigma_G = 0.6$).

We note that, in both cases, option-based strategies significantly improve hedging effectiveness.

These results are in line with the intuition that, as convexity increases, the relationship between the input price and the spread becomes more nonlinear, leading to a large improvement when hedging with options.

Overall, we find that hedging with forward contracts is effective when risk comes from supply shocks but not demand shocks, and is somewhat effective when risk comes from demand shocks but not supply shocks. With both supply and demand shocks, hedging with forward contracts is not a particularly effective hedging strategy. Option-based strategies, while more effective than forward-based ones, also tend to be ineffective when both supply and demand shocks are present.

TABLE 8

Hedging Effectiveness: Static Strategy

Table 8 presents the hedging effectiveness of forward-based strategies against the standard deviation of supply and demand shocks when the time step is very small.

Std. Dev. of Demand	Std. Dev. of Supply						
	0.0	0.10	0.20	0.30	0.40	0.50	0.60
0.0	–	1.00	1.00	0.99	0.99	0.99	0.99
0.10	0.94	0.18	0.06	0.39	0.63	0.77	0.85
0.20	0.76	0.48	0.07	0.01	0.18	0.37	0.53
0.30	0.60	0.47	0.20	0.02	0.01	0.11	0.25
0.40	0.50	0.43	0.25	0.09	0.00	0.01	0.08
0.50	0.44	0.39	0.27	0.14	0.04	0.00	0.01
0.60	0.39	0.37	0.28	0.17	0.08	0.02	0.00

TABLE 9

Hedging Effectiveness: Dynamic Strategy

Table 9 presents the effectiveness of a dynamic hedging strategy with an annual horizon and daily adjustments against the annualized standard deviation of supply and demand shocks.

Std. Dev. of Demand	Std. Dev. of Supply						
	0.0	0.10	0.20	0.30	0.40	0.50	0.60
0.0	–	1.00	1.00	0.99	0.99	0.99	0.99
0.10	0.95	0.19	0.19	0.43	0.59	0.69	0.75
0.20	0.82	0.49	0.19	0.18	0.29	0.41	0.50
0.30	0.70	0.54	0.30	0.19	0.19	0.25	0.33
0.40	0.62	0.53	0.36	0.24	0.19	0.19	0.23
0.50	0.55	0.50	0.38	0.28	0.21	0.19	0.19
0.60	0.51	0.48	0.39	0.30	0.23	0.19	0.18

B. Dynamic Hedging

So far we have limited our analysis to discrete-time hedging with a single, annual, time step. Adjusting the hedge ratio over shorter intervals has the potential to improve hedging performance: Based on our model, the hedge ratio depends on the level of supply and demand, both of which change over time. We investigate this possibility for two cases: First, by determining the improvement in hedging efficiency when the time step is shorter. Second, by considering the impact of a dynamic hedging strategy where, rather than a hedge with a single, annual, time step, we rebalance the hedge daily and evaluate hedging effectiveness over a year. In both cases, we only use forward contracts to hedge; results for the option-based hedging strategy are similar.

Table 8 presents the results when the time step is short: one trading day. The results show that hedging effectiveness improves, especially with a single source of uncertainty. However, with uncertainty in both supply and demand, hedging effectiveness remains low. We note that, similar to the case of an annual time step, there are combinations of standard deviations of supply and demand shocks for which hedging effectiveness is close to 0.

The results with a short time step suggest that rebalancing frequently has benefits. Considering the effectiveness of dynamic hedging for the base case set of parameters, we find that, over an annual horizon, the strategy that rebalances the hedge monthly achieves a hedging effectiveness of 28%. The effectiveness increases to 32% when we rebalance weekly, and to 49% when we rebalance daily. Comparing with the hedging effectiveness of 16% for the strategy maintaining the same hedge over the entire year, it is clear that effectiveness improves.

We explore the improvement further in Table 9, where we present the hedging effectiveness of the dynamic hedging policy with daily adjustments for a range of demand and supply volatility parameters. Similar to the case of Table 8, we find that improvement is most evident with a single source of uncertainty. While these results are encouraging, frequent hedging has challenges, associated with the accurate estimation of the supply and demand shocks, and potential costs. Nonetheless, a well-calibrated structural model that allows the accurate estimation of the hedge, has the potential to improve effectiveness.

VIII. Conclusions

There is a substantial literature that examines the motivations for firms to reduce the variance of their cash flows by hedging. This literature implicitly assumes that firms can in fact effectively hedge. In this article, we model the economic factors that influence the effectiveness of hedging and conclude that, under reasonably plausible conditions, hedging the risks associated with uncertain input prices only modestly reduces the variance of profits.

As our model illustrates, hedging the risk associated with uncertain input costs tends to be ineffective for two reasons. The first is that there are typically two sources of uncertainty (supply shocks and demand shocks) and only one hedging instrument. The second is that the nonlinearity of the cost function implies that the optimal hedge ratio changes as the level of production changes. While we have shown that using options can address nonlinearity, and improves hedging effectiveness with a single source of uncertainty, hedging remains ineffective when there is uncertainty about both supply and demand.

To a large extent, our model provides a best-case scenario. In reality, the challenges associated with effective hedging may be even greater. For tractability reasons, we have assumed that the parameters are all known and that firms are identical. Effective hedging becomes less effective if firms have imperfect information about parameters governing the supply and demand functions as well as imperfect information about the production functions of their competitors. In addition, if the availability of production capacity varies materially over time, as firms take their production off-line due to maintenance or accidents, or if production costs change stochastically due to, for example, changes in other unhedgeable costs, a substantial portion of the variation in the spread between output and input prices will be very difficult to hedge.

Although our analysis is pessimistic about the effectiveness of hedging, we do provide insights that can help improve the effectiveness in some situations. For example, our model suggests that variance-minimizing hedge ratios are likely to change over time as the relative variance of supply and demand shocks change. We have also shown that the relative magnitudes of the variances of supply and demand shocks influence the volatilities of the input commodity and the output product. One might be able to use information from the implied volatilities of option prices on both input and output to come up with more effective hedging strategies.

Finally, firms can improve hedging by using multiple instruments to hedge. In the case of a refiner, they can hedge both the inputs and the outputs. In some cases, it might also be possible to come up with more effective hedges using other commodities. For example, airlines often hedge their fuel exposure with forward purchases of crude oil rather than jet fuel. It is possible that crude oil provides a better profit hedge than jet fuel because crude oil prices are less sensitive to demand shocks than jet fuel prices. For other firms, operational hedges and vertical integration may more effectively reduce a firm's risk exposure than a financial hedge.

While our article focuses on commodity price risk, our results are applicable more broadly. One example is the exposure to exchange rate risk. There is a large literature that describes what is often referred to as the "currency exposure puzzle," which is the observation that the stock prices and earnings of exporters tend to be only modestly correlated with currency changes (see, e.g., Bartram, Brown, and Minton (2010) and Hoberg and Moon (2017)). Our model can provide a potential explanation for this puzzle, beyond those already considered in the literature. The idea is that, similar to the setup in our article, the strength of a country's currency can change for reasons that are roughly related to demand shocks and supply shocks. Consider a country with two sectors: mining and tourism. The country's currency may strengthen because of a boost in tourism, which is orthogonal to the demand for the resources produced in the country. Such a shock may increase the cost of mining without increasing the price of the commodity; thus, the country's mining firms may become less competitive internationally and their stock prices may decline. However, a second possibility is that the currency strengthens because of a positive shock to the demand for the country's resources, in which case, the strengthening of the currency may be associated with increases in their mining stock prices. The two offsetting channels may result in a low correlation between the mining firms' stock prices and the exchange rate. As in our model, if the nature of the shocks is not known in advance, one may not be able to predict the direction of the currency exposure.

Appendix A. Equilibrium Quantity Produced and Prices for the Competitive Case

Solving equation (4), the equilibrium quantity produced is given by

$$Q^* = \frac{X_d - P_{I,g} - X_s}{\gamma_d + \gamma_s + \lambda}. \tag{A-1}$$

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

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

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

$$\text{(A-3)}\qquad \begin{aligned} \text{var}(P_c) &= \left(\frac{\gamma_d+\lambda}{\gamma_d+\gamma_s+\lambda}\right)^2 \text{var}(X_s) + \left(\frac{\gamma_s}{\gamma_d+\gamma_s+\lambda}\right)^2 \text{var}(X_d), \\ \text{var}(P_g) &= \left(\frac{\gamma_d}{\gamma_d+\gamma_s+\lambda}\right)^2 \text{var}(X_s) + \left(\frac{\gamma_s+\lambda}{\gamma_d+\gamma_s+\lambda}\right)^2 \text{var}(X_d), \\ \text{var}(P_g-P_c) &= \frac{\lambda^2}{(\gamma_d+\gamma_s+\lambda)^2}(\text{var}(X_d)+\text{var}(X_s)), \\ \text{var}(Q^*) &= \frac{1}{(\gamma_d+\gamma_s+\lambda)^2}(\text{var}(X_d)+\text{var}(X_s)), \\ \text{cov}(P_g-P_c,P_c) &= \frac{-\lambda(\lambda+\gamma_d)\text{var}(X_s)+\lambda\gamma_s\text{var}(X_d)}{(\gamma_d+\gamma_s+\lambda)^2}, \\ \text{cov}(P_g-P_c,P_g) &= \frac{\lambda(\lambda+\gamma_s)\text{var}(X_d)-\lambda\gamma_d\text{var}(X_s)}{(\gamma_d+\gamma_s+\lambda)^2}. \end{aligned}$$

A.1. *Proposition 1*

From equation (A-3) we note that when $(\lambda+\gamma_d)\text{var}(X_s) > \gamma_s\text{var}(X_d)$, the numerator of the covariance, $\text{cov}(P_g-P_c,P_c)$, is a negative number. Given that the denominator is always positive, the sign is determined by the relative size of $\text{var}(X_s)$ and $\text{var}(X_d)$ normalized to the relative elasticity of demand and supply functions. A similar argument applies to $\text{cov}(P_g-P_c,P_g)$.

The remaining results in the proposition follow from equation (A-3), which shows that

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

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

A.2. *Proposition 2*

The proof of Proposition 2 follows directly from equation (A-3).

A.3. *Proposition 3*

Proposition 3 follows from the formula for the optimal hedge ratio and from equation (A-3)

$$\text{(A-4)}\qquad -\frac{\text{cov}(P_g-P_c,P_c)}{\text{var}(P_c)} = -\frac{-\lambda(\lambda+\gamma_d)\text{var}(X_s)+\lambda\gamma_s\text{var}(X_d)}{(\lambda+\gamma_d)^2\text{var}(X_s)+\gamma_s^2\text{var}(X_d)}.$$

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

From equation (A-4), 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., $\text{var}(X_s)=0,$) and demand is uncertain, the optimal hedge is to sell λ/γ_s forward contracts. The remaining parts of the proposition follow from equation (A-4) in a similar way.

A.4. *Proposition 4*

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.

$$R^2 = \frac{\left(-(\lambda+\gamma_d)\text{var}(X_s)+\gamma_s\text{var}(X_d)\right)^2}{\left((\lambda+\gamma_d)^2\text{var}(X_s)+\gamma_s^2\text{var}(X_d)\right)\left(\text{var}(X_s)+\text{var}(X_d)\right)}. \tag{A-5}$$

From equation (A-5), it follows 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 variances of supply and demand are such that the hedge ratio is 0; that is, when $(\lambda+\gamma_d)\text{var}(X_d)=\gamma_s\text{var}(X_s)$, the effectiveness of the hedge is also 0. In this case, there is no hedging benefit in using 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 0.

Appendix B. Equilibrium Quantity Produced and Prices When the Producer is a Monopolist

For the case of a firm with market power, the optimal quantity produced is given by the solution to the maximization problem in equation (7).

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

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

$$\begin{aligned} P_c &= X_s+\gamma_s Q_M^* = \frac{(2\gamma_d+\lambda)X_s+\gamma_s X_d-\gamma_s P_{I,g}}{2\gamma_d+\gamma_s+\lambda}, \\ P_g &= X_d-\gamma_d Q_M^* = \frac{(\gamma_s+\gamma_d+\lambda)X_d+\gamma_d X_s+\gamma_d P_{I,g}}{2\gamma_d+\gamma_s+\lambda}, \\ P_g-P_c &= X_d-X_s-(\gamma_s+\gamma_d)Q_M^* = \frac{(\lambda+\gamma_d)(X_d-X_s)+(\gamma_d+\gamma_s)P_{I,G}}{2\gamma_d+\gamma_s+\lambda}. \end{aligned} \tag{B-2}$$

B.1. *Proposition 5*

Comparing Q^* and Q_M^* from equations (A-1) and (B-1), we observe that the numerators of the two equations are the same but the denominator is greater in the case with market power. Thus, the level and the variance of quantity are both smaller in the case with market power.

B.2. *Proposition 6*

The hedge ratio for a firm with market power is given by

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

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

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

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 0.

References

Bartram, S. M. "Corporate Hedging and Speculation with Derivatives." *Journal of Corporate Finance*, 57 (2017), 9–34.

Bartram, S. M.; G. W. Brown; and B. A. Minton. "Resolving the Exposure Puzzle: The Many Facets of Exchange Rate Exposure." *Journal of Financial Economics*, 95 (2010), 148–173.

Bessembinder, H., and M. L. Lemmon. "Equilibrium Pricing and Optimal Hedging in Electricity Forward Markets." *Journal of Finance*, 57 (2002), 1347–1382.

Breeden, D. T., and S. Viswanathan. "Why Do Firms Hedge? An Asymmetric Information Model." *Journal of Fixed Income*, 25 (2015), 7–25.

Brown, G. W., and K. B. Toft. "How Firms Should Hedge." *Review of Financial Studies*, 15 (2002), 1283–1324.

Carter, D. A.; D. A. Rogers; B. J. Simkins; and S. D. Treanor. "A Review of the Literature on Commodity Risk Management." *Journal of Commodity Markets*, 8 (2017), 1–17.

Casassus, J.; P. Liu; and K. Tang. "Economic Linkages, Relative Scarcity, and Commodity Futures Returns." *Review of Financial Studies*, 26 (2012), 1324–1362.

Cooper, J. C. "Price Elasticity of Demand for Crude Oil: Estimates for 23 Countries." *OPEC Energy Review*, 27 (2003), 1–8.

Dale, S. "New Economics of Oil." *Oil and Gas, Natural Resources, and Energy Journal*, 1 (2016), 365.

DeMarzo, P. M., and D. Duffie. "Corporate Incentives for Hedging and Hedge Accounting." *Review of Financial Studies*, 8 (1995), 743–771.

Dybvig, P. H.; P. J. Liang; and W. J. Marshall. "The New Risk Management: The Good, the Bad, and the Ugly." *Federal Reserve Bank of St. Louis Review*, 95 (2013), 273–292.

Fehle, F., and S. Tsyplakov. "Dynamic Risk Management: Theory and Evidence." *Journal of Financial Economics*, 78 (2005), 3–47.

Froot, K. A.; D. S. Scharfstein; and J. C. Stein. "Risk Management: Coordinating Corporate Investment and Financing Policies." *Journal of Finance*, 48 (1993), 1629–1658.

Gilje, E. P., and J. P. Taillard. "Does Hedging Affect Firm Value? Evidence from a Natural Experiment." *Review of Financial Studies*, 30 (2017), 4083–4132.

Graham, J. R., and C. W. Smith. "Tax Incentives to Hedge." *Journal of Finance*, 54 (1999), 2241–2262.

Guay, W., and S. P. Kothari. "How Much Do Firms Hedge with Derivatives?" *Journal of Financial Economics*, 70 (2003), 423–461.

Hirshleifer, D. "Residual Risk, Trading Costs, and Commodity Futures Risk Premia." *Review of Financial Studies*, 1 (1988a), 173–193.

Hirshleifer, D. "Risk, Futures Pricing, and the Organization of Production in Commodity Markets." *Journal of Political Economy*, 96 (1988b), 1206–1220.

Hirshleifer, D. "Seasonal Patterns of Futures Hedging and the Resolution of Output Uncertainty." *Journal of Economic Theory*, 3 (1991), 304–327.

Ho, T. S. "Intertemporal Commodity Futures Hedging and the Production Decision." *Journal of Finance*, 39 (1984), 351–376.

Hoberg, G., and S. K. Moon. "Offshore Activities and Financial vs Operational Hedging." *Journal of Financial Economics*, 125 (2017), 217–244.

Hughes, J.; C. Knittel; and D. Sperling. "Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand." *Energy Journal*, 29 (2008), 113–134.

Jin, Y., and P. Jorion. "Firm Value and Hedging: Evidence from US Oil and Gas Producers." *Journal of Finance*, 61 (2006), 893–919.

Kamara, A. "Production Flexibility, Stochastic Separation, Hedging, and Futures Prices." *Review of Financial Studies*, 6 (1993), 935–957.

Lievenbrück, M., and T. Schmid. "Why Do Firms (not) Hedge? Novel Evidence on Cultural Influence." *Journal of Corporate Finance*, 25 (2014), 92–106.

Liu, W. "Modeling Gasoline Demand in the United States: A Flexible Semiparametric Approach." *Energy Economics*, 45 (2014), 244–253.

Mello, A. S., and J. E. Parsons. "Hedging and Liquidity." *Review of Financial Studies*, 13 (2000), 127–153.

Routledge, B. R.; C. S. Spatt; and D. J. Seppi. "The 'Spark Spread': An Equilibrium Model of Cross-Commodity Price Relationships in Electricity." gSIA Working Papers 1999-15, Carnegie Mellon University (1998).

Shapiro, A. C., and S. Titman. "An Integrated Approach to Corporate Risk Management." In J. Stern and D. Chew, eds. *The Revolution in Corporate Finance*. New York, NY: Basil Blackwell (1986), 331–354.

Smith, C. W., and R. M. Stulz. "The Determinants of Firms' Hedging Policies." *Journal of Financial and Quantitative Analysis*, 20 (1985), 391–405.

Strebulaev, I. A., and T. M. Whited. "Dynamic Models and Structural Estimation in Corporate Finance." *Foundations and Trends in Finance*, 6 (2012), 1–163.