How Firms Should Hedge

Gregory W. Brown
University of North Carolina

Klaus Bjerre Toft
Goldman Sachs International

Substantial academic research explains *why* firms should hedge, but little work has addressed *how* firms should hedge. We assume that firms can experience costly states of nature and derive optimal hedging strategies using vanilla derivatives (e.g., forwards and options) and custom "exotic" derivative contracts for a value-maximizing firm facing both hedgable (price) and unhedgable (quantity) risks. Customized exotic derivatives are typically better than vanilla contracts when correlations between prices and quantities are large in magnitude and when quantity risks are substantially greater than price risks. Finally, we discuss how our model may be applied in practice.

In the last 20 years, financial derivative use by nonfinancial corporations has become commonplace [see Bodnar, Hyat, and Marston (1998)]. Recent research has explained why firms use derivatives to manage risk.[1] As in other areas of corporate finance, violations of the assumptions underlying the Modigliani and Miller (1958) hypotheses result in value-increasing financial strategies. Risk management is no exception; taxes, bankruptcy costs, agency costs, etc. have all been employed to explain why firms should hedge. Surprisingly, little research has attempted to explain *how* firms should hedge.[2] For example, the current literature has not addressed the problem of when a firm will be better off using options over forwards, or perhaps foregoing both

---

Opinions in this paper are not those of Goldman Sachs International. The second author's contribution to this work was done while he was on the faculty at the University of Texas at Austin. The authors thank Dong-Hyun Ahn, Keith Brown, John Butler, Jennifer Carpenter, Dave Chapman, Jeff Fleming, Bin Gao, Jay Hartzell, Jim Hodder, Patrick Jaillet, Anthony Lynch, Barbara Ostdiek, Hui Ou-Yang, Neil Pearson, Laura Starks, Sheridan Titman, and Robert Whitelaw for their valuable assistance and suggestions. The authors also thank participants at the 1999 meetings of the American Finance Association, the Chicago Risk Management Conference, the 1998 meetings of the International Association of Financial Engineers and the Financial Management Association, the IAFE/IEEE 1998 Conference on Computational Intelligence in Financial Engineering, the Chicago Board of Trade 1997 fall research seminar, the University of Texas Institute for Computational Finance, and the Danske Bank Symposium (Odense University). Finally, we thank workshop participants at Dartmouth College, New York University, the Federal Reserve Board, University of Maryland, Rice University, University of North Carolina at Chapel Hill, University of Georgia, University of Florida, Emory University, and University of Wisconsin at Madison. Address correspondence to Gregory W. Brown, Department of Finance, Kenan-Flagler School of Business, University of North Carolina at Chapel Hill, Campus Box 3490, McColl Building, Chapel Hill, NC 27599-3490, or e-mail: gregwbrown@unc.edu.

[1] Research suggests that firms are not systematically speculating with derivatives [see Hentschel and Kothari (2001)]. This research also details the breakdown by derivative type and market for financial and nonfinancial firms.

[2] Some recent research on value-at-risk has considered alternatives to simple linear hedges; for example, Ahn et al. (1999) consider minimizing value-at-risk with put options. Petersen and Thiagarajan (2000) examine alternatives to derivatives. Mello, Parsons, and Triantis (1995) derive the optimal foreign-exchange hedging strategy for a firm with international flexibility in its production location.

The Review of Financial Studies Fall 2002 Vol. 15, No. 4, pp. 1283–1324

and entering into a customized contract with an over-the-counter derivatives dealer. In this article we address this question by constructing a simple model of a nonfinancial firm and allowing this firm to enter into a variety of derivative contracts. Ultimately our goals are to compare the relative effectiveness of different hedging tactics and identify firm characteristics most important for determining how the firm should structure its risk management strategy.

As noted, a substantial body of research has already provided theoretical justifications for corporate risk management. For example, Smith and Stulz (1985) show that value-maximizing firms should hedge if they face a convex corporate tax schedule or deadweight costs associated with financial distress.[3] These deadweight costs can go beyond the direct costs associated with bankruptcy or the loss of future tax benefits from debt financing. Stakeholders such as customers, suppliers, and employees may anticipate financial distress and therefore seek to reduce their long-term dependence on firms with a high likelihood of bankruptcy [see Shapiro and Titman (1986)]. Other indirect costs associated with financial distress include agency costs such as those associated with Myers' (1977) underinvestment problem. Capital market imperfections also motivate hedging strategies. When it is costly to access external capital markets, Froot, Scharfstein, and Stein (1993) suggest that firms use risk management to reduce the expected cost of financing future investments.[4]

The above studies suggest the existence of significant deadweight costs in certain states of nature. The goal of our analysis is not to expand on the reasons why firms hedge, but instead to take these as given and derive the optimal hedging strategy for a value-maximizing firm.[5] As is the case in practice, we also assume that firms face multiple sources of uncertainty of which some are unhedgable. A simple example is that of a wheat farmer: A variety of derivative securities make it possible for the farmer to manage exposure to wheat prices. While it is relatively easy for the farmer to hedge a given quantity, it is more difficult to anticipate (and hedge) the produced quantity. Another example is that of a personal computer manufacturer that produces domestically and sells in an overseas market. It is relatively easy for the firm to forecast and hedge sales measured in foreign currency in the

[3] This argument is supported by Graham and Smith (1999), who show that the current tax code on average results in a convex effective tax schedule.

[4] Another strain of literature bases its analysis on managerial incentives. Stulz (1984) shows that corporate hedging policies may be driven by risk-averse managers in a world where their compensation is proportional to the value of the firm's assets. When these managers maximize their personal utility, the firm will in most situations hedge. DeMarzo and Duffie (1995) and Breeden and Viswanathan (1996) argue that high-quality managers may have the incentive to hedge so outsiders can observe their superior skills. Furthermore, DeMarzo and Duffie (1991) show that shareholders of firms with valuable proprietary information may prefer that the firm maintain the information asymmetry and hedge for the shareholders so as to preserve the value of the private information.

[5] For simplicity, we assume that these deadweight costs can be expressed as a function of the firm's profits. The parameters of this function are chosen to reflect the aggregate incentive to hedge at a firmwide level. As such, they will differ from firm to firm.

short term. In the longer term, foreign sales are more uncertain, and therefore also more difficult to hedge.

The assumptions of costly states of nature and unhedgable risks permit us to derive a one-period model of the optimal hedging strategy using different types of derivative contracts. In fact, these two assumptions represent the bare minimum for obtaining nontrivial hedging strategies. Without any deadweight costs, the firm has no reason to hedge. Without any unhedgable risk, the optimal hedging strategy is to simply sell the entire exposure forward, which results in a perfect hedge. We intentionally concentrate on this simplest setting for two reasons. First, it provides for a surprisingly rich set of optimal hedging strategies. Second, it allows us to more easily identify the factors most (and least) important for corporate risk management. In this sense, our model provides a basic framework that can be easily expanded to incorporate additional features such as multiple future periods, market views (and convenience yields), and alternative hedging strategies that do not involve derivatives (e.g., foreign debt).

We analyze two general types of hedging strategies. First, we assume that the corporate risk management policy only allows for the use of "vanilla" derivative contracts such as forwards and simple options. In the case of forwards, the firm's optimal hedging strategy can be expressed in closed form. Our analysis of the optimal forward hedge leads to a series of intuitive results. For example, when prices are negatively correlated with produced quantities, the firm should typically hedge less than its expected exposure. We also find that firm-specific factors, such as the volatility of future sales, and market-specific factors, such as the volatility of the hedgable price risk, have large impacts on the optimal forward hedge. We show that the minimum-variance forward hedge is a special case of our more general model, but since the minimum-variance hedge ignores production technology and the fundamental reason why a firm is hedging, it always leads to deadweight costs that are at least as great as those of our optimal hedge. We also discuss how a firm could use put options or a portfolio of vanilla derivatives as its hedging instrument. For example, a long position in put options is often superior to selling forward contracts when price and quantity risks are negatively correlated.

The second type of strategy, and the most significant contribution of this article, involves the derivation of a closed-form solution for the optimal "exotic" payoff function. This function describes the payoff the firm should choose if it can contract in any fairly priced derivative. We call it an exotic derivative because it typically cannot be replicated with forwards and a finite number of vanilla options. We show that there is a very close relationship between the optimal number of forward contracts and this optimal exotic hedge: At the forward price of the traded commodity, the optimal forward contract and the optimal custom hedge have identical "deltas." At values different from the forward price, the exotic hedge fine tunes the firm's exposure by adding (or subtracting) convexity.

Since our results are expressed in closed form, we can also easily identify which firm types benefit most from buying a nonlinear exotic hedge. We show that price and quantity correlation, the degree of price and quantity volatility, and the ratio of these risks are the primary determinants of the optimal hedge's convexity. For example, firms should typically buy convexity (i.e., options) when correlation is negative. However, when correlation is positive, the optimal custom hedge usually (but not always) requires the firm to sell convexity. The exact degree of convexity is determined by price and quantity risk, and to a lesser degree the relative convexity of the deadweight cost function. Typically, high levels of quantity risk lead to more "optionality" in the optimal hedge.

A major advantage to this approach of determining an optimal hedge is that it allows for a simple, but powerful comparison between different hedging instruments. Because the firm is seeking to reduce deadweight costs with different hedging strategies, we can quantify the relative effectiveness of various alternatives. For example, we find that firms can benefit most from nonlinear exotic payoffs when the correlation between price and quantity is negative and quantity risk is large. If correlation between price and quantity is negligible, forward contracts are typically very effective hedging tools. When correlation is positive, exotic derivatives offer additional gains over forwards or options alone, and these gains increase with greater quantity risk and less price risk.

While we believe this article is the first to concentrate on how a value-maximizing corporation with unhedgable risks should structure its hedging program, researchers have analyzed similar problems in other areas. For example, our analysis is related to the work by Leland (1980), Brennan and Solanki (1981), and Carr, Jin, and Madan (2001), who determine optimal payoff functions for investors with certain preferences and beliefs. Leland (1980) uses a general utility specification and finds that an investor should hold a strictly convex payoff schedule if her coefficient of risk aversion increases more rapidly than that of the representative agent. Brennan and Solanki (1981), and Carr, Jin, and Madan (2001) determine optimal payoff functions for specific utility functions.

Another set of related research examines investors that can hedge a non-traded exposure by continuously transacting in a correlated and traded asset, for example, a futures contract on a similar asset. (In general, dynamic trading strategies in complete markets are equivalent to our "exotic" derivatives.) For example, Duffie and Richardson (1991) determine the optimal dynamic trading strategy for an investor with quadratic utility and access to a correlated futures contract. Similar to our results, Duffie and Richardson find that a dynamic hedge can result in a significantly lower variance of terminal wealth for certain parameter values. Other research that examines similar problems includes Adler and Detemple (1988).

Also related to our model is the research on optimal portfolio choice when some assets are nontraded and markets are incomplete. Svensson and Werner (1993) examine the portfolio decision and implicit hedging strategy when an investor has a nontraded, exogenous, and stochastic income. They show the solution for the optimal portfolio of traded assets contains a component that hedges the investors nontraded income risk. Duffie and Zariphopoulou (1993), He and Pagès (1993), Cuoco (1997), and Duffie et al. (1997) also consider optimal portfolio choice when markets are incomplete.

However, there exists two subtle but important differences between this prior research and the model presented here. First, all of the aforementioned solve a utility maximization problem for a single (or representative) agent, whereas we focus exclusively on a value-maximizing firm. The differences in the hedging problems have been pointed out by Froot and Stein (1998) and examined in detail by Brown and Khokher (2001). The difference is intuitive if we consider the simpler case of complete markets. If markets are complete and a firm experiences deadweight costs associated with low-profit states, the firm will hedge exposure to a risky asset completely. This compares to a risk-averse investor that will optimally retain some exposure to the risky asset for speculative purposes.

Second, instead of considering an additive, but correlated, nontraded risk, we consider a firm that faces a "quantity risk" where the price and quantity of production is uncertain. This multiplicative risk better characterizes the problem of a nonfinancial firm and leads to solutions that are distinct from those in the existing literature on mean variance hedging and portfolio choice. We differentiate our work from that of Froot and Stein (1998) in two other ways. First, we conceptually focus on a nonfinancial firm instead of a financial institution; second, we concentrate on the security design (instead of investment and capital structure) aspects of hedging.

Most similar to our model are hedging problems with nontraded (quantity) risk, which have been analyzed in the agricultural economics literature. For example, Rolfo (1980) derives the optimal futures hedging strategy for a sovereign entity by assuming that the country's utility can be expressed as a function of total revenue. Moschini and Lapan (1995) and Lapan, Moschini, and Hanson (1991) analyze the optimal hedge for a risk-averse farmer that can use options as well as futures. Others have solved the "uncertain quantity" hedging problem for firms wishing to hedge the variance of total revenues [see Siegel and Siegel (1988)].

The article proceeds as follows. Section 1 presents the model and solutions for optimal hedging policies. Section 2 analyzes optimal hedges in different scenarios and the model's sensitivity to changes in assumptions. Section 3 discusses a real-world application of the model, general empirical implications, and how our work relates to existing empirical results. Section 4 concludes.

1. Optimal Corporate Hedging Policies

1.1 General model specification

Most generally, we consider a price-taking firm that realizes an uncertain total revenue at a future date. The firm earns profits from operations which we define as the product of price and quantity minus the costs of production. There are two sources of uncertainty. First, the quantities of goods produced (and/or used in production) are uncertain and are not hedgable. Second, the prices of output (and/or input) goods are unknown today, but are hedgable by way of derivative securities. We assume that the firm faces deadweight costs that are a function of future states of nature. These costs (which could result from direct and indirect bankruptcy costs, costly external capital, a convex tax schedule, and agency costs, etc.) are exogenous and can be expressed as a function of the firm's net profits (i.e., profit from operation plus the profit from the hedge). The firm may therefore have an incentive to use derivatives to modify its cash flows in future states.

Consider a firm that produces a single commodity. At a future date, $t = 1$, it will realize an uncertain production quantity, q. While we will explicitly consider firms producing an uncertain quantity, the model allows inputs to production to be modeled as negative production quantities. Similarly the model can be expanded to include multiple produced and consumed uncertain quantities. At the future date the price, p, of the produced quantity is also uncertain. Profits from operations are defined by the function $f(p, q)$. Again, we assume that output prices can be hedged (e.g., the price of wheat or foreign exchange rates), but that the produced quantity cannot (e.g., bushels of wheat produced or sales revenues in foreign currency). The firm determines the optimal hedge by using a risk-adjusted joint density of p and q defined by $h(p, q)$.

The firm is value maximizing, but faces an exogenous deadweight cost function $C(P)$, where P is the uncertain profit at $t = 1$ net of the payoff from the firm's derivatives transactions. These deadweight costs need only be from the perspective of the shareholder. As discussed, we assume these costs can be accurately summarized as a function of the firm's net profits, P. This specification can capture most of the popular motivations for why a firm hedges. For example, a convex tax schedule can be represented as a piecewise linear cost function. Financial distress can be represented as deadweight costs in low profit states. Costly external capital can be represented by specifying costs in profit states where additional funds must be raised from outside the firm. (We discuss the impact of these different reasons why the firm hedges on how the firm hedges in Section 2.4.)

To hedge the marketable risks the firm can enter into a contract at $t = 0$ that pays $x(p; \mathbf{a})$ at $t = 1$, where the vector $\mathbf{a}$ represents the parameters of the contract (e.g., number of contracts, strike prices, etc.). The fairly priced contract $x(p; \mathbf{a})$ costs the firm $X(\mathbf{a})$. Without loss of generality, we assume the risk-free interest rate is zero, such that,

$$X(\mathbf{a}) = \int_p x(p; \mathbf{a}) g^*(p)\, dp, \tag{1}$$

where $g^*(p)$ is the risk-neutral pricing density of p. The firm's net profit function (before deadweight costs), $P(p, q; \mathbf{a})$, can be written as

$$P(p, q; \mathbf{a}) = f(p, q) + x(p; \mathbf{a}) - X(\mathbf{a}). \tag{2}$$

The derivative's cost is paid at $t = 1$, that is, all derivative securities are deferred payment instruments. (This prevents the firm from using derivatives to transfer cash across time.) We assume the firm maximizes expected net economic profit, that is, profits net of deadweight costs. Formally it solves the maximization problem:

$$\max_a \pi(\mathbf{a}) = \int_{p,q} \{f(p, q) + x(p; \mathbf{a}) - C[P(p, q; \mathbf{a})]\} \times h(p, q)\, dq\, dp - X(\mathbf{a}). \tag{3}$$

Hence the firm takes the joint distribution of (p, q) as given and chooses the parameters $\mathbf{a}$ of available derivative contracts to maximize expected net economic profit, $\pi(\mathbf{a})$. If we assume that the matrix $\partial^2 \pi(\mathbf{a})/\partial \mathbf{a} \partial \mathbf{a}'$ is negative semidefinite, then value maximization is equivalent to solving the first-order condition,

$$\frac{\partial \pi(\mathbf{a})}{\partial \mathbf{a}} = \int_{p,q} \left\{ \frac{\partial x(p; \mathbf{a})}{\partial \mathbf{a}} - \frac{\partial C(P)}{\partial P} \frac{\partial P(p, q; \mathbf{a})}{\partial \mathbf{a}} \right\} \times h(p, q)\, dq\, dp - \frac{\partial X(\mathbf{a})}{\partial \mathbf{a}} = 0. \tag{4}$$

For situations where the marginal density of p, $g(p) = \int_q h(p, q)\, dq$ equals the risk-neutral density, we can use the derivative pricing constraint [Equation (1)] to simplify the first-order condition to

$$-\int_{p,q} \frac{\partial C(P)}{\partial P} \frac{\partial P(p, q; \mathbf{a})}{\partial \mathbf{a}} h(p, q)\, dq\, dp = 0. \tag{5}$$

In this case, maximization of the risk-adjusted expected net economic profit is equivalent to value maximization. This case will be the focus of the subsequent analysis. However, in the more general setting, we can allow the firm to incorporate a view on the level of future commodity prices, that is, $g(p) \neq g^*(p)$.[6]

[6] See Stulz (1996) and Brown and Khokher (2001) for analysis of the impact of managerial views on optimal corporate hedging strategies. For the remainder of the analysis we assume $g(p) = g^*(p)$.

Finally, note that we take as predetermined the investment and cost structure of the firm. In addition, since we have not explicitly modeled the capital structure of the firm, we can think of debt service requirements as a part of the firm's net profit function. While in the long run these factors are most likely determined jointly with the firm's hedging policy, we prefer to think of this model as describing the optimal hedging strategy for the next period. For example, the computer manufacturer's factory is in place, it has made its investment in research and development, and has already determined its optimal capital structure. Since these factors are both expensive and time-consuming variables to adjust, the firm looks to its risk management program to determine the optimal short run hedging policy conditional on its near-term forecasts of price, demand, market conditions, etc. In fact, the inflexibility of investment and capital structure decisions is one potential source of the deadweight costs that drive the risk management program. Taking these factors as predetermined allows us to concentrate on the decision of how the firm should hedge in the short run, without attempting to say anything about how the firm may find it optimal to adjust its level of investment, capital structure, and product mix in the long run.

1.2 The one-product firm's operating environment

The hedging implications of the preceding model are most easily interpreted by considering a specific example of a one-product firm. We operationalize our model by specifying a production function, a joint density for (p, q), assuming a deadweight cost function $C(P)$, and restricting the set of available derivative contracts. Production costs are assumed to be linear in quantity, with a known per unit variable cost, s_1, and a known fixed cost, s_2, so that profits from operations are

$$f(p, q) = pq - s_1 q - s_2, \qquad s_1, s_2 \in R_+. \tag{6}$$

We assume that the joint density of (p, q) is bivariate-normal (with correlation coefficient ρ, expected price level, μ_p,[7] expected quantity level, μ_q, standard deviation of the price, σ_p, and standard deviation of the quantity, σ_q). We show later that our conclusions are relatively insensitive to the exact distributional assumption.

For now, we specify the exogenous deadweight cost function as exponential of the form

$$C(P) = c_1 e^{-c_2 P}, \qquad c_1, c_2 \in R_{++}. \tag{7}$$

This cost function is consistent with a firm that experiences high costs when profits are low ("bad" states of nature) and low costs when profits are large

[7] Note that our assumption that the risk-free interest rate is zero implies that μ_p equals $p(t = 0)$.

("good" states of nature). For example, indirect bankruptcy costs at $t = 1$ (affecting revenues at time $t > 1$) could impact the hedging decision in a way that is well approximated by an exponentially declining cost function. Another example is a firm confronting costly access to external capital markets [Froot, Scharfstein, and Stein (1993)] where external financing costs increase exponentially in the amount of funds raised.

The parameter c_1 measures the overall level of deadweight costs, while c_2 controls the slope and curvature. If c_2 is small, the deadweight costs associated with low-profit states are only slightly larger than those incurred in good states of nature. Conversely, if c_2 is large, the deadweight costs in states with low profits are much larger than those in good states of nature. For example, given reasonable production cost and distributional parameters, values of $c_1 = 0.1$ and $c_2 = \{2, 5, 8\}$ result in expected deadweight costs of about $\{1.3\%, 2.9\%, 4.2\%\}$ of expected revenues if no hedging is undertaken.[8] Finally, note that $C_{PP}(P) > 0$ for all P.

We can now solve the firm's maximization problem for a specific set of derivative contracts. First, we determine the optimal hedge ratio for a firm that uses forwards as its only risk management tool. This specification admits a closed-form solution for the optimal hedge when the firm is risk neutral. We also discuss the optimal position for a firm that uses vanilla options or a portfolio of forwards and options as its hedge. Next, we derive the optimal custom "exotic" derivative by maximizing expected profits over the space of all possible derivative contracts.

1.3 Optimal vanilla hedges

Derivatives with linear payoffs, such as forward contracts, are by far the most popular financial instruments for risk management.[9] Consequently we start by assuming that $x(p; \mathbf{a})$ has a payoff of the form

$$x(p; a) = a(p - \mu_p) \qquad a \in R. \tag{8}$$

The parameter a is the number of forward contracts the firm buys (so $a < 0$ implies the firm sells forwards). By construction the cost of entering into the contract $X(a) = 0$. For $g(p) = g^*(p)$, the optimal number of forward contracts must satisfy the first-order condition[10]

$$\frac{\partial \pi(a)}{\partial a} = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} (p - \mu_p) e^{-c_2[pq - s_1 q - s_2 + a(p - \mu_p)]} h(p, q)\, dq\, dp = 0. \tag{9}$$

[8] We assume price and quantity volatilities of 20%, correlation between price and quantity of -0.5, expected price and quantity standardized to 1.0, variable production costs $s_1 = 0.25$, and fixed costs $s_2 = 0.4$. While the costs may appear small, note that they are a percent of expected revenues (not profits). For example, if a firm has $1 billion in foreign sales, these costs could range from $13 million to $42 million.

[9] Hentschel and Kothari (2001) report that nonfinancial firms in their sample (425 large U.S. firms during 1990–1993) hold an average notional value of $246 million in forwards and $347 million in swaps (which can be considered a sequence of forwards). These two derivative classes represent about 90% of the average firm's total derivative position.

[10] We have proven analytically that the solution to this first-order condition is indeed a maximum.

Appendix A shows that this first-order condition can be expressed in closed form as

$$\frac{\partial \pi(a)}{\partial a} = P_1(a)e^{P_2(a)} = 0, \tag{10}$$

where $P_1(a)$ and $P_2(a)$ are linear and quadratic functions of a defined as

$$P_1(a) = \frac{-c_1c_2^2\sigma_p^2 a + c_1c_2[c_2\mu_q\sigma_p^2 - (1-\rho^2)(s_1-\mu_p)c_2^2\sigma_q^2\sigma_p^2 + (s_1-\mu_p)c_2\rho\sigma_q\sigma_p]}{[1+2c_2\rho\sigma_q\sigma_p - (1-\rho^2)c_2^2\sigma_q^2\sigma_p^2]^{\frac{3}{2}}}, \tag{11}$$

$$P_2(a) = \frac{\left\{\begin{array}{c} \frac{1}{2}c_2^2\sigma_p^2a^2 + c_2[c_2\mu_q\sigma_p^2 - (1-\rho^2)(\mu_p-s_1)c_2^2\sigma_q^2\sigma_p^2 + (s_1-\mu_p)c_2\rho\sigma_q\sigma_p]a \\ +c_2\left[\mu_q(s_1-\mu_p)+s_2+c_2\left(\begin{array}{c}\frac{1}{2}\mu_q^2\sigma_p^2 + (\frac{1}{2}(s_1+\mu_p^2)-\mu_p s_1)\sigma_q^2 \\ +(\mu_q(s_1-\mu_p)+2s_2)\rho\sigma_q\sigma_p\end{array}\right)\right] \\ -c_2^2s_2(1-\rho^2)\sigma_q^2\sigma_p^2 \end{array}\right\}}{1+2c_2\rho\sigma_q\sigma_p-(1-\rho^2)c_2^2\sigma_q^2\sigma_p^2}. \tag{12}$$

Equation (10) reduces to the linear equation $P_1(a) = 0$, which we can invert into a closed-form solution for the optimal number of forward contracts

$$a^* = -\mu_q - \rho\frac{\sigma_q}{\sigma_p}(\mu_p - s_1) + c_2\sigma_q^2(1-\rho^2)(\mu_p - s_1). \tag{13}$$

The first term on the right-hand side shows that the firm should hedge its entire production with forwards when the produced quantity is known with certainty. We denote this the "naive" forward hedge. The second and third terms of Equation (13) are adjustments to the naive hedge that depend on firm-specific parameters. This solution shows that determining the optimal number of forward contracts is nontrivial. Optimal hedge ratios depend on the correlation between price and quantity, the volatility of price and quantity, the relative severity of costly states of nature, and variable costs of production.

This optimal hedge can be compared to the forward position that minimizes the variance of the firm's revenues. Under our distributional assumptions, the minimum variance hedge is given by

$$a^*_{mv} = -\mu_q - \rho\frac{\sigma_q}{\sigma_p}\mu_p. \tag{14}$$

The expected value maximizing and minimum variance hedges equal the naive hedge when there is no quantity risk ($\sigma_q = 0$). However, there exist important differences between these two hedges when the produced quantity is uncertain. The minimum variance hedge does not take into account production costs. In addition, the effect of deadweight costs is ignored.

If instead of using forwards, the firm restricts itself to using vanilla put options, an implicit equation for the optimal number of put options can be derived.[11] Utilizing this equation the optimal put option notional value and

[11] The derivation, not presented here, is available from the authors upon request.

strike price can be found numerically. More generally, it is possible to solve for the optimal hedging position using any combination of vanilla derivatives with numerical optimization techniques. (Subsequently we numerically integrate Equation (5) and employ a generalized reduced gradient method for finding optimal parameter values.) Numerical techniques also allow for alternative deadweight cost functions and distributional assumptions. As we show next, most of the intuition from the model can be gleaned from analyzing the optimal forward and exotic hedges, so we concentrate on these securities.

1.4 The perfect exotic hedge

While vanilla derivative contracts are popular hedging instruments, there is no reason to believe that using only a limited number of forwards and options yields the first-best hedge. Consequently we generalize the available derivative contract so that the firm can create any state-contingent payoff it chooses. We are thus solving for the "perfect" exotic hedge. This allows us to answer several interesting questions: What are the properties of the optimal derivative contract? How much better can a firm hedge if it uses customized exotic derivatives? Is it easy to approximate the optimal derivative by combining a few vanilla instruments, or is there a genuine need for truly exotic types of payoffs?

To answer these questions we need to modify the general model described in Section 1.1. Since we are interested in determining the optimal payoff function rather than the parameters of a prespecified payoff profile, we need to modify the optimization problem to maximize expected profit over all possible payoff functions $x(p)$, where x is a real function of price.[12] The general optimization problem for the single-product firm is

$$\max_{x(p)} \pi(x(p)) \qquad \text{s.t.} \qquad \int_{-\infty}^{\infty} x(p)g(p)\,dp = 0, \tag{15}$$

where

$$\pi(x(p)) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} pq - s_1 q - s_2 - c_1 e^{-c_2(pq - s_1 q - s_2 + x(p))} h(p,q)\,dp\,dq. \tag{16}$$

This functional optimization problem results in an optimal payoff function with a simple quadratic form

$$x(p) = \alpha_2 p^2 + \alpha_1 p + \alpha_0, \tag{17}$$

[12] This problem is mathematically analagous to the continuous time portfolio problem examined by Svensson and Werner (1993). Specifically, a similar problem would be obtained if we considered the case where $g(p) \neq g^*(p)$, and in the Svensson and Werner context, "labor income" is specified as $pq - s_1 q - s_2$ at $t = 1$. We thank the referee for pointing out this important relation.

where

$$\alpha_2 = -\rho\frac{\sigma_q}{\sigma_p} + \frac{1}{2}(1-\rho^2)c_2\sigma_q^2,$$

$$\alpha_1 = -\mu_q + (\mu_p + s_1)\rho\frac{\sigma_q}{\sigma_p} - (1-\rho^2)s_1c_2\sigma_q^2,$$

$$\alpha_0 = -\alpha_2\left(\sigma_p^2 + \mu_p^2\right) - \alpha_1\mu_p.$$

The parameters of this optimal payoff function allow us to determine exactly which firms should use nonlinear hedging tools.

While we have relegated the derivation of the optimal payoff function to Appendix B, outlining the proof highlights the genesis of the quadratic nature of the optimal payoff function. The derivation is done in two steps. First, we remove the dependence on the *nonhedgable quantity* variable by integration over q. This step effectively reduces the incomplete market problem with nonhedgable quantity risk to a complete market problem where we are left with an expected profit function that only depends on the *hedgable* price variable and the derivative payoff function $x(p)$ (which we desire to determine). The perfect hedge $x(p)$ must therefore undo all the remaining uncertainty in the expected profit resulting from hedgable price uncertainty. The exact form of $x(p)$ is determined in the second step, where we form the Lagrangian and solve the associated first-order conditions for all p and a Lagrange multiplier. The origin of the quadratic form of $x(p)$ becomes apparent when noting that the first-order conditions for $x(p)$ for all p [Equation (B.4)] can be reexpressed as

$$\lambda - 1 = c_1c_2e^{\alpha_2c_2p^2+\alpha_1c_2p+A-c_2x(p)}. \tag{18}$$

This equation shows that the only way the right-hand side can be a constant is if the term inside the exponential is a constant and, as such, $x(p)$ must be a quadratic function.

The slope of the payoff function for a given price measures the local exposure provided by the hedging program. For example, if the optimal payoff function is linear ($\alpha_2 = 0$), and the slope $\alpha_1 + 2\alpha_2 p = -1$, the exposure is identical to that obtained by selling a single forward contract. In general, the slope of the optimal payoff function equals

$$\begin{aligned}\frac{dx(p)}{dp} &= \left(-2\rho\frac{\sigma_q}{\sigma_p} + (1-\rho^2)c_2\sigma_q^2\right)p - \mu_q + (\mu_p + s_1)\rho\frac{\sigma_q}{\sigma_p} - (1-\rho^2)s_1c_2\sigma_q^2 \\ &= -\mu_q - \rho\frac{\sigma_q}{\sigma_p}(2p - \mu_p - s_1) + (1-\rho^2)c_2\sigma_q^2(p - s_1). \end{aligned} \tag{19}$$

Consider the slope of the optimal exotic at the point where $p = \mu_p$. At this price level the slope is identical to the hedge ratio of the optimal forward hedge [Equation (13)], indicating that the forward hedge approximates the optimal exotic hedge by creating identical exposures at the forward price, μ_p.

As prices deviate from μ_p, the absolute difference in slopes between the exotic and the forward hedge increases at a rate equal to the convexity of the exotic hedge. In this context, we define convexity as the second derivative of the payoff function,

$$\frac{d^2x(p)}{dp^2} = -2\rho\frac{\sigma_q}{\sigma_p} + (1-\rho^2)c_2\sigma_q^2. \tag{20}$$

The convexity of the optimal contract depends only on price and quantity risk, correlation, and the cost parameter, c_2. It is easy to see why the optimal hedge tends to be convex when price and quantity are negatively correlated. We know that all hedgable uncertainty is eliminated by the price hedge [see Equation (18)]. Because a negative correlation between price and quantity means that, conditional on price, the expected revenue is a concave function, the hedge must offset this function and as such the hedge must be convex. The opposite argument holds when price and quantity are positively correlated (ignoring the effect of the exponential cost function). Mathematically, if correlation satisfies

$$\rho < \frac{-1+\sqrt{1+(c_2\sigma_p\sigma_q)^2}}{c_2\sigma_p\sigma_q}, \tag{21}$$

the firm should buy convexity. Note that the right-hand side of Equation (21) is always positive.[13] This implies that a firm facing negative price-quantity correlation should always buy a convex hedge; this result is independent of all other parameters. Also, for negative correlation, the convexity of the optimal custom hedge increases in both quantity risk and the cost parameter c_2, while convexity decreases in price risk.

2. Comparison of Optimal Hedging Strategies

2.1 Setup

We now turn our attention to examining the solutions derived in the previous section. By carefully analyzing the factors that determine the optimal hedge parameters, it is possible to gain insight into which firm-specific attributes are most important for determining the optimal hedge's defining characteristics. We start with a base case of reasonable parameter estimates and then analyze

[13] This highlights another difference between our model and existing work. For example, Svensson and Werner (1993) find that if (nontraded) income is uncorrelated with rates of return on traded assets, the income hedge portfolio is also zero. Hence the nontraded income has no effect on the holdings of the traded assets. Our result shows that even when correlation is zero, the firms optimal hedging strategy is nontrivial.

the effect of perturbing these initial values:

- Expected quantity and price are normalized to 1.0, so expected total revenues will be approximately 1.0 (depending on the price-quantity correlation). This also allows for the interpretation of the optimal forward hedge parameter (a^*) as a hedge ratio.
- Base-case volatilities for price and quantity are 0.20, or 20%.
- The exponential deadweight cost function has level parameter $c_1 = 0.1$ and curvature parameter $c_2 = 5.0$. These result in an expected deadweight cost of about 3.4% of revenues (for zero correlation and other base-case parameters).
- Base-case fixed costs are 0.4 and variable costs are 0.25, so that expected profits, π, for the firm are about 0.3 (again, depending on the price-quantity correlation).

The remainder of this section contains four parts. The first part compares different hedging strategies and analyzes the determinants of the optimal convexity of the exotic hedge. Second, we look at implications for hedging at different time horizons. Next, solutions with alternative deadweight cost functions and log-normally distributed variables are examined. Finally, we provide an example of an application for the multiasset case.

2.2 Shape and efficiency of optimal hedges

Figure 1 plots optimal hedges for our set of base-case parameters with a correlation of -0.5. The dotted line plots the naive forward hedge and the lightest-weight solid line plots the optimal forward hedge ($a^* = -0.51$). The medium-weight solid line plots the net payoff to the optimal number ($a^* = 0.88$) of at-the-money put options,[14] while the heavy solid line plots the perfect exotic hedge. As indicated by Equation (19), the slope of the optimal forward hedge is identical to that of the optimal custom hedge when price equals μ_p ($p = \mu_p = 1.0$).

As expected, the optimal hedge tends to pay positive amounts in low-price states and negative amounts in high-price states. Since Figure 1 assumes a negative price-quantity correlation, the convexity of the optimal custom contract is quite pronounced. It is also evident that the optimal at-the-money put position is chosen so that it resembles the optimal exotic hedge, especially in low-price states where deadweight costs are more likely to be large.

We can now compare the degree of hedging effectiveness achieved from using forwards, options with a single strike, and the perfect exotic derivative security. We do this by constructing an efficiency rating that compares the decrease in expected deadweight costs from a hedge to the reduction in

[14] Formulas for the optimal position using put options are available from the authors upon request.

Panel A: Derivative payoff by price

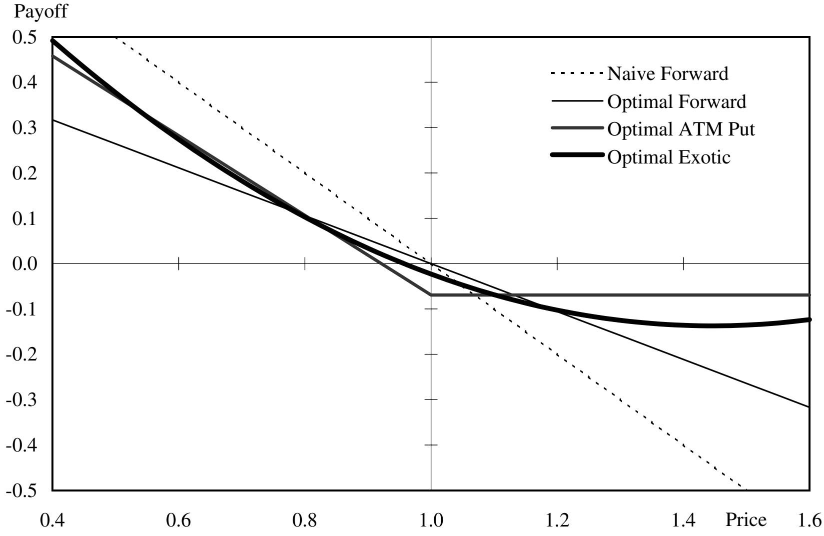

Panel B: Deadweight costs and efficiency

Hedge	Hedge ratio	E[DWC]	Efficiency
No hedge	0.00	2.959%	0.0%
Naive forward	−1.00	2.912%	9.5%
Optimal forward	−0.51	2.495%	92.3%
Optimal ATM put option	0.88	2.482%	94.9%
Optimal exotic	NA	2.456%	100.0%

Figure 1
Optimal hedging with an exponential deadweight cost function (negative correlation)
Expected price and quantity are standardized to 1.0 with volatilities of 20%, and the correlation between price and quantity equals −0.5. We assume that variable production costs are $s_1 = 0.25$ and fixed costs are $s_2 = 0.4$. The exponential cost function has parameters $c_1 = 0.1$ and $c_2 = 5.0$. Panel A graphs payoffs to the naive forward hedge, the optimal forward hedge, the optimal at-the-money (ATM) put option hedge, and the optimal exotic hedge as a function of the realized price level. Panel B tabulates the hedge ratio, expected deadweight cost as a percentage of expected revenues (E[DWC]), and relative efficiency for the various hedging scenarios. Relative efficiency of a hedge measures the expected reduction in deadweight costs relative to the largest possible reduction achieved from implementing the optimal exotic hedge.

expected costs obtained from the optimal custom hedge. The efficiency of hedge i, ζ_i, is defined as

$$\zeta_i = \frac{E[C_{NH}] - E[C_i]}{E[C_{NH}] - E[C_{OPT}]}, \tag{22}$$

where $E[C_i]$, $E[C_{OPT}]$, and $E[C_{NH}]$ are the expected deadweight cost of hedge i, the optimal custom hedge, and no hedge, respectively. By construction, the maximum efficiency rating any hedge can obtain is 100%, and an efficiency less than 0% indicates that the hedge in question results in higher expected deadweight costs than those expected from not hedging at all. Panel B of Figure 1 shows the expected deadweight costs and efficiency ratings of various hedges. The naive hedge has an efficiency $\zeta_{Naïve} = 9.5\%$, implying that simply selling expected production forward is an ineffective hedge. In contrast, the optimal forward hedge performs relatively well, reducing the hedgable deadweight costs by 92.3%. The hedge can be improved further by using the at-the-money put options ($\zeta_{Put} = 94.9\%$).

As noted already, firms are not usually limited to using forwards or options, and in most circumstances may create a hedge that consists of both. If a simple portfolio of forwards and put options with a high efficiency can be constructed in place of a highly customized instrument, this may be preferred by the firm due to higher liquidity, lower monitoring and transactions costs, or other institutional reasons. We can approximate the optimal exotic in panel A by letting the firm enter into both forward contracts and put options with a single strike. An almost optimal hedge is obtained by selling 0.26 forwards and buying 0.52 puts with a strike price of 0.99.[15] This portfolio results in an efficiency rating of 99.1%. Therefore, depending on available contracts, the firm may be able to construct a highly efficient hedge using forwards and options with a limited number of striking prices. This is consistent with empirical evidence indicating that the majority of firms use both linear contracts and options in their hedging programs [see Tufano (1996) and Bodnar, Hyat, and Marston (1998)].

Figure 2 shows the results of the same optimization problem discussed above, but with a positive correlation of 0.5. Panel A plots the naive and optimal forward hedges, the optimal at-the-money put option hedge, and the optimal custom hedge. The most striking difference from the negative correlation case is the shape of the optimal custom hedge. In contrast to the previous example, the custom hedge is concave. This follows directly from Equation (20), which shows that the convexity of the optimal hedge is negative for these parameters. However, note that the magnitude of the convexity is less in the positive correlation case, suggesting that linear hedges are relatively more efficient. As indicated by Equation (19) the custom payoff is more steeply sloped near $p = \mu_p$ than in the negative correlation case. In general, the firm should hedge considerably more in most price states to compensate for the increased exposure associated with the positive price-quantity correlation.

In the case of positive correlation, gains from hedging are much greater. Panel B of Figure 2 shows that expected deadweight costs can be reduced from 3.84% when the firm does not hedge to 1.93% when the firm hedges optimally. The approximately linear custom hedge explains the high efficiency (99.3%) of the optimal forward hedge. Likewise, because the naive forward hedge is fairly close to the optimal forward hedge, it is also an

[15] We solve this problem numerically as described in Section 1.3.

Panel A: Derivative payoff by price

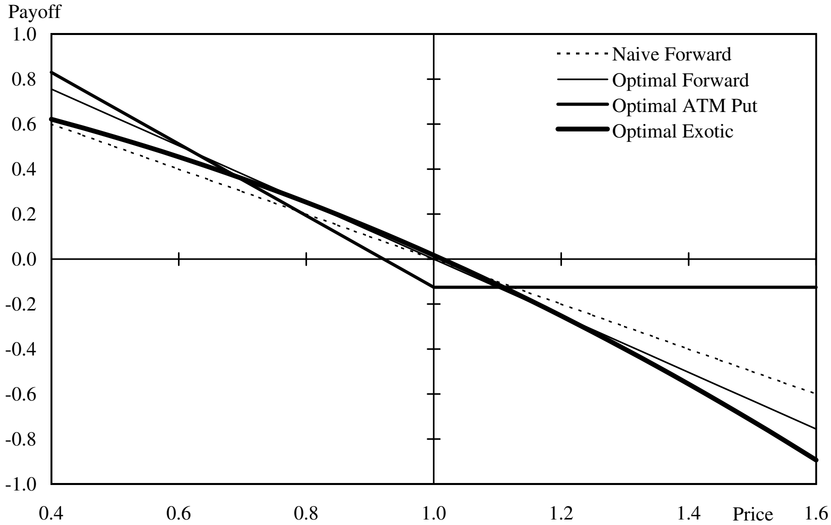

Panel B: Deadweight costs and efficiency

Hedge	Hedge ratio	E[DWC]	Efficiency
No hedge	0.00	3.843%	0.0%
Naive forward	−1.00	2.003%	96.3%
Optimal forward	−1.26	1.945%	99.3%
Optimal ATM put option	1.60	2.344%	78.5%
Optimal exotic	NA	1.932%	100.0%

Figure 2
Optimal hedging with an exponential deadweight cost function (positive correlation)
Expected price and quantity are standardized to 1.0 with volatilities of 20%, and the correlation between price and quantity equals 0.5. We assume that variable production costs are $s_1 = 0.25$ and fixed costs are $s_2 = 0.4$. The exponential cost function has parameters $c_1 = 0.1$ and $c_2 = 5.0$. Panel A graphs payoffs to the naive forward hedge, the optimal forward hedge, the optimal at-the-money (ATM) put option hedge, and the optimal exotic hedge as a function of the realized price level. Panel B tabulates the hedge ratio, expected deadweight cost as a percentage of expected revenues (E[DWC]), and relative efficiency for the various hedging scenarios. Relative efficiency of a hedge measures the expected reduction in deadweight costs relative to the largest possible reduction achieved from implementing the optimal exotic hedge.

efficient approximation to the optimal custom hedge (96.3%). In contrast, the optimal at-the-money put option hedge is only 78.5% efficient. Panel A reveals why put options are inferior to forwards in this setting: a long position in put options has a convex payoff while the optimal custom hedge is concave. We find this last result particularly interesting given that deadweight costs are a convex function of profits.

As noted, the convexity of the optimal exotic hedge depends only on correlation, price and quantity volatilities, and the shape of the deadweight cost function (c_2). Figure 3 shows how reasonable variations in these parameters

Panel A: Various deadweight costs

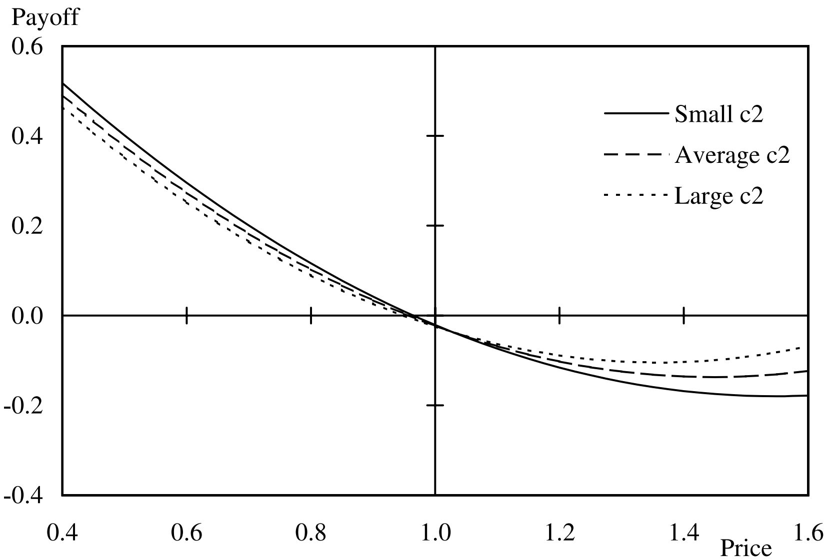

Panel B: Various quantity risks

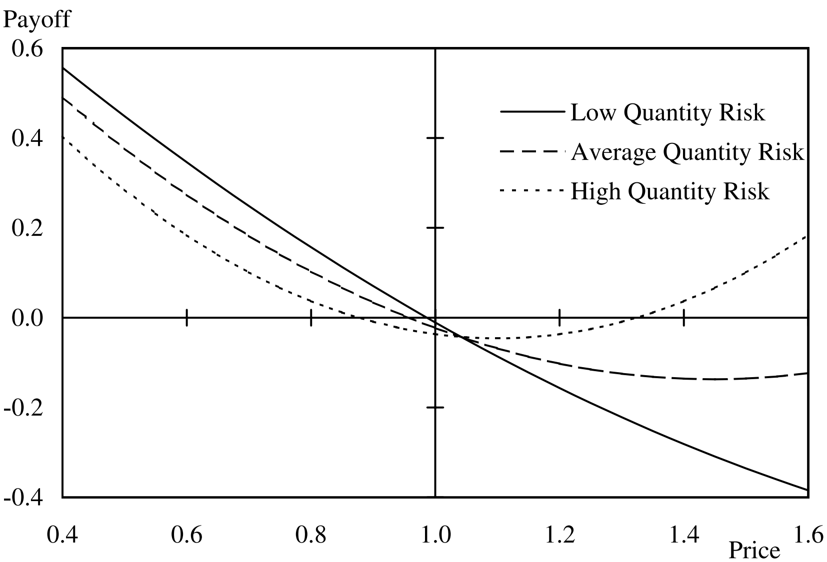

Panel C: Various price risks

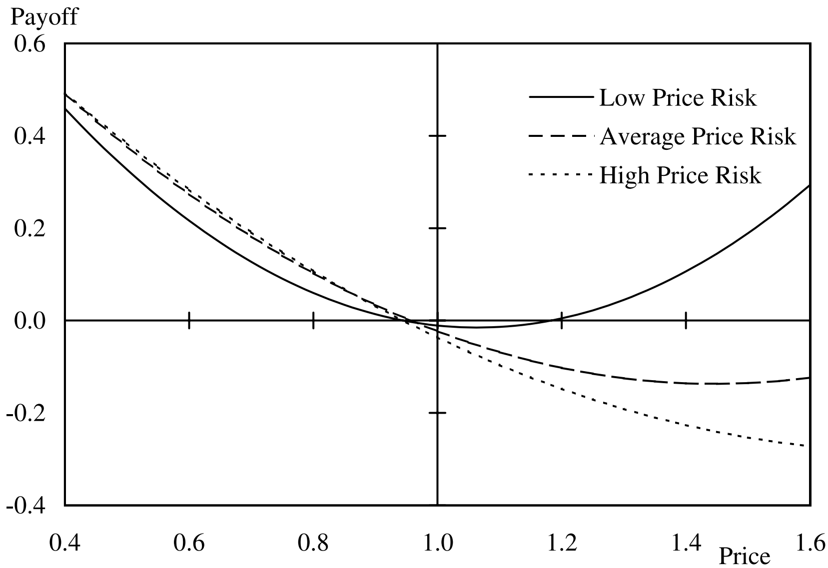

Panel D: Efficiency of forward hedge for different scenarios

Scenario	Efficiency of forward hedge
Panel A: Deadweight costs	
Small c_2	94.0%
Average c_2	92.3%
Large c_2	90.8%
Panel B: Quantity risk	
Low	99.3%
Average	92.3%
High	38.1%
Panel C: Price risk	
Low	46.0%
Average	92.3%
High	96.0%

Figure 3
Optimal custom hedges for various deadweight costs and volatilities
This figure shows optimal exotic hedges for different deadweight cost parameters and levels of price and quantity volatility holding other parameters fixed. As a base case, we determine the optimal hedge when expected price and quantity are both 1.0 and the correlation between price and quantity is -0.5. We assume that production costs are $\{s_1, s_2\} = \{0.25, 0.4\}$, and the deadweight cost function has parameters $\{c_1, c_2\} = \{0.1, 5.0\}$. Panel A graphs the optimal hedge for small (2), average (5), and large (8) values of c_2. Panel B graphs the optimal exotic hedge for low (10%), average (20%), and high (30%) quantity volatility. Panel C graphs the optimal exotic hedge for low (10%), average (20%), and high (30%) price volatility. Panel D tabulates the relative efficiency of the optimal forward hedge for each of the cases presented in panels A, B, and C.

affects the optimal payoff function. Panel A plots exotic hedges for $c_2 = \{2, 5, 8\}$. With the exception of the payoff in relatively high-price states, large variation in the convexity of the deadweight cost function has little effect on the optimal custom hedge. This suggests that firms may not need to know the exact parameterization of their deadweight costs to benefit substantially from hedging. (Optimal exotic hedges with alternative deadweight cost specifications are discussed below.)

All of the previous examples have used a volatility of 20% for both price and quantity. Panel B of Figure 3 shows optimal custom hedges for low (10%), average (20%), and high (30%) quantity risk. Low quantity risk results in an optimal hedge that is nearly linear for all relevant price states. This is an intuitive result since when quantity risk is zero, markets are complete and the optimal hedge is linear. When quantity risk is large, convexity becomes very important. In fact, with high quantity risk, the payoff of the optimal custom contract resembles that of a straddle or strangle. This is confirmed by examining Equation (20), which shows that optimal convexity depends quadratically on quantity risk.

Panel C illustrates the effect of price risk when correlation is negative. For high price volatility, the optimal hedge is well approximated by a linear hedge. As price risk decreases, the optimal custom hedge becomes more convex. As in the previous example, a strangle-type option position can be a good approximation to the optimal custom hedge. The intuition is derived from the ability to hedge price risk: As price risk increases relative to quantity risk, the hedge becomes more linear, since the unhedgable risk is less important. But as price risk decreases, the unhedgable component of total risk increases and the convexity of the optimal hedge increases. Equation (20) confirms this; when correlation is negative, a decrease in price risk increases the convexity of the optimal hedge.

Finally, Panel D reports the efficiency of optimal forward hedges for each of the cases shown in Figure 3. As is evident from the plots in panels B and C, the efficiency of the optimal forward hedge is high when quantity risk is low and when price risk is high. However, when the quantity risk is substantial, or the price risk is small, the efficiency of the best linear hedge drops dramatically. This occurs because the optimal forward hedge ratio is small. Note that this does not imply that only small benefits can be achieved from hedging or that the magnitude of the optimal hedge is small in all price states. It only implies that in these cases nonlinear hedges are essential for effective risk management.

Unlike the minimum variance hedge in Equation (14), our results indicate that production costs, or equivalently, operating leverage, can have a substantial effect on the optimal hedge. Figure 4 illustrates this by plotting the optimal forward hedge as a function of correlation for different variable costs, s_1.[16] If the firm's variable costs are negligible and the expected profit margin

[16] Recall that with our assumptions convexity is independent of operating costs.

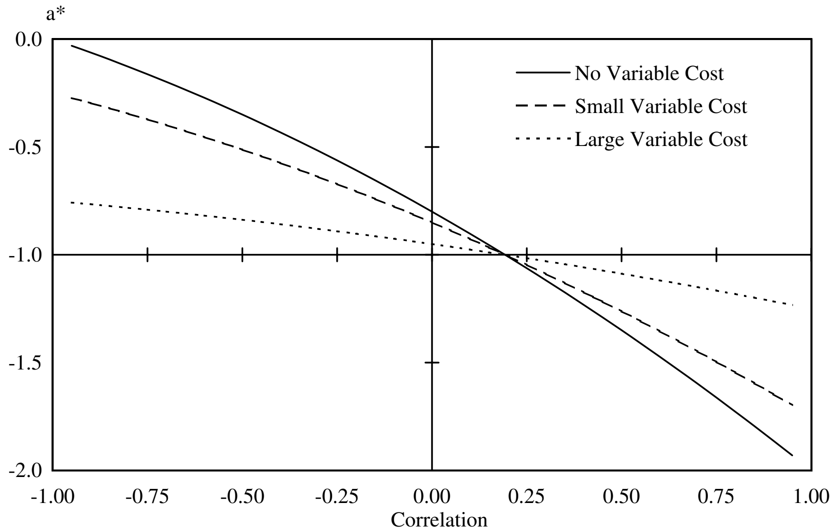

Figure 4
Optimal hedging and operating leverage
This figure plots the optimal number of forward contracts (a*) that the firm should buy for different levels of variable costs. Correlation between price and quantity is plotted on the horizontal axis and the forward hedge ratio is plotted on the vertical axis. The axes cross at zero correlation and $a^* = -1.0$ (the naive forward hedge). For the base case, expected price and quantity are both 1.0 with volatilities of 20%. We assume that variable production costs are $s_1 = 0.25$ and fixed costs are $s_2 = 0.4$. Variable costs of 0.0, 0.25, and 0.75 correspond to no, small, and large variable cost scenarios, respectively. Fixed costs are set to keep expected profits identical across scenarios.

is large (solid line), then the optimal hedge is very sensitive to changes in correlation. On the other hand, if the variable costs equal the expected price (i.e., expected contribution margin, $\mu_p - s_1$, is zero), the optimal hedge is identical to the naive hedge [by inspection of Equation (13)]. At first glance this may seem counterintuitive because the expected profit from operations is zero, yet the firm remains fully exposed to price fluctuations that are most efficiently managed using the naive hedge. The optimal hedge does not depend on the price-quantity correlation when $s_1 = \mu_p$. This happens because the benefits (a marginal reduction in variability) associated with "quantity hedging" in price states with negative marginal profits ($p < s_1$) are exactly offset by the costs (a marginal increase in variability) associated with "quantity hedging" in price states with positive marginal profits ($p > s_1$).[17]

The sensitivity of the optimal hedge to changes in variable costs provides another interesting insight. It is generally believed that less profitable firms have a greater incentive to hedge than highly profitable firms because they are more likely to reach states of nature with significant deadweight costs. For

[17] This result follows from the exponential nature of the deadweight cost function. It does not necessarily hold for more general cost functions. However, the qualitative comparative static is the same.

example, Tufano (1996) argues that gold mining firms with high extraction costs (cash costs) should hedge more than firms with low extraction costs. Our model indicates that this argument is only correct when the correlation between price and quantity is sufficiently small. When

$$\rho > -\frac{1}{2c_2\sigma_p\sigma_q} + \sqrt{1 + \frac{1}{4(c_2\sigma_p\sigma_q)^2}} \tag{23}$$

the optimal hedge decreases (a^* increases) when variable costs increase. Furthermore, when

$$\rho = -\frac{1}{2c_2\sigma_p\sigma_q} + \sqrt{1 + \frac{1}{4(c_2\sigma_p\sigma_q)^2}} \tag{24}$$

the optimal hedge is independent of the variable costs s_1. In Figure 4 this occurs near $\rho = 0.2$. This scenario is then consistent with Tufano (1996), who finds that hedging activity in the gold mining industry is statistically unrelated to extraction costs.[18]

Finally, we note that Equation (13) does not include the fixed cost parameter, s_2. This parameter does not affect the optimal hedge when $g(p) = g^*(p)$ because it cancels in the first-order condition, Equation (10). Intuitively the firm only hedges the relative costliness of states and since the fixed costs are by definition identical in all states, different levels of fixed costs do not result in different optimal hedges.[19]

Overall, this section has shown that the most important factors in constructing optimal hedges are the correlation between price and quantity and the respective volatilities of price and quantity risk. Production technology (i.e., operating leverage) is important for determining the size, but not the convexity, of hedges. A surprising result is that variation in the curvature of the deadweight cost function is relatively unimportant. Furthermore, our analysis indicates that heuristic hedging strategies such as "hedge expected output" or "minimize the variance of revenues" are inferior to the value maximizing strategy. Finally, and in contrast to some implications of related research, optimal hedge ratios depend crucially on the operating characteristics of the firm as well as prevailing market conditions.

2.3 Hedge horizon

A simple extension to our model makes it possible to analyze the hedging horizon's impact on the optimal exotic payoff function. If we assume that price and quantity uncertainty are well described by Brownian motions, then

[18] Note that it is reasonable to assume that mining firms increase their production quantities when the price of the produced commodity increases. Tufano (1996) finds weak evidence of this in the gold mining industry.

[19] Again, this finding follows from the exponential nature of the cost function $C(P)$.

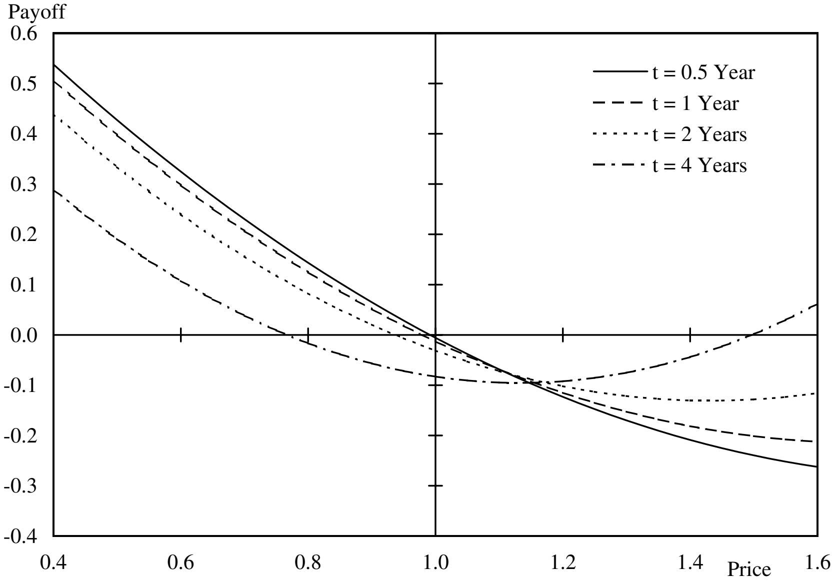

Figure 5
Optimal exotic hedges for different time horizons
This figure shows optimal exotic hedges for different hedging time horizons, holding other parameters fixed. As a base case, we determine the optimal hedge when expected price and quantity are both 1.0 and the correlation between price and quantity is -0.5. We assume that production costs are $\{s_1, s_2\} = \{0.25, 0.4\}$, and the deadweight cost function has parameters $\{c_1, c_2\} = \{0.1, 5.0\}$. Price and quantity volatilities are specified as a function of time. Specifically price and quantity volatility are $0.20t^{0.5}$.

we can model different hedging horizons by multiplying all (annual) volatilities by $\sqrt{t}$ (where t is measured in years).

Figure 5 shows optimal exotic hedges for the negative correlation case when the hedging horizon varies from two quarters to four years. There are two important effects to consider. First, note that the hedges “rotate” counterclockwise as the horizon lengthens, or equivalently the notional values of the optimal forward hedges decrease. Simply put, the firm hedges less when the exposure it is hedging is farther in the future. This can be observed analytically by referring to Equation (19). Multiplying the volatility terms by $\sqrt{t}$ results in a decrease in hedging whenever the slope of the optimal hedge is negative, since the new factors cancel in the second term.

The second effect is also obvious from Figure 5; the optimal exotic hedge becomes more convex as the hedging time horizon increases. This implies that similar firms with long hedging horizons should use more options than those with short hedging horizons. The convexity effect can also be verified analytically, in this case by inspecting Equation (20). Multiplying the volatility terms by $\sqrt{t}$ results in an increase in convexity. When correlation is negative, this will always lead to an increase in the optimal “optionality” of the hedge. In cases where the optimal hedge is concave, the optimal hedge will become less concave and eventually convex as the time horizon increases.[20] The basic conclusion of this analysis is intuitively satisfying and consistent with empirical evidence suggesting limited long-term hedging by nonfinancial firms [see, for example, Bodnar, Hyat, and Marston (1998) or Hentschel and Kothari (2001)]. Firms may be reluctant to lock in a large hedge for the distant future because of the difficulty in making accurate exposure (quantity) forecasts. Likewise, what hedge they do undertake will more likely be composed of long positions in options to prevent large payouts that could be possible with linear instruments [see Mello and Parsons (2000)].

2.4 Alternative cost functions and price-quantity distributions

One potential shortcoming of our closed-form solutions is the reliance on an exponential deadweight cost function. While some theoretical deadweight costs may be well approximated by an exponential function (e.g., costly external capital in low-profit states, indirect bankruptcy costs, etc.), others may not (e.g., taxes, costly external capital for a firm with large investment opportunities in high-profit states, a fixed cost to distress, etc.). To determine how sensitive the optimal hedging strategy is to changes in the deadweight cost, we replace the exponential cost function with three alternative deadweight cost specifications and then solve for the optimal exotic hedge numerically.[21]

First we consider a firm that can experience a fixed cost to distress in low-profit states.[22] Panel A of Figure 6 shows that there is not much difference between the optimal custom hedge for this cost function and that of the exponential cost case (the same hedge as shown in Figure 1). Again, the payoff looks similar to that of a portfolio of put options. The optimal exotic for this cost scenario is slightly less convex than for the exponential cost

[20] This raises the interesting possibility of firms optimally having both long and short positions in options of different maturities. At a minimum it suggests the need for further research into multiperiod derivative choice models.

[21] More specifically, we descretize the state space and solve an optimization problem that approximates a continuous derivative payoff function. We approximate the state space using a grid of equally spaced price and quantity states. We assume that price and quantity are jointly normally distributed, but we use a discrete approximation to the density. We allow the firm to transact in Arrow–Debreu securities (that pay $1 if and only if a given discrete price state occurs at $t = 1$). The firm can buy or sell these securities with payoff states equal all possible price states but is subject to the zero initial cost constraint.

[22] We approximate bankruptcy costs with a smooth hyperbolic tangent function to facilitate the numerical analysis. Specifically we let $C(P) = \gamma_1(\tanh(\gamma_2 P + \gamma_3) - E[P])$, where γ_1 determines the maximum level of the cost, γ_2 determines the steepness of the function near the bankruptcy trigger, and γ_3 sets the bankruptcy trigger. $E[P]$ is expected profits, and for our example, $\{\gamma_1, \gamma_2, \gamma_3\} = \{0.1, 20, 0.1\}$. These parameters imply a fixed distress cost of 0.1 for negative profits. Because the deadweight cost function is not strictly convex, a pathological solution exists where the firm wishes to sell an infinite amount of payoff in the highest price state financing positive payoffs in other price states. Since we are unaware of a counterparty that would buy such a contract, we restrict the exotic payoff function in this section so that the maximum (minimum) payoff in any price state is 2.0 (-2.0). This leads to the solutions shown.

function. Hedging is still effective at reducing expected deadweight costs; expected costs are reduced by about 58%.

Deadweight costs may also occur in "good" states of nature. For example, if a firm faces a convex tax schedule based on profits, higher profits will lead to higher tax rates [see Smith and Stulz (1985) and Graham and Smith (1999)]. To model this cost structure we specify the deadweight cost function equal to a tax rate of 35% on profits beyond 0.3, that is, $C(P) = 0.35\max(P - 0.3, 0)$. Panel A shows that, again, the optimal custom derivative is convex in price and qualitatively similar to the optimal exotic for the exponential cost function. Expected deadweight costs are reduced by roughly 23%, from 0.0312 to 0.0240. As a final example, we assume a deadweight cost function that is a composite of exponential costs, fixed bankruptcy costs, and the corporate tax schedule described above. Panel B shows that the payoff to the optimal custom derivative still resembles the derivative payoff in the much simpler exponential case.[23]

There is an intuitive explanation for why the optimal exotic is relatively insensitive to the specific cost function. The deadweight cost as a function of price is an average of costs across all quantity states. This has the effect of "smoothing" irregular cost functions. In the case of the fixed deadweight costs, the resulting function (of price) looks very much like the exponential function in high probability states. This reinforces the logic presented in the previous sections, where we identified correlation and volatility as the primary determinants of the optimal exotic hedge's shape. In other words, the qualitative features of the optimal hedge are primarily determined by the firm's other characteristics.

This facet of our model has considerable practical ramifications. It implies that firms do not need to know their deadweight cost function exactly to benefit significantly from a hedging program. For example, suppose a firm mistakenly believes it faces an exponential deadweight cost function, when it actually confronts an unknown composite deadweight cost function. As a consequence, the firm undertakes the wrong optimal hedge (the hedge shown in Figure 1 instead of the "combination" hedge in Figure 6). The efficiency of this wrong optimal custom hedge is nevertheless 98.8%, much better than the naive forward hedge (44.9%) and even greater than the correct optimal forward hedge (95.0%).

Another assumption of the closed-form solutions is the joint normal density of price and quantity. As a consequence, we determine our results' sensitivity to changes in distributional assumptions. As an example, we repeat the numerical procedure described above for the exponential cost function and log-normally distributed prices and quantities. Panel B of Figure 6 shows the optimal exotic hedges that result from this procedure for three different

[23] Experiments comparing the optimal derivative for different cost functions, but using other parameter values (e.g., ρ, σ_p, σ_q), yielded similar conclusions. These results, not presented here, are available upon request.

Panel A: Optimal hedges with alternative deadweight cost functions

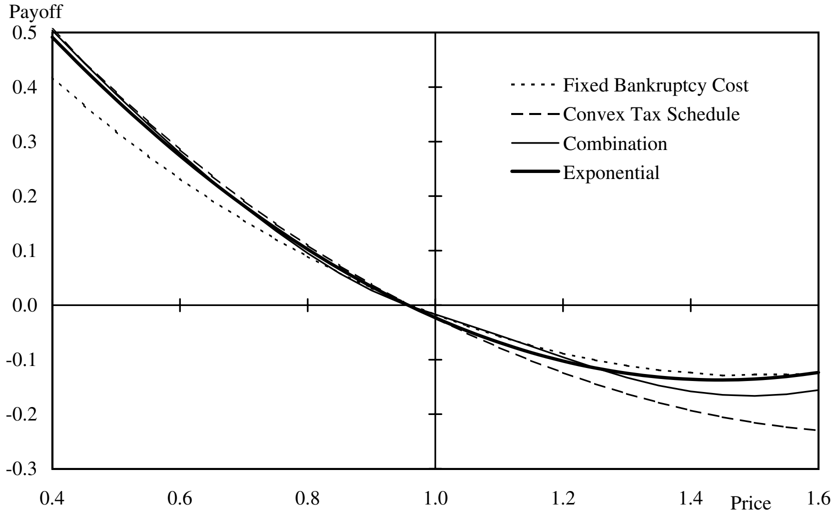

Panel B: Optimal exotic hedges for lognormal distribution

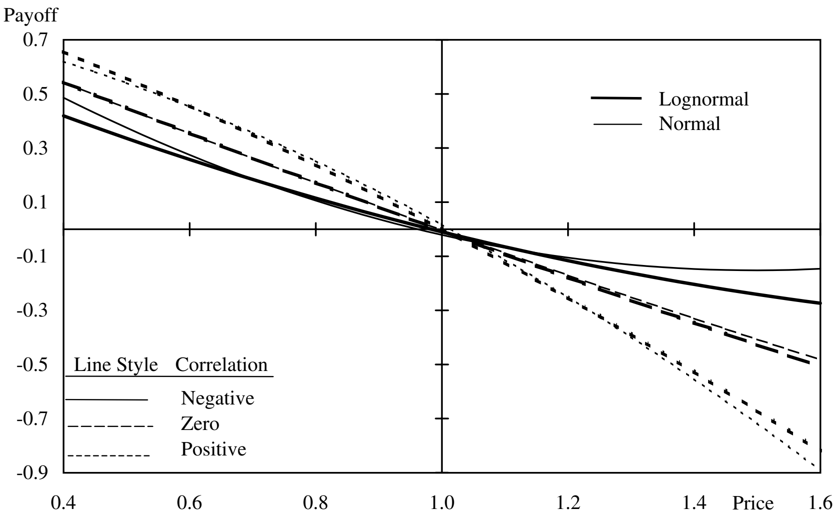

Figure 6
Robustness of results
This figure shows optimal exotic hedges for alternative deadweight cost functions (panel A) and lognormally distributed price and quantity (panel B). In panel A, expected price and quantity are both 1.0, volatilities are 20%, and the correlation between price and quantity is -0.5. We assume that variable production costs are $s_1 = 0.25$ and fixed production costs are $s_2 = 0.4$. The panel shows the optimal exotic hedges associated with a fixed deadweight cost, a convex tax schedule, and a combination of exponential, fixed, and tax-related deadweight cost functions. The optimal hedge using only the exponential deadweight cost is shown as a reference. Panel B shows optimal exotic hedges when price and quantity are jointly lognormally distributed as compared to jointly normally distributed. Parameter values are as in panel A, with the exponential deadweight cost function having parameters $\{c_1, c_2\} = \{0.1, 5.0\}$. This figure plots cases with positive (0.5), zero, and negative (-0.5) correlation.

correlations between price and quantity. The other parameters of the problem are held constant (and are the same as those used in Figures 1 and 2). First, consider the negative correlation case (solid lines). The darker solid line shows the optimal exotic hedge when prices and quantities are jointly log-normal, while the lighter solid line plots the optimal exotic hedge as shown in Figure 1 for comparison (the normal case). First, note that the log-normal and normal hedges are qualitatively very similar. In most price states the magnitudes of the two hedges are approximately equal. The largest deviations occur in high-price states, where the normal hedge is notably less in magnitude. Nevertheless, the convexity of the log-normal hedge is positive, as is the convexity of the normal hedge. It is safe to say, however, that the convexity of the log-normal hedge is less than that of the normal hedge.

These traits also hold for the zero and positive correlation cases. In each case the approximate magnitudes (notional values) of the normal and log-normal hedges are similar. Likewise, for the positive correlation case, convexity is clearly negative for both hedges, and nearly linear hedges are optimal when correlation is zero. From the figure (and confirmed by other cases) it appears that the magnitude of convexity for the log-normal case is always less than for the normal case. The most important conclusion to be drawn from this analysis is that it is not our specific distributional assumption that is driving the qualitative features of our results. The intuition concerning the magnitude and convexity of the optimal hedge holds for the log-normal case as well as for the normal case.

2.5 Multiple price risks

The analysis so far has focussed on a firm with a single product sold in a single market. While this model may suit a sole-commodity producer such as a gold mining firm or a farmer, it will not accurately describe the exposures of a large multinational corporation. For example, if a domestic manufacturer sells its products in one foreign market, then it probably sells them in other foreign markets as well. This gives rise to multiple, correlated currency risks. Alternatively, a firm may have random input prices as well as random output prices. In this section we briefly discuss each of these cases.

First, consider the case of a multinational manufacturer. The firm produces a single good domestically and sells it in two foreign countries. For simplicity, assume that revenues in each country are perfectly correlated (with mean 1.0 and standard deviation 0.2). However, correlations between the two exchange rates (price 1 and price 2) and the level of foreign revenues are different. Specifically we assume a correlation between revenues and price 1 of -0.5 and a correlation between revenues and price 2 of 0.5.[24] Finally, the correlation between price 1 and price 2 is 0.5. The problem is closed by

[24] We pick opposite correlations to illustrate the generalization to multiple dimensions.

specifying reasonable deadweight cost parameters, $\{c_1, c_2\} = \{0.1, 5.0\}$, and production cost parameters, $\{s_{11}, s_{12}, s_2\} = \{0.25, 0.25, 0.4\}$.

The firm is allowed to contract in a zero-price two-asset quadratic derivative with a net payoff function of

$$X(p_1, p_2) = \kappa_{12} p_1^2 + \kappa_{11} p_1 + \kappa_x p_2 p_1 + \kappa_{22} p_2^2 + \kappa_{21} p_2 + \kappa_0. \tag{25}$$

We solve the firms hedging problem numerically, as described in Section 1.3, to find the profit maximizing set of parameters $\boldsymbol{\kappa} = \{\kappa_{12}, \kappa_{11}, \kappa_x, \kappa_{22}, \kappa_{21}, \kappa_0\}$.

An interesting facet of this problem is the addition of the term allowing the firm to hedge the product of the two prices ($\kappa_x p_2 p_1$). We refer to this as a "cross-hedge," since (in this case) the firm will be hedging the foreign exchange cross-rate. The need for cross-hedging adds a new dimension to the firms optimal hedging strategy, since it may no longer be possible for the firm to closely approximate the optimal exotic hedge with a portfolio of vanilla instruments—the exotic hedge is now *truly* exotic. This follows from the "quanto" nature of the cross-hedging term. If κ_x is considerably different from zero, then for the firm to effectively hedge it must construct a derivative that pays off in domestic currency an amount that depends on the product of two exchange rates. Hence the payoff resembles a "quanto-product forward" contract.

Panel A of Figure 7 shows the optimal hedging strategy for this example. Note the shape of the payoff function in each of the separate price dimensions. A slice in the price 1 dimension reveals a convex hedge, as might be predicted from the negative correlation between price 1 and foreign revenues. Likewise, a slice in the price 2 dimension reveals a concave hedge, as might be predicted from the positive correlation between price 2 and foreign revenues. Less obvious is the significant cross-hedging component.

To examine the importance of the cross-hedging term, we consider alternative strategies for a firm that does not desire to (or cannot) enter into a cross-hedge. One strategy would be to treat each exposure separately using the one-asset analysis in Section 1.4. This would mean applying Equation (17) to price 1 and price 2 separately. Panel C of Figure 7 shows that this would be only a partially effective method of hedging (efficiency of 46.4%). Another strategy would be to undertake the optimization described above, but simply omit any cross-hedging instruments from the final hedge portfolio. The third line of panel C shows this could be a costly mistake, resulting in a large *negative* efficiency rating. A final, and most preferred, strategy is to solve the optimization problem again while omitting the cross-hedging term from the payoff function. This reoptimization procedure results in a substantially improved efficiency rating of 81.9%. These results indicate that cross-hedging can have an important effect on the effectiveness of a hedge.

A contrasting case of multiple price risks is a firm that is exposed to both input and output price risks. For example, an oil refiner is exposed to the price

Panel A: Optimal hedging strategy for firm with revenue in 2 currencies

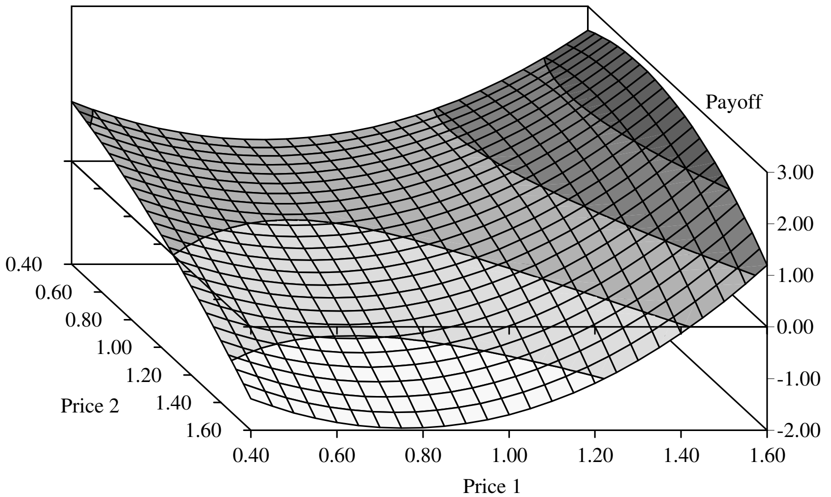

Panel B: Optimal hedging strategies for an oil refiner

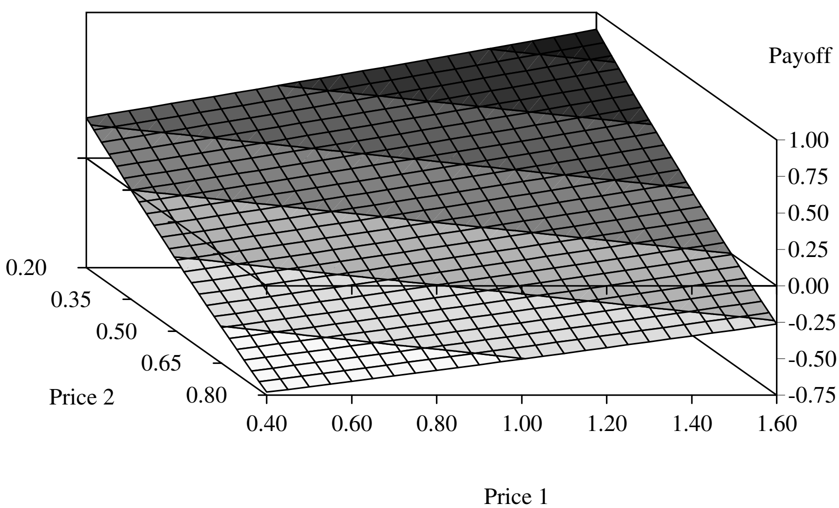

Panel C: Efficiency of alternative hedges

	Panel A		Panel B	
Hedging strategy	E[DWC]	Efficiency	E[DWC]	Efficiency
No hedge	0.076	0.0%	0.053	0.0%
No cross hedge (assume 1-D parameters)	0.066	46.4%	0.028	98.6%
No cross hedge (ignore cross terms)	0.112	−166.9%	0.029	95.6%
No cross hedge (reoptimize)	0.058	81.9%	0.028	99.9%
Optimal quadratic hedge	0.054	100.0%	0.028	100.0%

Figure 7
Hedging with multiple price risks
Panel A shows the payoff of the optimal hedge with two random output prices. The correlations between price and the random quantity are -0.5 for price 1 and 0.5 for price 2. The correlation between prices is 0.5. Expected prices and quantities are 1.0 with volatilities of 20%. Variable costs are 0.25 and fixed costs are 1.0. Panel B shows the payoff of the optimal hedge where the correlation between the (input) price 1 and the random quantity is -0.2 and the correlation between the (output) price 2 and the random quantity is -0.2. The correlation between prices is 0.9. The expected input price is 0.5, the expected output price is 1.0, and the expected quantity is 1.0 (all with volatilities of 20%). Fixed costs are 0.2, and the deadweight cost parameters are $\{c_1, c_2\} = \{0.1, 5.0\}$. Panel C shows expected deadweight costs (E[DWC]) and efficiency ratings for approximate cross-hedging strategies for the hedging scenarios in panels A and B. The second row, labeled "No cross hedge (assume 1-D parameters)," reports values for a hedge where each price risk is treated separately using the analysis in Section 2. The third row, labeled "No cross hedge (ignore cross term)," reports values for a hedge that was optimized with the cross-term, but the hedge omits the cross-term. The fourth row, labeled "No cross hedge (reoptimize)," reports values for a hedge that was optimized after omitting the cross-term.

of crude oil as a major input and the price of gasoline as a major output. The correlation structure of these price risks differs substantially from the case of the multinational manufacturer. Specifically one would expect that the correlation between the price of crude oil and the price of gasoline is very high. We assume it to be 0.9. One would also expect that there is a negative correlation between the price of crude oil (or gasoline) and the quantity consumed. We assume correlations of -0.2 between the random quantity and each of the prices.[25]

Panel B of Figure 7 shows the optimal hedging strategy for the oil refiner. In this case, the optimal hedge more closely resembles a plane. As one would expect from the small correlations between prices and quantity, the curvature of the surface is minimal. In addition, the strong positive correlation between the input price and output price provides a good natural hedge for the firm and in the process reduces the need for cross-hedging. As shown in panel C, and in contrast to the previous case, each of the approximate hedging strategies that do not involve cross-hedging is a highly efficient hedge.

The results of these two cases illustrate that the hedging problem for a firm with multiple risks is also highly dependent on the market factors specific to the firm. Finally, we note that for the two cases presented here, we assumed a single uncertain quantity. In practice, firms often have several uncertain

[25] Again we assume perfect correlation between the two quantities. We also assume that the expected price of gasoline is 1.0 (price 1) and the expected price of crude oil is 0.65 (price 2), both with volatilities of 20%.

quantities. For example, foreign revenues in different countries may not be highly correlated, especially if the rates of exchange are not. Likewise, an oil refiner produces several products from crude oil, such as heating oil and polymers, in addition to gasoline. Relative prices of these outputs affect the quantities of each produced. However, these more realistic problems may also be solved in the manner described above.

3. Applications and Empirical Implications

3.1 Applying the model in practice

Quantitative models have had less success in penetrating corporate finance compared to investment and derivative applications. There are many possible explanations, but two probable reasons are (1) corporate finance models tend to simplify problems in order to highlight their objective or remain tractable, and (2) input parameters to the models are often difficult or impossible to estimate. While our model is more practical than many, it still faces these challenges. This section describes the implementation of our model for managing exchange-rate risk at a Fortune 100 corporation.

HDG Inc. (pseudonym) is a U.S.-based manufacturer of durable equipment.[26] Approximately half of the company's $10 billion in 1997 revenues are from foreign sales. HDG actively uses foreign exchange (FX) derivatives to hedge its exposure to exchange-rate movements. For example, in 1997 the company undertook about $15 billion (notional) in FX derivative trades and held roughly $3 billion (notional) in FX derivatives at fiscal year-end. In the spring of 1998, one of the authors adapted the model presented here for use by the FX risk management group (part of the centralized corporate treasury).

As a practical matter, HDG has a fairly constrained FX hedging policy. For example, corporate policy set forth by the board of directors sets minimum and maximum notional values (as a percent of exposure) for hedges. In addition, to obtain hedge accounting treatment for its FX derivatives, the company separates exposures by currency and by fiscal quarter (henceforth, a currency-quarter) and prefers plain vanilla options and forwards. With direct FX exposures in about 25 currencies and a typical hedging horizon of four quarters, HDG can have as many as 100 separate "hedges" in place. The parameters of the FX derivatives are chosen by the FX risk management group so as to provide a most effective hedge for that currency-quarter. In practice, the goal of implementing the model is to quantify the notion of a "most effective hedge" and translate this into a specific FX derivative position.

[26] The company has requested that its identity not be revealed. A detailed description of HDG and its foreign exchange risk management program is presented in Brown (2001).

Determining the qualities of an effective hedge is akin to specifying the functional form of $C(P)$. Because HDG does not wish to aggregate its exposures across currencies, the interpretation of $C(P)$ differs from that in the model presented here. Specifically $C(P)$ as it applies to HDG's hedging practice is best thought of as a penalty function that quantifies the relative desirability of different cash flow states for each currency-quarter instead of as a firmwide deadweight cost function. For example, a primary goal of HDG's hedging program is to prevent downside surprises to earnings from FX movements.[27] Through discussions with treasury management, it was determined that an appropriate cost function would penalize variation in USD cash flow, but with a larger penalty for downside variation. Specifically the functional form

$$C(P) = \beta\big[\alpha \max(0, E[P]-P)^2 + (1-\alpha)\max(0, P-E[P])^2\big] \quad \alpha \in [0,1], \beta \in (0,\infty) \qquad (26)$$

was chosen so that $\alpha > 0.5$ weights variation below the mean more than variation above the mean and β determines the overall cost of variation. As was the case with other deadweight cost functions, variation in α or β has a relatively small impact on the makeup of the optimal hedge.

A more challenging issue for implementing the model is estimation of other unobserved parameters, such as the volatility of foreign revenues and the correlation between foreign revenues and exchange rates. Estimating these parameters from historical (time-series) accounting data is confounded by the rapid growth in some of HDG's foreign markets. To remove much of the growth trend, we chose to estimate parameters from percentage deviations from internal forecasts at appropriate horizons.[28] For example, the variance of foreign revenues at the three-quarter horizon is estimated using the percentage difference between quarterly revenue forecasts at the three-quarter horizon and realized values for that quarter. Correlations are calculated with similar forecast data: we use the percentage change in forecasted revenues and the (end of quarter) percentage change in exchange rates.[29] Estimated correlations are near zero for most markets (0.04 on average) and ranged from a low of -0.26 to a maximum of 0.34. The volatilities of foreign revenues depend much more on the foreign market. Annualized revenue volatilities range from a low of 15.8% (Spain) to 116.0% (Hong Kong). The remaining parameters are more easily estimated. The expected exchange rate is assumed to be the forward rate, exchange rate volatilities are estimated from

[27] Again, see Brown (2001) for a detailed discussion of HDG's motivations for FX risk management.

[28] HDG's foreign revenue forecasts are surprisingly unbiased with an (average across all currencies) mean absolute error of only 2.2% from 1996:Q1 to 1998:Q2.

[29] Because correlations are statistically difficult to estimate, we use all the data for each currency to estimate a single correlation for each currency instead of a correlation for each forecast horizon.

option-implied volatilities, and costs of production are available from internal forecasts.

To determine optimal hedging strategies, the preceding parameters are combined with HDG policy constraints for the notional value of the hedge portfolio and the moneyness of options (e.g., at-the-money forward plus or minus 5%). The model is solved numerically as described in Section 1.3 for a variety of alternative hedging strategies, including the optimal exotic hedge.

In most cases the currency-quarter of interest would already have a hedge portfolio in place and HDG risk managers are interested in the effectiveness of the current hedge and the optimal method for updating the hedge. For example, should the company sell out of its current options and construct a new portfolio or can a nearly equivalent hedge portfolio be created by simply adding new positions to the existing portfolio (perhaps at less expense)? By using the optimal exotic as a benchmark, it is easy to quantify the answer to this question and others such as the cost of the policy constraints on notional value and type of hedging instruments.

In general, we find that a highly efficient hedge can be constructed from only two options (with differing strike prices) and forward contracts. As suggested in Section 2.2, near-term hedges with only forwards are usually very efficient. Consistent with the findings in Section 1.4, multiple options are most important when foreign revenues are volatile, as was the case with most of HDG's smaller and newer foreign markets. Unfortunately this finding points out a potential shortcoming of the model. Several of HDG's smaller markets were still recovering from the Asian financial crisis and liquidity in the FX options market was limited. Without explicitly incorporating transaction costs for options with different strike prices, it is difficult to assess the true costs and benefits of hedging with options.

HDG's desire to separate hedges by currency and quarter simplifies the optimization problem greatly by reducing the dimension of the state space. However, as noted in Section 2.5, this comes at a cost. Ideally HDG would consider all of its exposures both across currencies and through time when determining the optimal hedging strategy. For HDG, this is a difficult numerical problem (since $N \simeq 100$). Two experiments allow us to estimate the cost of HDG's constrained program without having to solve the complete problem.

First, the optimal hedges for a single currency, but for all forecasted quarters, are determined by solving the multiperiod problem. For this case we switch to a simulation method that generates correlated price and foreign revenue paths. This method has the added benefit of being able to sample the joint process at a high frequency (e.g., daily) and then evaluate a richer set of derivative contracts with path-dependent features such as Asian or barrier options. The results of this experiment indicate that portfolios of vanilla options and forwards with varying maturity dates are still highly effective hedging tools for most reasonable parameter values. Asian features provide a clear advantage only when foreign revenues and exchange rates are highly correlated. Finally, barrier features typically reduce hedge efficiency slightly, but can appreciably trim up-front option premiums. This suggests that (unlike HDG) if a hedging program were given a limited budget for option premiums, it should consider the use of barrier options.

A second experiment expands the problem across a larger set of currencies to measure the costs to HDG of not cross-hedging. The results of these experiments confirm the results from Section 2.5. As long as the correlations between USD exchange rates are large, highly efficient hedges with derivatives based on single USD-based exchange rates are possible. This experiment also provides for measuring the benefits of basket options and other derivatives based on multiple exchange rates (that implicitly provide a cross-hedge). As with Asian options, basket features can reduce premiums, but only increase efficiency significantly for specific correlation scenarios. In this case, when correlations across currencies are low. In sum, HDG's policy of treating each currency-quarter separately does not appear to significantly reduce the effectiveness of their hedging program. However, the constraints implied by the policy do increase the up-front expenditure on option premiums. From a cash-efficiency standpoint, this may imply an important opportunity cost to the firm.

3.2 Empirical implications

Our hedging model has a host of testable implications. Many of these are new; others are consistent with prevailing theories of corporate risk management. Qualitatively our model is consistent with some stylized empirical facts. First, corporations on average only hedge a fraction of their exposures. Stulz (1984) suggests this is due to firms "selectively hedging," with a view on future market levels. Our model provides an alternative explanation, in that the optimal hedge ratio will typically be less than 1.0 if, on average, correlation between price and quantity is negative. Second, firms often use a variety of derivative types. Our model suggests that there is a value maximizing motivation for this behavior.

The current empirical evidence concerning the validity of existing theories is mixed. There is evidence that firms use risk management to reduce the probability of financial distress, ensure the availability of internally generated funds, minimize expected tax liability, and reduce the underinvestment problem.[30] Tufano (1996) reports evidence that the exact nature of senior

[30] Nance, Smith, and Smithson (1993) find evidence that larger firms with less interest coverage, many growth opportunities, and few hedging substitutes hedge more. Some of these findings are supported by Berkman and Bradbury (1996) in a sample of firms from New Zealand. Géczy, Minton, and Schrand (1997) find that currency derivatives are used by large firms with tight financial constraints, large growth opportunities, and extensive currency exposures. Dolde (1995) finds that financial leverage is also related to hedging activity. Other studies such as Mian (1996) find that distress costs cannot explain hedging activity.

management's financial claims on the firm's stock affects their hedging strategy; managers who hold more options hedge less than managers with large stock positions.

Most of the empirical literature has tested univariate relationships as predicted by theories that focus on one specific determinant of corporate hedging incentives. Often proxies for a determinant are themselves endogenously determined by exogenous firm characteristics and the firm's operating environment. For example, univariate distress cost models predict that a firm with higher leverage is more likely to benefit from hedging because it is closer to a state of financial distress. This argument ignores the fact that firms operating in stable environments (e.g., consumer nondurables) are more likely to be highly leveraged than those in risky and rapidly changing business climates (e.g., high tech or software development). Because the highly leveraged firms operate in stable environments, the higher endogenous leverage may not be sufficiently large to make distress a significant concern. Thus there is a need to specify and test relationships between exogenous variables such as fundamental business and price risk (and perhaps most importantly correlations between variables fundamental to the firms' profits) and the endogenous hedging policy decision.

For example, our model suggests the following testable relationships:

> When fundamental business risks (e.g., variations in sales) and price risks are uncorrelated, forwards (and swaps) are efficient risk management tools. This suggests that firms with low correlation between sales volume and price risk should be less likely to use options or exotic derivatives. This may help explain the popularity of linear hedging instruments over nonlinear derivatives.
>
> Firms with negative price-quantity correlation are more likely to benefit from options or exotic derivatives. The analysis in Section 2.2 indicates that although the overall gains from hedging are lower when the firm faces negative correlation, the benefit of using a nonlinear strategy is proportionately greater (in terms of efficiency). Firms with negative (positive) price-quantity correlation are more likely to buy (sell) options. Because we are able to determine exactly which firms should buy and sell convexity, we also know which firms should buy and sell options as part of their hedging strategy.
>
> Firms with relatively large quantity risk or small price risk and significant positive or negative price-quantity correlation should hedge more with options. Consequently firms with more stable production quantities or sales will have different optimal hedging strategies than firms with more volatile quantities. For example, if we assume negative price-quantity correlation, a producer of oil, for which production is less volatile, may hedge more with forwards than a producer of wheat, for which production is more volatile.

The optimal hedge for firms with smaller contribution margins will be less sensitive to price-quantity correlation. Nevertheless, there should not be a direct relationship between variable costs and the convexity of the hedge. We also predict firms with higher levels of fixed investment (such as some utilities and manufacturers) should on average hedge price risk more than firms with lower levels of fixed investment (such as many service-oriented corporations).

If, as seems likely, production quantities become less certain the farther they are in the future, then firm's hedging decisions may also depend on when revenues are realized. Unless price-quantity correlation is considerably positive, firms will tend to hedge less with forwards as revenues move farther into the future. This is generally consistent with findings and stylized facts that firms primarily hedge near-term risks [see Bodnar, Hyat, and Marston (1998)].

4. Conclusion

Optimal hedging by a value maximizing firm is a substantially more complex task than implied by previous research. We have shown that simply selling expected output forward is rarely, if ever, the optimal risk management strategy. We have discovered this by analyzing a simple model of a value maximizing firm. In our model we have

- used value maximization instead of minimum variance (or some other objective) as the firm's objective,
- exposed the firm to unhedgable (quantity) risks,
- expanded the set of available contracts with which the firm can hedge to include options and custom derivatives, and
- allowed production technology to enter into the hedging decision.

Because of the significant amount of research on why firms should hedge, we do not propose yet another explanation, but instead simply assume that firms find some states of nature to be relatively more or less costly (as a result of market imperfections such as bankruptcy costs, financing costs, etc.). We assume that these are deadweight costs to the firm's shareholders and that these costs are accurately summarized as a function of the firm's profits. This allows us to concentrate on the more practical and, to the best of our knowledge, unexplored question of how value maximizing firms should hedge.

We find many factors that materially affect the optimal hedge. These include the volatility of the marketable good's price, the volatility of the quantity of that good, and most importantly the correlation between the price and the quantity. In fact, ignoring the impact of price-quantity correlation can lead a firm to sell forward contracts when it should have bought them. Also of critical importance are the costs of production (or equivalently operating leverage). To our surprise, the functional form of the deadweight costs (the fundamental reasons why a firm should hedge) is less important in determining the qualitative features of the optimal hedge.

Our model has specific implications regarding nonlinear hedging strategies. We show that the slope of the optimal exotic derivative at the forward price is equal to the optimal number of forward contracts the firm should buy. As the price level deviates from the forward price, the optimal exotic derivative adds convexity to the linear payoff of the forward contract to construct the value maximizing payoff function. Since we are able to derive an analytical solution for the perfect exotic derivative, we can isolate which firm types (or characteristics) require nonlinear hedges.

We find that firms with negative price-quantity correlation often benefit substantially from options and exotic derivatives. High quantity volatility or low price volatility magnifies these advantages. Firms with positive price-quantity correlation will profit less from nonlinear payoffs, and firms with negligible correlation might as well use only forward contracts. We are also able to deduce that firms with negative price and quantity correlation generally benefit from buying options, and firms with positive correlation typically benefit from selling options.

We are also left with some new unanswered questions. The specific case we analyze strips the firm to its basics: a single-product, price-taking company with linear production costs making a one-period hedging decision. A more detailed model of the firm should explicitly include financing decisions and hedging's effects on capital structure (e.g., probability of distress). Also, it could be informative to investigate how empirical distributions change the firm's optimal choice of hedging instrument, since these distributions typically have higher moments that deviate significantly from the normal and lognormal distributions analyzed in this paper. We leave these extensions to future research.

Appendix A: Derivation of the Optimal Forward Hedge

The marginal net profit function for a firm that faces an exponential deadweight cost function of the form $C(P) = c_1 e^{-c_2 P}$ and uses forward contracts to manage price risk can be written as

$$\frac{\partial \pi(a)}{\partial a} = c_1 c_2 \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} (p - \mu_p) e^{-c_2 [pq - s_1 q - s_2 + a(p - \mu_p)]} h(p, q)\, dq\, dp, \tag{A1}$$

where $h(p, q)$ denotes the bivariate normal density function

$$h(p, q) = \frac{1}{2\pi \sigma_p \sigma_q \sqrt{1 - \rho^2}} e^{\frac{-1}{2(1-\rho^2)} \left[\frac{(p-\mu_p)^2}{\sigma_p^2} - 2\rho \frac{(p-\mu_p)}{\sigma_p} \frac{(q-\mu_q)}{\sigma_q} + \frac{(q-\mu_q)^2}{\sigma_q^2} \right]}. \tag{A2}$$

Collecting terms with similar powers of p and q reduces the marginal net profit function to

$$\frac{\partial \pi(a)}{\partial a} = A e^{d_6(a)} \int_{-\infty}^{\infty} (p - \mu_p) e^{-d_1 p^2 + d_2(a) p} \int_{-\infty}^{\infty} e^{-d_4 q^2 + d_3 pq + d_5 q} dq\, dp, \tag{A3}$$

where

$$
\begin{aligned}
A &= \frac{c_1 c_2}{2\pi\sigma_p\sigma_q\sqrt{1-\rho^2}}, & B &= \frac{1}{2(1-\rho^2)},\\
d_1 &= \frac{B}{\sigma_p^2}, & d_2(a) &= 2B\left(\frac{\mu_p}{\sigma_p^2}-\frac{\rho\mu_q}{\sigma_p\sigma_q}\right)-c_2 a,\\
d_3 &= c_2-\frac{2B\rho}{\sigma_p\sigma_q}, & d_4 &= \frac{B}{\sigma_q^2},\\
d_5 &= c_2 s_1+2B\left(\frac{\mu_q}{\sigma_q^2}-\frac{\rho\mu_p}{\sigma_p\sigma_q}\right), & d_6(a) &= c_2(s_2+a\mu_p)-B\left(\frac{\mu_p^2}{\sigma_p^2}-\frac{2\rho\mu_p\mu_q}{\sigma_p\sigma_q}+\frac{\mu_q^2}{\sigma_q^2}\right).
\end{aligned}
$$

One further substitution, $d_7 = d_5 - d_3 p$, allows us to write the integral over q in a common form

$$\frac{\partial\pi(a)}{\partial a} = Ae^{d_6(a)}\int_{-\infty}^{\infty}(p-\mu_p)e^{-d_1p^2+d_2(a)p}\int_{-\infty}^{\infty}e^{-d_4q^2+d_7q}\,dq\,dp. \tag{A4}$$

We can now integrate over q,

$$\frac{\partial\pi(a)}{\partial a} = Ae^{d_6(a)}\int_{-\infty}^{\infty}(p-\mu_p)e^{-d_1p^2+d_2(a)p}\sqrt{\frac{\pi}{d_4}}\,e^{\frac{d_7^2}{4d_4}}\,dp. \tag{A5}$$

Substituting for d_7 and rearranging leads to

$$\frac{\partial\pi(a)}{\partial a} = A\sqrt{\frac{\pi}{d_4}}\,e^{d_6(a)+\frac{d_5^2}{4d_4}}\int_{-\infty}^{\infty}(p-\mu_p)e^{-\frac{1}{2}\left[\left(2d_1-\frac{d_3^2}{2d_4}\right)p^2+\left(-2d_2(a)+\frac{d_3d_5}{d_4}\right)p\right]}dp. \tag{A6}$$

This integral has a finite solution when $\frac{2d_1-d_3^2}{2d_4} > 0$. This restriction can be restated as

$$Q_2(c_2) = 1+2c_2\rho\sigma_p\sigma_q-(1-\rho^2)c_2^2\sigma_p^2\sigma_q^2 > 0. \tag{A7}$$

It implies that the cost parameter c_2 cannot be arbitrarily large. It must be specified such that $Q_2(c_2) > 0$. The cost parameter c_2 must therefore be between the two roots of the equation $Q_2(c_2) = 0$. This implies that c_2 must satisfy

$$\frac{-1}{(1+\rho)\sigma_p\sigma_q} < c_2 < \frac{1}{(1-\rho)\sigma_p\sigma_q}. \tag{A8}$$

All numerical examples presented in the main body of the text satisfy this restriction. With $Q_2(c_2) > 0$ we can proceed by performing the final integration over p. After some simplifications the marginal net profit equation can be restated as a function of the choice variable a (the number of forward contracts):

$$\frac{\partial\pi(a)}{\partial a} = P_1(a)\exp[P_2(a)], \tag{A9}$$

where $P_1(a)$ and $P_2(a)$ are linear and quadratic functions of a defined as

$$P_1(a) = \frac{-c_1c_2^2\sigma_p^2a+c_1c_2\left[c_2\mu_q\sigma_p^2-(1-\rho^2)(s_1-\mu_p)c_2^2\sigma_q^2\sigma_p^2+(s_1-\mu_p)c_2\rho\sigma_q\sigma_p\right]}{\left[1+2c_2\rho\sigma_q\sigma_p-(1-\rho^2)c_2^2\sigma_q^2\sigma_p^2\right]^{\frac{3}{2}}}, \tag{A10}$$

$$P_2(a) = \frac{\left\{ \begin{array}{c} \frac{1}{2}c_2^2\sigma_p^2 a^2 + c_2\left[c_2\mu_q\sigma_p^2 - (1-\rho^2)(\mu_p - s_1)c_2^2\sigma_q^2\sigma_p^2 + (s_1-\mu_p)c_2\rho\sigma_q\sigma_p\right]a \\ +c_2\left[\mu_q(s_1-\mu_p)+s_2+c_2\left(\begin{array}{c} \frac{1}{2}\mu_q^2\sigma_p^2 + \left(\frac{1}{2}(s_1+\mu_p^2) - \mu_p s_1\right)\sigma_q^2 \\ +(\mu_q(s_1-\mu_p)+2s_2)\rho\sigma_q\sigma_p \end{array}\right)\right] \\ -c_2^2 s_2(1-\rho^2)\sigma_q^2\sigma_p^2 \end{array} \right\}}{1+2c_2\rho\sigma_q\sigma_p - (1-\rho^2)c_2^2\sigma_q^2\sigma_p^2}. \tag{A11}$$

Finally, the optimal hedging policy can be determined as the solution a^* to the simplified first-order condition

$$\frac{\partial \pi(a)}{\partial a} = P_1(a)\exp[P_2(a)] = 0. \tag{A12}$$

Appendix B: Derivation of the Perfect Exotic Hedge

Let us begin by writing the firm's expected net profit as a function of the optimal derivative strategy $x(p)$. Because we calculate the firm's expected profit net of the cost of the derivative strategy, we can, without loss of generality, assume that the derivative strategy $x(p)$ is a zero cost contract. The expected net profit can then be written as

$$\pi(x(p)) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} pq - s_1 q - s_2 - c_1 e^{-c_2(pq - s_1 q - s_2 + x(p))} h(p,q)\, dp\, dq. \tag{B1}$$

We can simplify the net profit function by using the definition of the bivariate normal density function, $h(p,q)$, integrating over q, and integrating all terms that do not depend on $x(p)$ over p. This allows us to write the expected net profit as

$$\pi(x(p)) = \rho\sigma_q\sigma_p + \mu_p\mu_q - s_1\mu_q - s_2 - \int_{-\infty}^{\infty} K e^{-\frac{1}{2}(Ap^2+Bp+C) - c_2 x(p)}\, dp, \tag{B2}$$

where

$$\begin{aligned} K &= \frac{c_1}{\sqrt{2\pi}\sigma_p}, \\ A &= \frac{1}{\sigma_p^2} + 2c_2\rho\frac{\sigma_q}{\sigma_p} - (1-\rho^2)c_2^2\sigma_q^2, \\ B &= \frac{-2\mu_p}{\sigma_p^2} + 2c_2\mu_q - 2(\mu_p+s_1)c_2\rho\frac{\sigma_q}{\sigma_p} + 2(1-\rho^2)s_1c_2^2\sigma_q^2, \\ C &= \frac{\mu_p^2}{\sigma_p^2} - 2c_2(\mu_q s_1 + s_2) + 2c_2 s_1\mu_p\rho\frac{\sigma_q}{\sigma_p} - (1-\rho^2)s_1^2c_2^2\sigma_q^2. \end{aligned}$$

We can now maximize the firm's expected net profit subject to the constraint that the chosen derivative payoff is costless. Form the Lagrangian

$$L = \pi(x(p)) - \lambda\int_{-\infty}^{\infty} x(p)g(p)\, dp, \tag{B3}$$

take the partial derivative with respect to $x(p)$ for all p and for λ, and set these expressions equal to zero,

$$\frac{\partial L}{\partial x(p)} = g(p)\,dp + c_2 K e^{-\frac{1}{2}(Ap^2+Bp+C)-c_2 x(p)}\,dp - \lambda g(p)\,dp = 0 \qquad \forall p, \tag{B4}$$

$$\frac{\partial L}{\partial \lambda} = -\int_{-\infty}^{\infty} x(p) g(p)\,dp = 0. \tag{B5}$$

We can now solve for the optimal derivative contract. First, solve $x(p)$ for an arbitrary $\lambda > 1$. Then use the zero price constraint to solve for λ. Some algebra shows that the optimal derivative security has the payoff function

$$x(p) = \alpha_2 p^2 + \alpha_1 p + \alpha_0, \tag{B6}$$

where

$$\begin{aligned}
\alpha_2 &= -\rho\frac{\sigma_q}{\sigma_p} + \frac{1}{2}(1-\rho^2)c_2\sigma_q^2,\\
\alpha_1 &= -\mu_q + (\mu_p + s_1)\rho\frac{\sigma_q}{\sigma_p} - (1-\rho^2)s_1 c_2 \sigma_q^2,\\
\alpha_0 &= -\alpha_2\left(\sigma_p^2 + \mu_p^2\right) - \alpha_1\mu_p.
\end{aligned}$$

References

Adler, M., and J. Detemple, 1988, "On the Optimal Hedge of a Nontraded Cash Position," *Journal of Finance*, 43, 143–153.

Ahn, D., J. Boudoukh, M. Richardson, and R. Whitelaw, 1999, "Optimal Risk Management Using Options," *Journal of Finance*, 54, 359–375.

Berkman, H., and M. Bradbury, 1996, "Empirical Evidence on the Corporate Use of Derivatives," *Financial Management*, 25, 5–13.

Bodnar, G., G. Hyat, and R. Marston, 1998, "1998 Wharton Survey of Financial Risk Management by US Non-Financial Firms," *Financial Management*, 27, 70–91.

Breeden, D., and S. Viswanathan, 1996, "Why Do Firms Hedge? An Asymmetric Information Model," working paper, Duke University.

Brennan, M., and R. Solanki, 1981, "Optimal Portfolio Insurance," *Journal of Financial and Quantitative Analysis*, 16, 279–300.

Brown, G., 2001, "Managing Foreign Exchange Risk with Derivatives," *Journal of Financial Economics*, 60, 401–448.

Brown, G., and Z. Khokher, 2001, "Corporate Hedging, Market Imperfections and Speculative Motives," working paper, University of North Carolina at Chapel Hill.

Carr, P., X. Jin, and D. Madan, 2001, "Optimal Investment in Derivative Securities," *Finance and Stochastics*, 5, 33–59.

Cuoco, D., 1997, "Optimal Consumption and Equilibrium Prices with Portfolio Constraints and Stochastic Income," *Journal of Economic Theory*, 72, 33–73.

DeMarzo, P., and D. Duffie, 1991, "Corporate Financial Hedging with Proprietary Information," *Journal of Economic Theory*, 53, 261–286.

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

Dolde, W., 1995, "Hedging, Leverage, and Primitive Risk," *Journal of Financial Engineering*, 4, 187–216.

Duffie, D., W. Fleming, H. M. Soner, and T. Zariphopoulou, 1997, "Hedging in Incomplete Markets with HARA Utility," *Journal of Economic Dynamics and Control*, 21, 753–782.

Duffie, D., and H. Richardson, 1991, "Mean-Variance Hedging in Continuous Time," *Annals of Applied Probability*, 1, 1–15.

Duffie, D., and T. Zariphopoulou, 1993, "Optimal Investment with Undiversifiable Income Risk," *Mathematical Finance*, 3, 135–148.

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

Froot, K., and J. Stein, 1998, "Risk Management, Capital Budgeting and Capital Structure Policy for Financial Institutions: An Integrated Approach," *Journal of Financial Economics*, 47, 55–82.

Géczy, C., B. Minton, and C. Schrand, 1997, "Why Firms Use Currency Derivatives," *Journal of Finance*, 52, 1323–1354.

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

He, H., and H. Pagès, 1993, "Labor Income, Borrowing Constraints, and Equilibrium Asset Prices," *Economic Theory*, 3, 663–696.

Hentschel, L., and S. P. Kothari, 2001, "Are Corporations Reducing or Taking Risks with Derivatives?" *Journal of Financial and Quantitative Analysis*, 36, 93–118.

Lapan, H., G. Moschini, and S. Hanson, 1991, "Production, Hedging, and Speculative Decisions with Options and Futures Markets," *American Journal of Agricultural Economics*, 73, 66–74.

Leland, H., 1980, "Who Should Buy Portfolio Insurance?," *Journal of Finance*, 35, 581–594.

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

Mello, A., J. Parsons, and A. Triantis, 1995, "An Integrated Model of Multinational Flexibility and Financial Hedging," *Journal of International Economics*, 39, 27–51.

Mian, S., 1996, "Evidence on Corporate Hedging Policy," *Journal of Financial and Quantitative Analysis*, 31, 419–439.

Modigliani, F., and M. Miller, 1958, "The Cost of Capital, Corporate Finance, and the Theory of Investment," *American Economic Review*, 30, 261–297.

Moschini, G., and H. Lapan, 1995, "The Hedging Role of Options and Futures Under Joint Price, Basis, and Production Risk," *International Economic Review*, 36, 1025–1049.

Myers, S., 1977, "Determinants of Corporate Borrowing," *Journal of Financial Economics*, 5, 147–175.

Nance, D., C. Smith, and C. Smithson, 1993, "On the Determinants of Corporate Hedging," *Journal of Finance*, 48, 267–284.

Petersen, M., and S. R. Thiagarajan, 2000, "Risk Management and Hedging: With and Without Derivatives," *Financial Management*, 29, 5–30.

Rolfo, J., 1980, "Optimal Hedging Under Price and Quantity Uncertainty: The Case of a Cocoa Producer," *Journal of Political Economy*, 88, 100–116.

Shapiro, A., and S. Titman, 1986, "An Integrated Approach to Corporate Risk Management," in J. Stern and D. Chew, (eds.), *The Revolution in Corporate Finance*, Basil Blackwell, Oxford.

Siegel, D. R., and D. F. Siegel, 1988, *Futures Markets*, Dryden Press, Orlando, FL.

Smith, C., and R. Stulz, 1985, "The Determinants of Firms' Hedging Policies," *Journal of Financial and Quantitative Analysis*, 20, 391–402.

Stulz, R., 1984, "Optimal Hedging Policies," *Journal of Financial and Quantitative Analysis*, 19, 127–140.

Stulz, R., 1996, "Rethinking Risk Management," *Journal of Applied Corporate Finance*, 9, 8–24.

Svensson, L., and I. Werner, 1993, "Nontraded Assets in Incomplete Markets: Pricing and Portfolio Choice," *European Economic Review*, 37, 1149–1168.

Tufano, P., 1996, "Who Manages Risk? An Empirical Examination of Risk Management Practices in the Gold Mining Industry," *Journal of Finance*, 51, 1097–1137.