Optimal Hedging with Higher Moments

CHRIS BROOKS
ALEŠ ČERNÝ
JOËLLE MIFFRE*

This study proposes a utility-based framework for the determination of optimal hedge ratios (OHRs) that can allow for the impact of higher moments on hedging decisions. We examine the entire hyperbolic absolute risk aversion family of utilities which include quadratic, logarithmic, power, and exponential utility functions. We find that for both moderate and large spot (commodity) exposures, the performance of out-of-sample hedges constructed allowing for nonzero higher moments is better than the performance of the simpler OLS hedge ratio. The picture is, however, not uniform throughout our seven spot commodities as there is one instance (cotton) for which the modeling of higher moments decreases welfare out-of-sample relative to the simpler OLS. We support our empirical findings by a theoretical analysis of optimal hedging decisions and we uncover a novel link between OHRs and the minimax hedge ratio, that is the ratio which minimizes the largest loss of the hedged position. © 2011 Wiley Periodicals, Inc. Jrl Fut Mark 32:909–944, 2012

*Correspondence author, EDHEC Business School, 393 Promenade des Anglais, 06202, Nice, France. Tel: +33-(0)4-93-18-32-55, Fax: + 33-(0)4-93-83 08-10. e-mail: joelle.miffre@edhec.edu

- *Chris Brooks is a Professor of Finance at the ICMA Centre, University of Reading, England, United Kingdom.*

- *Aleš Černý is a Professor of Finance at Cass Business School, City University London, London, United Kingdom.*

- *Joëlle Miffre is a Professor of Finance at EDHEC Business School, Nice, France.*

1. INTRODUCTION

There is now indisputable evidence to suggest that the return distributions of risky assets depart from normality.[1] Under some fairly weak assumptions concerning the shape of investor utility functions, Scott and Horvarth (1980) show that investors are concerned not just with the mean and variance of asset returns but also with the distribution's higher moments, exhibiting a preference for larger odd moments and smaller even ones. Importantly, Kraus and Litzenberger (1976, 1983) and Harvey and Siddique (2000) have made it clear that systematic risk related to skewness is priced by the market.[2] Moments higher than the third have also been taken into account in the asset pricing literature.[3] With the noticeable exceptions of Kallberg and Ziemba (1983), Post, Vliet, and Levy (2008), and arguably also Jondeau and Rockinger (2006), the general message from these studies is that higher moments do matter in terms of asset pricing and that a failure to account for them may lead to sub-optimal asset allocation decisions.

Another, almost entirely separate strand of finance literature, has looked at the hedging decisions of risk-averse investors.[4] A large number of empirical studies[5] have been concerned with the estimation of the optimal hedge ratio (OHR), defined as the optimal number of futures contracts to employ per unit of the spot asset to be hedged.[6] An easy way to calculate this number of futures contracts is to employ the OLS hedge ratio, which is simply measured as the slope coefficient of an OLS regression of spot returns on futures returns. This implies a static risk management strategy that involves a one-off decision on the optimal hedge and might therefore yield suboptimal hedging decisions in periods of high basis volatility. To overcome this problem, quite a large literature has developed, which models the OHR within a conditional

[1]For example, deviations from normality have been observed for emerging stock market indices (Harvey, 1995), hedge fund indices (Agarwal & Naik, 2004), individual hedge funds (Brooks & Kat, 2002), relative-strength strategies (Harvey & Siddique, 2000), and futures contracts (Christie-David & Chaudhry, 2001).
[2]Along the same lines, examples of studies that extend the existing literature to incorporate skewness include Hong, Tu, and Zhou (2007) on measuring the economic significance of incorporating asymmetries into investment decisions, Chunhachinda, Dandapani, Hamid, and Prakash (1997) on international portfolio decision in the presence of skewness, Barone-Adesi, Gagliardini, and Urga (2004) on incorporating co-skewness into asset pricing models, and Post et al. (2008) on risk aversion and skewness preference.
[3]Of note for example are Kallberg and Ziemba (1983) on the impact of different utility functions on asset allocation, and Jondeau and Rockinger (2006) on optimal portfolio allocation.
[4]It is generally accepted that privately held, owner-managed firms are risk-averse. Listed companies, too, can act as risk averters in the presence of capital market imperfections, i.e. outside the Modigliani–Miller paradigm (Brown & Toft, 2002; Froot, Scharfstein, & Stein, 1993).
[5]The existing theoretical treatment of optimal hedging relies either on joint normality or log-normality (Brown & Toft, 2002; Moschini & Lapan, 1995) or on a specific decomposition, additive or multiplicative, of hedgeable and nonhedgeable risks (Benninga, Eldor, & Zilcha, 1983; Briys, Crouhy, & Schlesinger, 1993). An exception is Harris and Shen (2003), who develop an OHR estimator that is robust to leptokurtosis.
[6]See, for example, Cechetti, Cumby, and Figlewski (1988).

framework, taking into account the dynamics between the spot and futures returns (for example, Brooks, Henry, & Persand, 2002; or Miffre, 2004). These studies have mainly employed models from the multivariate generalized autoregressive conditionally heteroscedastic (MGARCH) family. They have reached conflicting results on the out-of-sample hedging effectiveness of conditional minimum variance hedge ratios, even before taking into account the additional costs involved with continually buying and selling futures contracts so as to rebalance the hedged portfolio when the model suggests. At best, MGARCH models have led to very modest improvements in cross-hedging efficiency when evaluated on an out-of-sample basis. Hence, the benefits of active risk management strategies ought to be viewed with caution.

Almost without exception, empirical studies on the determination of OHRs at best assume that investors have two-moment (quadratic) utility functions or that the distribution of returns on the hedged portfolio is normal, so that the mean and variance alone are sufficient to determine the hedge ratio optimally.[7] In a slight generalization, Levy (1969) shows that a cubic utility function can be employed where investor preferences depend on skewness. However, it is not at all obvious, when one is released from the constraint of the mean–variance framework, why one should stop at skewness, for in addition to an aversion to negative skewness, rational investors should possess an aversion to positive excess kurtosis as well. Even less plausibly, many studies focus on minimum variance hedging, where the mean, as well as any moments of order higher than the second, is ignored. Such an assumption concerning the mean will only be appropriate if investors are infinitely risk-averse, or if the expected return is zero.

Clearly then, if return distributions depart from normality and/or mean returns are nonzero, hedging strategies that assume normality might lead to sub-optimal hedging decisions. This study therefore develops a new methodology for estimating OHRs within a utility-based framework that allows for investors to have nonzero preferences for higher moments. The approach that we propose has many advantages, including that it

1. does not employ (notoriously unreliable) estimates of higher moments,

2. does not impose a parametric distribution on returns and is therefore not subject to parameter uncertainty,

---

[7]A slightly weaker assumption than return normality is that the spot and corresponding futures returns are drawn from a multivariate elliptical distribution. In such circumstances, even if the spot returns are skewed and/or leptokurtic, the magnitude or otherwise of these higher moments is not affected by hedging with futures and thus optimally, they should not enter into the hedger's objective function.

3. does not require the futures market to be unbiased, and
4. permits fast and reliable numerical implementation.[8]

While designing our utility-based hedge ratio, we measure, for the first time, the loss of welfare that may be incurred if one uses OLS hedge ratios in nonquadratic utility functions. By doing so, we draw together the literatures on hedging with futures, and on utility maximization with higher moments.

The extant literature concerning the impact of higher moments on hedging is very sparse. We are aware of only three studies (Gilbert, Jones, & Morris, 2006; Harris & Shen, 2006; Yamada & Primbs, 2004) that study the impact of higher moments on hedging. There are, however, important differences between these studies and this study. First, Yamada and Primbs (2004) and Harris and Shen (2006) focus on value-at-risk and conditional value-at-risk, whereas we focus on utility-based hedging. Second, while we use futures contracts as hedging instruments, Yamada and Primbs (2004) use options and Harris and Shen (2006) consider cross hedging with currencies. Third, and possibly most importantly, Yamada and Primbs (2004) and Harris and Shen (2006) focus on parameterization up to the fourth moment, Gilbert, Jones, and Morris (2006) derive and apply a partial equilibrium model of hedging that allows for skewness (but not kurtosis) in the hedger's utility function, whereas we provide a more general utility-based framework that determines optimal hedges using all moments of the return distribution. Our approach therefore encompasses that of Yamada and Primbs (2004), Gilbert et al. (2006) and Harris and Shen (2006).

Interestingly, Harris and Shen (2006) show, using a set of daily currency exposures, that minimum variance hedging is likely to reduce the out-of-sample variance of the hedged portfolio, but the skewness and kurtosis are likely to fall and rise, respectively. This result indicates that the benefit of hedging may be overstated since these higher moments move in exactly the opposite directions to those preferred by a rational utility maximizer of the form described in the theoretical literature. This provides a strong motivation for developing a utility-based framework for determining the OHR, and we are the first (to our knowledge) to do so, since such an approach will automatically take these higher moments into account when estimating the hedge ratio and assessing its effectiveness.[9]

---

[8]Our approach, based on Newton's optimization method and detailed in Appendix A, is able to deal with general nonparametric distributions of returns and, apart from strict concavity, does not impose restrictions on the utility function. The method exhibits quadratic convergence, doubling the number of digits of accuracy at each iteration, and it is thus extremely fast. Apart from computing OHRs, our procedure is also suitable for solving optimal investment problems with a large number of assets because the computational effort grows only quadratically with the number of assets. By contrast, the use of co-skewness and co-kurtosis (see, for example, Harvey, Liechty, Liechty, & Müller, 2010; Jondeau & Rockinger, 2006) makes the computational time grow with the third and fourth power of the total number of assets, respectively.

[9]Similarly, Brooks and Kat (2002) observed that hedge funds, while demonstrating impressive performance on mean-variance grounds, also typically have less desirable higher moment values than traditional asset classes.

In anticipation of our results, we find that for moderate and large commodity exposures, the out-of-sample performance of hedges constructed allowing for nonzero higher moments is better than the performance of the simpler OLS hedge ratio. Hence, it seems that for most of our cross-section (namely, for six of the seven commodities studied), higher moments do matter when it comes to determining OHRs: the welfare of hedgers is generally maximized when higher moments are taken into account.

The picture is, however, not totally uniform as we uncover cases for which the modeling of higher moments hurts more than it helps: namely, for cotton, the simpler OLS hedge ratio maximizes the welfare of hedgers more than a more sophisticated hedge ratio that takes into account higher moments. The case of cotton is of interest as it highlights the limitations of our utility-based hedge ratio which works well when departures from normality are not too strong. The return distribution of cotton presents an extremely negative skew, which was induced by government incentives to stimulate production in the mid 1980s in China.[10] A hedger who is wary of the risk that cotton producers in China could again be given strong incentives to increase output, might see our utility-based hedge ratio as appealing. Any sharp drops in prices similar to those of 1986 will be better hedged within our approach than within the standard OLS. In normal circumstances, however, the extreme drops in prices driven by government-induced output are unlikely to reoccur, and thus a hedger taking into account skewness risk may end up giving too much weight to extreme events such as that of 1986. This explains why our utility-based hedge ratio fares worse under normal circumstances than standard OLS for cotton.

The remainder of the article is organized as follows. Section 2 presents the theoretical underpinning of our higher moment hedge ratio and provides a numerical illustration based on an airline company hedging its fuel exposure. Section 3 introduces the data set of seven commodities and Section 4 presents the empirical results. Finally, Section 5 concludes.

2. METHODOLOGY

2.1. Derivation

An agent who hedges a long spot position at time t using h_t futures contracts will receive the following change in wealth between times t and $t+1$, R_{t+1}, to the hedged position

---

[10]The very sharp decrease in the price of cotton during August 1986 (-58.46%) was driven by a sharp increase in stocks between 1984 and 1986. The excess supply was the result of incentives from the Chinese government where procurement in excess of standard volume was paid extra (United Nations Conference on Trade and Development report).

$$R_{t+1} = C_{t+1} - h_t F_{t+1} \tag{1}$$

where C_{t+1} and F_{t+1} denote the changes in the cash (spot) and futures prices, respectively, between times t and $t + 1$.

Although the literature on determining OHRs is now vast, traditionally, academic research has assumed that only the first two moments of the utility function are of concern to the investor. Under this assumption, and provided that the value of the hedged portfolio follows a pure martingale process, it is easy to show that the OHR is simply the ratio of the covariance between the cash and futures returns to the variance of the futures returns, equivalent to the OLS hedge.

In general, it is not obvious whether we should stop at the second, third, fourth, or nth moment. Thus, we are compelled to adopt a more general utility-based approach. To test the robustness of our results, we examine a whole family of utility functions including the logarithmic, exponential, power, and quadratic utility (the so-called hyperbolic absolute risk aversion (HARA) or LRT class, see Cass & Stiglitz, 1970; Ingersoll, 1987) as well as fourth moment polynomial approximations thereof. To emphasize the generality of our approach, we develop the main theoretical results for general utility functions and only later focus on the HARA class.

Definition 2.1: We call $U: \mathbb{R} \to [-\infty, \infty)$ with effective domain $\mathcal{D}_U$ (i.e. the set where U is finite) a utility function if

1. U is at least twice differentiable on the interior of $\mathcal{D}_U$,
2. $U'' < 0$ on the interior of $\mathcal{D}_U$,
3. the maximal domain $\overline{\mathcal{D}}_U$ on which U is strictly increasing has nonempty interior,
4. $\lim_{v \to -\infty} U'(v) = -\infty$ or $\lim_{v \to -\infty} U'(v) \leq 0$, where we set $U'(v) = -\infty$ for $v \notin \mathcal{D}_U$.

In cases when $\overline{\mathcal{D}}_U \subsetneq \mathcal{D}_U$, we define the inverse utility U^{-1} as taking values in $\overline{\mathcal{D}}_U$.

Fix a probability space $(\Omega, P, \mathcal{F})$ with finite sample space $\Omega := \{\omega_i: i = 1, \ldots, n\}$ $n \in \mathbb{N}$. Denote by X, Y two random variables representing the excess returns of the future contract and of the spot asset, respectively.[11] We denote their realizations concisely by X_i, Y_i with $i = 1, \ldots, n$.

---

[11]We suppress time subscripts throughout this section. The random variable X corresponds to the change in the futures index F_{t+1} and Y is interpreted as the change in the cash value C_{t+1}. The expectation $E(\cdot)$ is interpreted as the expectation at time t conditional on the information at that time.

Assumption 2.2: Throughout the study we assume that there is no arbitrage, i.e. there is a measure Q equivalent to P and such that $E^Q(X) = E^Q(Y) = 0$.

The next theorem is technical. It states that, in the absence of arbitrage, the spot position is either completely unhedgeable or the optimal hedge is well-defined. This result is no longer true, in general, with unbounded return distributions.

Theorem 2.3: Consider a utility U and an initial endowment $v \in \overline{\mathcal{D}}_U$ and define

$$u(v, \eta, \vartheta) := E(U(v + \eta Y + \vartheta X)) \in [-\infty, \infty).$$

The hedgeable set $I := \{\eta \in \mathbb{R} : \sup_{\vartheta \in \mathbb{R}} E(U(v + \eta Y + \vartheta X)) > -\infty\}$ is an interval containing an open neighborhood of zero. For every hedgeable spot position ($\eta \in I$), the maximizer in $\sup_{\theta \in \mathbb{R}} u(v, \eta, \vartheta)$ exists and is unique; we denote it by $\varphi(v, \eta)$. Here, η allows for a "wealth effect" of the spot position, or the physical exposure to the spot asset.

Define the *certainty equivalent* (CE) wealth increase in the standard way,

$$\mathrm{CE}(v, \eta, \vartheta) := U^{-1}(u(v, \eta, \vartheta)) - v$$

and denote its maximal value by $\widehat{\mathrm{CE}}(v, \eta) := \sup_{\vartheta \in \mathbb{R}} \mathrm{CE}(v, \eta, \vartheta)$. For $\eta \in I$, the quantity $\widehat{\mathrm{CE}}(v, \eta)$ is finite, whereas for $\eta \notin I$ we have $\widehat{\mathrm{CE}}(v, \eta) = -\infty$.

Our framework involves a series of one-period hedging decisions. Suppose that the hedger is long η units of the spot asset (i.e. η captures the physical exposure to the spot asset) and assume that this position is hedgeable, $\eta \in I$. If the investor does not hedge, she optimally continues to hold $\varphi(v, 0)$ futures contracts. This is known as a speculative position. The literature on optimal hedging typically assumes that the futures market is unbiased, $E(X) = 0$ in which case the speculative position is zero by Jensen's inequality. If the investor hedges optimally, her position in the futures changes to $\varphi(v, \eta)$. One can now define the OHR as the difference between the optimal futures position and the speculative position, per unit of commodity exposure,

$$\mathrm{OHR}(v, \eta) := -\frac{\varphi(v, \eta) - \varphi(v, 0)}{\eta}. \tag{2}$$

We use the standard convention whereby the hedge ratio signifies the number of futures contracts the investor *shorts* as a result of being *long* one unit of the spot asset, and hence the extra minus sign in Equation (2). In the case of backwardation or contango when $E(X) \neq 0$, the optimal futures position is nonzero even if the agent holds no spot assets, and therefore $\varphi(v, 0)$ does not constitute a hedge in itself. In such a case, only the incremental position

over and above the speculative holding $\varphi(v, 0)$ should be interpreted as the hedging position, which is reflected in Equations (2) and (3).

The welfare gain (WG) from a particular (not necessarily optimal) hedge h is defined as follows[12]

$$WG(v, \eta, h) = CE(v, \eta, \varphi(v, 0) - \eta h) - CE(v, \eta, \varphi(v, 0)). \tag{3}$$

If one wants to understand and compare optimal investment/hedging dictated by the various utility functions (i.e. those whose value depends on higher moments as opposed to quadratic utility), it is important to normalize the resulting portfolio by some measure of risk aversion. This insight goes back to Arrow (1971). The most convenient normalization factor turns out to be the Arrow–Pratt coefficient of risk aversion; see Arrow (1963), Pratt (1964), Kallberg and Ziemba (1983), and Samuelson (1970). We apply a similar normalization to the risk-adjusted performance measurement below.

Using the coefficient of local absolute risk aversion,

$$A(v) := -\frac{U''(v)}{U'(v)}$$

we define the normalized spot and futures positions as follows:

$$\lambda := A(v)\eta, \quad \theta := A(v)\vartheta. \tag{4}$$

Similarly, we define a normalized WG, which we call the *hedging potential* (HP),

$$HP(v, \eta, h) = \frac{WG(v, \eta, h)}{A(v)\eta^2}. \tag{5}$$

The normalization is performed to enable meaningful comparison of the hedging coefficients and welfare measurements across different utility functions. Essentially, we will see that the results are primarily driven by the values of λ and θ and only to a lesser extent by the shape of the utility function itself. HP is *robust* in the sense that it possesses a meaningful limit as η approaches 0 and this limit coincides across all utility functions when the futures market is unbiased.[13] For HARA utility functions, the HP is independent of the wealth level v. The role of the HP is clarified in the numerical example below (Section 2.2).

---

[12]Our measure of the welfare loss arising from using a second-best strategy is based on the CE as in Kallberg and Ziemba (1983) and Pulley (1983), by contrast to Simaan (1993), who uses a compensating variation in terminal wealth.

[13]In an unbiased futures market, the limiting value of the HP of the optimal hedge for $\eta \to 0$ is $\rho^2(X, Y)Var(Y)$, where ρ is the correlation coefficient. This holds for an arbitrary utility function, see Theorem 2.6.

To evaluate the normalized quantities, it is convenient to define a "normalized utility."

Definition 2.4: We say that $f: \mathbb{R} \to [-\infty, \infty)$ defined by

$$f(z) := c_1 U\left(v + \frac{z}{A(v)}\right) + c_2 \tag{6}$$

$$c_1 := \frac{A(v)}{U'(v)}, \quad c_2 := -c_1 U(v) \tag{7}$$

is a normalized utility to U at initial wealth $v \in \overline{\mathcal{D}_U}$.

The normalization in (6) ensures that f has a coefficient of absolute risk aversion equal to 1 at 0. The choice of c_1 and c_2 in (7) is for convenience only[14]; it achieves $f(0) = 0$ and $f'(0) = 1$. Denoting by γ the coefficient governing the shape of HARA utility, we show in Appendix B that the normalized HARA utility is given by

$$f_\gamma(z) := \begin{cases} \dfrac{(1 + z/\gamma)^{1-\gamma} - 1}{1/\gamma - 1} & \text{for } \gamma > 0, \gamma \neq 1 \\ \ln(1 + z) & \text{for } \gamma = 1 \\ \dfrac{|1 + z/\gamma|^{1-\gamma} - 1}{1/\gamma - 1} & \text{for } \gamma < 0 \\ 1 - e^{-z} & \text{for } \gamma = \pm\infty. \end{cases} \tag{8}$$

Conveniently, the normalized HARA utility depends only on the shape parameter γ and not on the initial wealth v or its own local risk aversion $A(v)$. Thus, we have obtained a very parsimonious representation of the entire HARA class, which makes it easier to present our numerical results. The literature also discusses fourth order polynomial approximations of different utility functions, obtained by the Taylor expansion.[15]

In the next theorem, we prove that one can compute the OHR and the HP generated by utility U by means of the normalized utility f.

---

[14]It is well known in utility theory that any values of $c_1 > 0$ and $c_2 \in \mathbb{R}$ would lead to the same ordinal expected utility.

[15]For the HARA utility class, the corresponding polynomial normalized utility reads

$$\tilde{f}_\gamma(z) = z - \frac{z^2}{2} + (1 + 1/\gamma)\frac{z^3}{6} - (1 + 1/\gamma)(1 + 2/\gamma)\frac{z^4}{24}. \tag{9}$$

For example, Jondeau and Rockinger (2006) use a special case of Equation (9) with $\gamma = \infty$. We have also analyzed hedging decisions using the polynomial approximations, but since they are materially the same as for the original HARA utilities, we do not report them in the study.

Theorem 2.5: Consider a utility U, initial endowment $v \in \overline{\mathcal{D}}_U$ and a corresponding normalized utility f. Define

$$a(\lambda, \theta) = f^{-1}(E(f(\lambda Y + \theta X)))$$

$$\alpha(\lambda) = \arg\max_{\theta \in \mathbb{R}} E(f(\lambda Y + \theta X))$$

$$\hat{h}(\lambda) = -\frac{\alpha(\lambda) - \alpha(0)}{\lambda} \tag{10}$$

$$g(\lambda, h) = \frac{a(\lambda, \alpha(0) - \lambda h) - a(\lambda, \alpha(0))}{\lambda^2}.$$

Then η is hedgeable for utility U if and only if λ is hedgeable for utility f and

$$\mathrm{OHR}(v, \eta) = \hat{h}(\lambda) \tag{11}$$

$$\mathrm{HP}(v, \eta, h) = g(\lambda, h) \tag{12}$$

where λ is the normalized spot position from Equation (4).

Please see Appendix C for proofs.

2.2. Airline Example

We now present an illustrative numerical example of an airline hedging its fuel exposure. This provides a physical illustration of the potential economic importance (or otherwise) of higher moments when hedging at the firm level (Brown & Toft, 2002) and demonstrates how the abstract methodology described in the previous section can be practically implemented. Given the losses that some airline companies faced in 2008–2009 as a result of hedging decisions that were with the benefit of hindsight very unfortunate, this example is worthwhile. In the following sections, we then proceed to a more detailed but also more stylized analysis using commodity data, an area which is much more familiar within the hedging literature, and which researchers in this area can more easily relate to.

Suppose that the book value of the company is $3.5 bn and the expected net income is $0.5 bn, giving projected book value $v = \$4.0$ bn. Assume that the expected fuel bill at current prices is $\eta = \$0.8$ bn and that the fuel bill uncertainty due to price variations dominates all the other uncertainty in the airline's revenues and expenses. Assume further that the airline does not wish to pass fuel cost increases onto its passengers.[16] Finally, assume that the local

---

[16]Figures based on Southwest Airlines (US) financial statement for 2000 (source: SEC 10-K filing for 2001). Unlike some other major airlines, Southwest did not apply fuel surcharges to fares—see Morrell and Swan (2006).

Optimal Hedge Ratios for Airline Fuel Exposure

γ	-1	-3	1	5	∞
OHR	0.9511	0.9487	0.9654	0.9520	0.9502

Note. Illustrative example for normalized exposure $\lambda = 1$. γ determines the shape of HARA utility function. Skewness and kurtosis of the futures (spot) return distributions are, respectively, 0.3, (0.2); 4.2, (5.2). OHR, optimal hedge ratio; HARA, hyperbolic absolute risk aversion.

relative risk aversion of the airline is moderate[17] at 5. Then the normalized exposure is $\lambda = 5 \times \frac{0.8}{4.0} = 1$.

To compute the optimal hedge, we compile data on monthly jet fuel price returns[18] to obtain a histogram for Y and obtain synchronized returns for prospective cross hedging with commodity futures (in this case light crude oil) to obtain the distribution of X. With the joint empirical distribution of X and Y, we then evaluate (11) for $\lambda = 1$, using different normalized utility functions as shown in Equation (8). The results are presented in Table I. We use five utility functions: quadratic ($\gamma = -1$), which takes into account only the first two moments, quartic ($\gamma = -3$) also including (co-)skewness and (co-)kurtosis, and exponential ($\gamma = \infty$), logarithmic ($\gamma = 1$) and fourth power hyperbolic ($\gamma = 5$), which involve all moments of the joint distribution in different proportions. We discuss the relationship of the OLS hedge ratio to the quadratic utility hedge ratio in Appendix D. Numerically, the OLS hedge is in this case indistinguishable from the quadratic utility hedge ($\gamma = -1$).

It is evident that the OHRs dictated by different utility functions are very similar to each other and to the OLS hedge ratio. One may nevertheless wonder about the welfare implication of using the OLS hedge when the hedger cares about higher moments of the return distribution. It might conceivably be the case that a small deviation from the OHR causes a large loss in the certainty equivalent wealth. In Table II, we therefore report (i) OLS HP, $g(\lambda, h_{OLS}) \times 1{,}200$, the normalized WG that results from using the second-best (i.e. the OLS) hedge ratio in each utility function, and (ii) OHR HP, $g(\lambda, \hat{h}_\gamma(\lambda)) \times 1{,}200$, the WG of the optimal hedge for each utility. Function g is defined via Theorem 2.5 with normalized utility given by (8). Since we use monthly data, the multiplication by 1,200 means that we interpret the WG as the percentage point increase in projected value per year.

[17]By comparison, Brown and Toft (2002) specify the welfare loss of an unhedged position between 3.9 and 12.7% of an *expected net income*. To obtain the same result in our example, the coefficient of relative risk aversion would need to be roughly between 2.5 and 8, corresponding to a normalized exposure λ between 0.5 and 1.6.

[18]Monthly returns on U.S. Gulf Coast kerosene prices in the period from April 1990 to April 2007. Source: U.S. Department of Energy, Energy Information Administration.

Normalized Welfare Gain (HP) from Using Optimal (Denoted OHR) and
OLS Hedge Ratios, Respectively

γ	-1	-3	1	5	∞
OLS HP	4.5626	4.3793	4.9474	4.4527	4.4045
OHR HP	4.5626	4.3793	4.9484	4.4527	4.4045

Note.   Illustrative example for normalized exposure $\lambda = 1$ in jet fuel. OHR, optimal hedge ratio; HP, hedging potential.

What do these figures mean in our airline example? The welfare impact of hedging as opposed to no hedging is substantial and for all utility functions, the HP represents 440–500 basis points. However, this figure is the normalized WG corresponding to a local risk aversion of unity and 100% exposure to the commodity. A simple conversion from (5) shows

$$\frac{\text{WG}(v, \eta, h)}{v} = \text{HP}(v, \eta, h) \times (A(v)v) \times \left(\frac{\eta}{v}\right)^2 = \frac{\lambda^2}{A(v)v} \text{HP}(v, \eta, h). \qquad (13)$$

Our company has relative risk aversion $A(v)v = 5$ and normalized exposure $\lambda = 1$, and therefore the actual WG translates to $500/5 = 100$ basis points from the projected book value of 4.0 bn, or 8% of expected net profit, which is roughly in the middle of the range [3.9%,12.7%] specified by Brown and Toft (2002).

The normalization above serves two purposes. First, if we fix the shape of the HARA utility function (say exponential), then two different companies with the same λ will have exactly the same optimal hedge. The first company may have a smaller fuel exposure and higher risk aversion and the second conversely higher physical exposure to fuel price fluctuations but a lower degree of risk aversion. Such invariance is embedded in the definition of HARA utility and can be shown algebraically. More importantly, even if one uses different utility functions (for example quadratic vs. exponential), the OHRs tend to be very similar, at least for small values of λ. This is no longer guaranteed by construction, but rather it is an empirical feature that comes out of our analysis. Determining the values of λ that are "small" must be conducted empirically and we examine this question below.

A similar statement applies for the optimal HP $g(\lambda, \hat{h}(\lambda))$. If we fix the shape of the HARA utility (for example as exponential), two different companies with the same λ will have identical percentage WG *per unit of relative risk tolerance and per square of relative commodity exposure*. Significantly, this quantity has a nondegenerate limit as λ approaches 0 and in an unbiased futures market, the limit coincides across *all* utility functions. If we consider

Fuel Expenditure for Southwest Airlines

Year	Book value bb,v	Fuel cost bn,η	Fuel costto book,η/v	λ		WG/v to HP	
				RRA = 5	RRA = 20	RRA = 5	RRA = 20
2000	3.45	0.49	0.14	0.7	2.8	0.1	0.4
2001	4.01	0.80	0.20	1.0	4.0	0.2	0.8
2002	4.42	0.77	0.17	0.9	3.5	0.2	0.6
2003	5.05	0.83	0.16	0.8	3.3	0.1	0.5
2004	5.52	1.00	0.18	0.9	3.6	0.2	0.7
2005	6.68	1.34	0.20	1.0	4.0	0.2	0.8
2006	6.45	2.14	0.33	1.7	6.6	0.6	2.2
2007	6.94	2.54	0.37	1.8	7.3	0.7	2.7

Note. Conversion factor between percentage welfare gain and the hedging potential equals $(\eta/v)^2 \times$ RRA, see Equation (13). RRA, relative risk aversion; WG, welfare gain; HP, hedging potential. Source: SEC filings 2000–2007.

different utility functions (for example quadratic vs. exponential), the normalized WG is no longer identical across utility functions for a fixed value of λ but it turns out empirically that for small values of λ, the normalized hedging performance tends to be similar across different utility functions.

Theorem 2.6: Consider arbitrage-free excess returns X, Y and assume the futures market is unbiased—that is, $E(X) = 0$. For any utility U and any $v \in \overline{\mathcal{D}}_U$ we have

$$\lim_{\eta \to 0} \text{OHR}(v, \eta) = \frac{Cov(X, Y)}{Var(X)} =: h_{OLS}$$

$$\lim_{\eta \to 0} HP(v, \eta, \text{OHR}(v, \eta)) = \rho^2(X, Y)Var(Y)$$

where $\rho(X, Y)$ is the correlation coefficient between X and Y, and h_{OLS} is the OLS hedge ratio.

Let us examine typical values of λ. In Table III, we detail the book value v and the fuel expenditure η for Southwest Airlines in the period 2000–2007. The values of λ range between 0.7 and 1.8 for moderate relative risk aversion of 5, and between 2.8 and 7.3 for high risk aversion of 20. In the last two columns, we show the conversion factor between the HP and the actual percentage WG, WG/v. The table suggests that in the airline industry, λ can be as high as 7.3, while the conversion factor may go up to 2.7. It is possible that in other sectors, λ might be even higher.

We now proceed to examine the impact of higher moments for $\lambda = 7.3$. The results are shown in Table IV. Compared with Table II, the difference

Normalized Welfare Gain (Hedging Potential, HP) from Using Optimal (denoted OHR) and
OLS Hedge Ratios, Respectively

γ	-1	-3	1	5	∞
OHR	0.9511	0.9628	–	1.0448	0.9963
OLS HP	3.8244	3.7398	–	7.0314	4.8426
OHR HP	3.8244	3.7404	–	7.0821	4.8527

Note.   Illustrative example for normalized exposure $\lambda = 7.3$ in jet fuel. OHR, optimal hedge ratio; HP, hedging potential.

between the OLS hedge ratio and the utility-based hedge ratios is more pro-
nounced. The same is true for the WG of utility-based hedging. Consider, for
example, the HARA utility with $\gamma = 5$. For $\lambda = 1$, there is no perceptible differ-
ence between the OLS HP and the potential of the optimal utility-based hedge.
For $\lambda = 7.3$, the story is very different—the increase in HP amounts to five
basis points. Using the conversion factor in Table III (bottom right corner), we
find that the actual WG is 2.7 times higher, amounting nearly to 14 basis
points per year.

2.3. Hedgeable Positions and Asymptotics for Large Commodity Exposure

We now consider the dependence of the optimal hedge on the normalized expo-
sure, λ, for $\lambda \to \infty$. It is clear that for any utility with effective support bounded
from below (i.e. utility functions such as log utility, which take the value of $-\infty$
for a finite argument), hedging will become infeasible for large enough values
of λ unless a perfect hedge is available. Hence for $0 < \lambda < \infty$, it is useful to
know the range of λ values that are hedgeable. This leads to the notion of the
minimax hedge ratio.

Definition 2.7: Consider random variables X,Y with realizations $\{X_i, Y_i\}_{i=1}^{n}$. We
call $\bar{h}$ the minimax hedge ratio if it solves the problem

$$\max_{h \in \mathbb{R}} \{ \min_{i \in \{1,\dots,n\}} (Y_i - hX_i) \}. \tag{14}$$

The optimization (14) can be written as a linear program which admits a
feasible solution under the no-arbitrage assumption. Since the value function
in (14) is bounded above by zero, it follows that the minimax hedge always
exists. The minimax hedge ratio itself needs not be unique but the minimax
return (the optimized value function) always is. We can now address the issue
of hedgeable positions for HARA utility functions with $0 < \gamma < \infty$.

OHR for Jet Fuel as a Function of Normalized Fuel Exposure λ.

$\gamma\backslash\lambda$	0	1	10	100	1000
−1	**0.9511**	0.9511	0.9511	0.9511	0.9511
−3	0.9468	0.9487	0.9696	1.0478	1.0323
−5	0.9463	0.9491	0.9849	1.1823	1.1404
−15	0.9458	0.9498	1.0093	1.3054	1.2770
∞	0.9456	0.9502	1.0275	1.3184	**1.3224**

Note. Parameter γ governs the shape of utility function—see Equation (8). OHR, optimal hedge ratio.

Theorem 2.8: Denote by w the minimax return, $w := \min_{i \in \{1,\ldots,n\}}(Y_i - \bar{h}X_i) \leq 0$. The normalized position λ is hedgeable for normalized HARA utility with $0 < \gamma < \infty$ if and only if $0 < 1 + \lambda w/\gamma$.

In the airline example, the minimax hedge ratio equals 1.3239 and the corresponding minimax return is -22.6%. Consider a firm with logarithmic HARA utility ($\gamma = 1$ in Equation 2.8). In this case, fuel price risk becomes unhedgeable if the normalized exposure exceeds $1/0.226 \approx 4.4$. As λ approaches the unhedgeable threshold of 4.4 from below, the optimal hedge increasingly resembles the minimax hedge. In Table I, we have used $\lambda = 1$ which means that extreme events do not play a significant role in the hedging decision and the optimal log utility hedge does not depart significantly from the OLS hedge ratio. In contrast, when $\lambda = 7.3 > 4.4$, we obtain an unhedgeable position for the HARA logarithmic utility ($\gamma = 1$), and hence the missing values in the fourth column of Table IV.

We can now look at the behavior of the OHR for large values of λ. As transpires from the previous discussion, this question is only interesting for $\gamma < 0$ and for $\gamma = \infty$ because for $0 < \gamma < = \infty$, any large enough commodity exposure becomes unhedgeable. In Table V, we report results for five utility functions: quadratic ($\gamma = -1$), which takes into account only the first two moments, quartic ($\gamma = -3$) also including (co-)skewness and (co-)kurtosis, and another two HARA utilities using the first 6 and the first 16 moments, respectively. Finally, we employ exponential utility ($\gamma = \infty$), which involves all moments of the joint distribution.

The OHRs in the first column ($\lambda = 0$) vary slightly due to the nonzero mean futures return (that is why Theorem 2.6 does not apply here). We observe that for $\lambda \leq 1$, there is relatively a little variation in the OHR across utility functions while for $\lambda \geq 10$, higher moments matter substantially. Somewhat surprisingly, for exponential utility, the outcome for large λ is very close to the minimax hedge (we prove this theoretically in Theorem 2.9 below). One can

TABLE VI
Summary Statistics for Spot and Future Commodity Returns

	Spot Commodity					Futures Contract					OLS Hedged Portfolio				
	Mean %	*Std %*	*Skew*	*Kurt*	*JB*	*Mean %*	*Std %*	*Skew*	*Kurt*	*JB*	*Mean %*	*Std %*	*Skew*	*Kurt*	*JB*
Corn	2.1	25.5	0.25***	3.83*	191*	−8.0***	21.8	0.49*	4.21*	240.2*	10.4*	12.0	−1.10*	5.06*	390*
Cotton	2.0	28.4	−0.71*	9.58*	1156*	−3.0	24.8	0.14	1.10*	15.8*	4.6	18.6	−6.48*	81.38*	83762*
Gold	4.0	18.8	1.11*	5.83*	500*	−3.6	18.7	0.30**	5.19*	350.1*	7.5*	5.2	1.15*	10.69*	1535*
Heating Oil	8.4	38.3	0.94*	5.16*	342*	15.8**	34.9	0.53*	2.99*	114.0*	−4.7	25.3	0.14	5.19*	306*
Silver	5.6	32.9	0.57*	10.88*	1537*	−13.5**	26.1	−0.07	2.39*	61.3*	11.7*	7.9	1.05*	6.39*	484*
Soybean Oil	3.4	28.1	0.70*	2.96*	138*	−5.5	25.6	0.36*	2.09*	62.6*	9.0*	10.0	0.40*	1.23*	28*
Sugar	7.4	37.4	0.99*	2.62*	138*	−9.8	41.9	0.35*	2.01*	58.3*	14.8*	20.4	0.00	2.96*	112*

Note. The table presents summary statistics for the spot returns, the futures returns and the out-of-sample OLS hedged portfolio returns. Mean and standard deviation (Std) have been annualized. Skew is the skewness of the series and Kurt is its kurtosis. JB, Jacque–Bera test of normality. One, two, and three stars indicate significance at 1, 5, and 10% levels, respectively.

view Table V as containing two polar cases—the OLS hedge ratio in the top left corner, and the minimax hedge ratio in the bottom right corner.

The minimax HR is an ultra-cautious hedging strategy concerned solely with the most extreme events captured by the data. On the other hand, the OLS hedge ratio by construction pays more attention to small returns which occur most of the time and contribute much toward the overall variance. This may, of course, backfire if there is a temporary divergence between the spot and futures markets—Metalgesselschaft is one notorious victim of such divergence. The choice of γ shifts the focus between "extreme-event" and "every-day" hedging strategies. With $\gamma = -1$, one adopts the "every-day" approach no matter how high the exposure to the commodity λ. With $\gamma = \infty$, the right strategy is determined endogenously as a function of both the commodity exposure λ and the size of extreme events captured in the data. Table VII provides empirical values of the OLS and minimax hedge ratios for the seven commodities, which we discuss subsequently.

Based on the results in Table V, one may conjecture that HARA hedge ratios have a well-defined limit for $\lambda \to \infty$. We capture the limit analytically in the following theorem.

Theorem 2.9: 1) Choose $\gamma \in (-\infty, 0)$ and denote by $\hat{h}_\gamma(\lambda)$ the OHR generated by the normalized HARA utility (8). Suppose the problem $\min_{h \in R} E(|Y - hX|^{1-\gamma})$ has a unique minimizer. Then

$$\lim_{\lambda \to \infty} \hat{h}_\gamma(\lambda) = \arg\min_{h \in R} E(|Y - hX|^{1-\gamma}).$$

2) Suppose the minimax hedge $\overline{h}$ is unique, then the exponential utility hedge ratio tends to $\overline{h}$ as $\lambda \to \infty$,

$$\lim_{\lambda \to \infty} \hat{h}_\infty(\lambda) = \overline{h}.$$

The first part of the theorem has a clear intuitive meaning. For integer $\gamma < 0$, the expected HARA utility includes a combination of co-moments of order 1 to $1 - \gamma$. Theorem 2.6 suggests that for small λ, only the first two moments matter. For $\lambda \to \infty$, the situation is completely the opposite: only the largest moment of the hedged return matters.

The second part of the theorem uncovers a novel link between exponential utility hedge ratios and the minimax hedge ratio. In a wider context, the result hinges on the asymptotic elasticity of the utility at $-\infty$. That is, in general the result holds for *any* utility which is finite valued on $\mathbb{R}$, bounded from above, and satisfies

$$\lim_{x \to -\infty} = \frac{f(x)}{xf'(x)} = 0.$$

<div align="center">

TABLE VII

Comparison of OLS and MINIMAX In-Sample Hedge Ratios

</div>

| | Hedge Ratio | | | | |
	OLS	MINIMAX	Minimax Loss %	(co)Skewness	Excess (co)Kurtosis
Corn	1.033	0.519	15.3	0.5 0.4 0.4 0.3	4.1 4.0 3.8 3.7 3.7
Cotton	0.862	−1.203	47.2	0.1 0.2 0.7 −0.7	1.1 1.0 1.3 −0.8 9.4
Gold	0.966	0.587	7.0	0.3 0.6 0.8 1.1	5.1 4.9 5.0 5.3 5.7
Heating Oil	0.825	0.911	31.9	0.5 0.4 0.6 0.9	2.9 1.3 1.5 2.7 5.0
Silver	0.927	0.525	8.4	−0.1 0.2 0.3 0.5	2.3 2.0 1.8 1.7 1.7
Soybean Oil	1.024	0.994	6.4	0.4 0.5 0.6 0.7	2.0 2.2 2.4 2.6 2.9
Sugar	0.748	0.625	21.8	0.3 0.6 0.8 1.0	2.0 1.8 1.9 2.1 2.6

Note. Minimax loss represents the worst case loss of a portfolio hedged with the minimax hedge ratio. Coskewness values correspond to sample versions of the following population moments:

$$\frac{E((X - \mu_X)^3)}{\sigma_X^3}, \quad \frac{E((X - \mu_X)^2(Y - \mu_Y))}{\sigma_X^2\sigma_Y}, \quad \frac{E((X - \mu_X)(Y - \mu_Y)^2)}{\sigma_X\sigma_Y^2}, \quad \frac{E((Y - \mu_Y)^3)}{\sigma_Y^3}.$$

Excess cokurtosis values correspond to sample versions of the following population moments:

$$\frac{E((X - \mu_X)^4)}{\sigma_X^4} - 3, \quad \frac{E((X - \mu_X)^3(Y - \mu_Y))}{\sigma_X^3\sigma_Y} - 3\rho_{XY}, \quad \frac{E((X - \mu_X)^2(Y - \mu_Y)^2)}{\sigma_X^2\sigma_Y^2} - 1 - 2\rho_{XY}^2,$$

$$\frac{E((X - \mu_X)(Y - \mu_Y)^3)}{\sigma_X\sigma_Y^3} - 3\rho_{XY}, \quad \frac{E((Y - \mu_Y)^4)}{\sigma_Y^4} - 3.$$

Here μ_X, μ_Y are the means and σ_X, σ_Y are the standard deviations of the futures and spot returns, respectively. Variable ρ_{XY} denotes the correlation between futures and spot returns. The slight discrepancies between the skewness and kurtosis values in this table and the corresponding values in Table VI are caused by the use of different estimators. This table uses consistent estimators of co-moments which are, however, not unbiased in a normal model. Table VI employs the so-called G_1, G_2 estimators which are unbiased in a normal model (Joanes & Gill, 1998). Similar estimators for normalized co-moments are not readily available.

In the HARA class, the exponential is the only utility function with this property.

We now proceed to a more detailed analysis using a set of commodity hedges using futures contracts, with the data described in the following section and the results presented and discussed in Section 4.

3. DATA

The data, downloaded from Datastream International, comprise end-of-month spot and futures prices on seven U.S. commodities: four agricultural commodity futures (corn, cotton, soybean oil, and sugar), one energy future (heating oil), and two metal futures (gold 100 oz, and silver 1,000 oz). To compile the time-series of futures prices, we collect the futures prices on all nearest and second nearest contracts. We hold the first nearby contract up to one month before maturity. At the end of that month, we roll our position over to the second nearest contract and hold that contract up to one month before maturity.

Returns are then computed as the changes of these settlement prices. The procedure is then rolled forward to the next set of nearest and second nearest contracts when a new sequence of futures returns is compiled. The process is repeated throughout the data set to generate a sequence of nearby maturity futures returns.

The data set covers the period from January 31, 1979 to September 30, 2004. Note that we include in our analysis some commodity futures and spot assets that started trading after January 1979 or that were delisted before September 2004. As a result, the sample spans shorter periods for some contracts (cotton, heating oil, and silver).[19]

Since, by construction, any practical hedging decision is made out-of-sample, the overall period is split into two sub-samples. The in-sample period covers approximately two thirds of the data set and is used for estimation. The out-of-sample period, used for forecasting and hedging decisions, covers the remaining one-third.

Table VI presents some summary statistics for the futures returns, the corresponding spot returns, and the hedged portfolio returns, where a time-invariant OLS hedge is employed. Most spot series are significantly leptokurtic and are positively skewed because events such as hurricanes or wars positively affect commodity prices. Most noteworthy, cotton is an exception since its return distribution is negatively skewed at the 1% level. This reflects the rise in cotton stocks that took place in the mid 1980s (and culminated at 11.4 million tons in 1985–1986) that came hand-in-hand with a very sharp fall in cotton prices (−58.46% in August 1986). The excess supply was in part driven by government incentives where procurement in excess of standard volume was given an extra 30% price bonus in China. Other incentives to encourage production included subsidizing inputs (e.g., fertilizers) or advancing cash ahead of procurement.

Hedging with futures is evidently very successful for the vast majority of series. Compared with the spot return variance, the hedged portfolio variance is on average around 73% lower, and for gold, the reduction in variance is over 90%. However, interestingly, the skewness falls for the hedged portfolio returns in five of the seven series compared with the spot skewness, whereas the kurtosis rises for five of the series. Thus, if we accept the premise that hedgers are indeed concerned with higher moments, then the effectiveness of the OLS hedge may be overstated by a consideration of only the reduction in variance.

The values of the minimax return (Section 2.3) for different commodities are shown in Table VII. Empirically, the smallest minimax losses (in absolute value) occur for soybean oil (at 6.4%), gold (at 7%), and silver (at 8.4%).

[19]For these commodities, the samples used were as follows: January 1980—September 2004 for cotton, January 1982—September 2004 for heating oil, and May 1981—October 2002 for silver.

The largest minimax losses occur for cotton and heating oil, which stand at 47.2 and 31.9%, respectively. There appears to be no firm link between the minimax hedge and the OLS hedge ratio across different commodities. Broadly, the two hedge ratios tend to have the same sign and similar magnitude, but cotton is again the notable exception to this rule of thumb in our collection of commodities.

4. EMPIRICAL RESULTS

In the interests of brevity, we report only the out-of-sample results rather than those in sample, since the former are of direct practical interest as they match the investor experience. We illustrated our higher moment hedging methodology using an airline company hedging its fuel exposure in Section 2.2. We will now replicate the same computations for the seven U.S. commodities described in the previous section. To avoid repetition, we refer the reader to Section 2.2 where we provided a thorough interpretation of the reported quantities.

Table VIII measures (i) the averages and standard deviations of the OHRs over the rolling out-of-sample period—that is, the OHRs obtained for each utility function, (ii) OLS HP, $g(1, h_{OLS}) \times 1,200$, the normalized WG that results from using the second-best (i.e. the OLS) hedge ratio in each utility function, and (iii) OHR HP $g(1, \hat{h}(1)) \times 1,200$, the WG of the optimal hedge for each utility. Function g is defined via Theorem 2.5 with normalized utility given by (8). The multiplication by 1,200 means we can interpret the WG as the percentage point increase in initial wealth per year. We consider HARA utility functions with baseline risk aversion $\gamma \in \{-3, -1, 1\ 5, \infty\}$. The framework allows us to examine a much wider range of parameters, but we have found that all utility functions with $|\gamma| > 5$ essentially behave like the exponential utility, $\gamma = \infty$.

We report results for two levels of normalized exposure: moderate, $\lambda = 1$; and large, $\lambda = 10$.[20] The labelling "moderate, large" refers to the empirically observed difference between the OLS HR and the optimal utility-based HR at the given level of λ and the size of the corresponding WG from switching from OLS to optimal utility-based HRs. The labeling does not imply that the impact of hedging is medium or large, respectively. In fact, in the Southwest airline example, the value of jet fuel as a fraction of projected book value was 20%, which, assuming relative risk aversion of 5, led to $\lambda = 1$. Even though $\lambda = 1$ is labelled as "moderate," the gain from hedging was about 1% of the projected book value. This amounts to 5% of the total fuel exposure in that example and

---

[20]We also compute the results for the case where $\lambda = 0$, which we term small commodity exposure. Since the findings are qualitatively identical to those for $\lambda = 1$, in the interests of brevity they are not reported.

<div align="center">

TABLE VIII

Out-of-Sample Hedging Performance, $\lambda = 1$

</div>

Commodity (in, out)		OLS	-1	-3	1	5	∞
			γ				
Corn	OLS HP		17.83	17.87	17.98	17.91	17.89
(204, 103)	OHR HP		17.69	17.89	18.43	18.07	18.00
	mean HR	1.00	0.99	1.00	1.04	1.01	1.01
	std HR	0.01	0.01	0.01	0.02	0.01	0.01
Cotton	OLS HP		19.85	19.83	19.83	19.82	19.82
(196, 99)	OHR HP		19.85	19.42	17.50	18.97	19.16
	mean HR	0.83	0.83	0.81	0.71	0.79	0.80
	std HR	0.02	0.02	0.02	0.02	0.02	0.02
Gold	OLS HP		3.63	3.62	3.61	3.62	3.62
(204, 103)	OHR HP		3.61	3.65	3.71	3.67	3.66
	mean HR	0.97	0.96	0.98	1.01	0.99	0.99
	std HR	0.00	0.00	0.00	0.01	0.01	0.01
Heating oil	OLS HP		−11.60	−11.75	−11.93	−11.84	−11.81
(180, 91)	OHR HP		−11.55	−11.39	−11.39	−11.36	−11.37
	mean HR	0.83	0.82	0.81	0.81	0.80	0.80
	std HR	0.01	0.01	0.01	0.01	0.01	0.01
Silver	OLS HP		14.11	14.12	14.14	14.12	14.12
(170, 86)	OHR HP		13.91	14.30	14.89	14.56	14.46
	mean HR	0.94	0.92	0.96	1.00	0.98	0.97
	std HR	0.01	0.01	0.01	0.01	0.01	0.01
Soybean oil	OLS HP		12.15	12.09	12.06	12.07	12.08
(204, 103)	OHR HP		12.12	12.11	12.18	12.12	12.12
	mean HR	0.99	0.99	1.00	1.02	1.01	1.00
	std HR	0.02	0.02	0.02	0.02	0.02	0.02
Sugar	OLS HP		2.20	2.17	2.15	2.16	2.17
(204, 103)	OHR HP		2.21	2.16	2.08	2.12	2.14
	mean HR	0.76	0.75	0.77	0.80	0.79	0.78
	std HR	0.01	0.01	0.01	0.02	0.01	0.01

Note. The values in parentheses below the series labels denote the number of in-sample and out-of-sample observations, respectively. Below that, means and standard deviations of the OLS hedge ratios are also presented in the first column. The entries for each asset in the remaining columns give first the hedging potentials (welfare gains) of the OLS and of the utility-based hedges, respectively, followed by the means and standard deviations of the hedge ratios for the OLS and utility-based hedges. Parameter γ determines the shape of HARA utility; see Equation (8) and Appendix B. OHR, optimal hedge ratio; HARA, hyperbolic absolute risk aversion; HP, hedging potential.

$40 million in monetary terms, which cannot be considered small by any standards. It is also important to bear in mind that λ can be simply large because the local risk aversion of the hedger is very high, irrespective of the size of the actual physical exposure to the commodity.

4.1. Moderate Commodity Exposure, $\lambda = 1$

What increase in welfare can be achieved if one uses historical hedge ratios to determine appropriate hedging strategies for future time periods? Hedgers are

assumed to update their information sets once a month and to re-estimate their OHRs accordingly. The new hedge ratios are then used as a basis for risk management over the following month. We calculate the resulting time series of returns according to Equation (1). The out-of-sample (ex post) HPs generated from different utility functions are reported in Table VIII for both the case when we inappropriately use the OLS hedge (OLS HP) and the case when we use the utility-based optimal hedge (OHR HP).

For moderate commodity exposures, there is some evidence ex post that the investment potential of the optimal hedge exceeds that of the OLS hedge. This is the case for example for corn, where the HPs of the OLS hedge average 17.9 across utility functions, whereas the HPs of the utility-based hedge are slightly higher at an average of 18.1. This suggests that modeling the hedge ratios with the true distribution and thus taking into account higher moments increases the welfare of the hedger out-of-sample by an average of 0.15% a year. The increase in welfare is more noticeable for a HARA utility with $\gamma = 1$ (for which the OHR HP exceeds the OLS HP by 0.45% a year). In this case, adopting a more sophisticated approach to determining the hedge ratio helps as it increases welfare by an incremental average of 0.45% a year compared with the OLS hedge. The results presented in Table VIII convey a similar picture for heating oil and silver, where the increases in welfare generated from modeling higher moments equal 0.37 and 0.34% a year, respectively, across the eight utility functions considered. Across commodities and utility functions, the maximum increase in welfare generated from explicitly taking higher moments into account is obtained for silver for a hedger with a HARA utility and $\gamma = 1$. In that case, utility-based hedging increases welfare by a substantial 0.75% a year.

The results for gold and soybean oil are not as pronounced, with increases in welfare that average 0.05% a year across utility functions. So for these two commodities, there is only a slight decrease in wealth that occurs from using OLS hedging. Finally, in the case of cotton, the HP of the OLS hedge ratio exceeds that of the OHR by 0.81% a year across the eight utility functions we considered. The results are particularly dramatic for a hedger with a HARA utility and $\gamma = 1$ (hedging with higher moments then decreases wealth by 2.33% a year). This suggests that, in this case, anything more sophisticated than OLS hedging decreases welfare relative to the simpler OLS hedge. The case of cotton is of interest as it highlights the limitations of the utility-based hedge ratio which works well when departures from normality are not too pronounced. In the case of cotton, the higher moment hedge ratio fails to improve welfare as the return distribution of cotton presents an extremely negative skew (Table VII). In normal circumstances, the extreme event that caused the negative skew[21] is

---

[21]The negative skew was due to a very sharp decrease in cotton prices in 1986 that followed incentives from the Chinese government to stimulate production.

unlikely to reoccur, making our higher moment hedge ratio too conservative and thus unsuitable for hedging normal price exposure. A hedger fearing a sharp drop in cotton prices (similar to that of the mid 1980s) would, however, be well advised to adopt the utility-based hedge ratio.

Bringing together the evidence of Table VIII, it appears in most cases that modeling the hedge ratios with the true distribution and thus taking into account higher moments does increase the out-of-sample welfare of a hedger with moderate commodity exposure. To put it differently, there is, for most commodities (such as silver, heating oil and corn mainly, but also for gold and soybean oil) some systematic loss in wealth that occurs from inappropriately using OLS hedging.

All else equal, a hedge ratio that is stable over time is preferable to one that is highly volatile in order to keep the transactions costs from rebalancing the hedged portfolio to a minimum. In order to investigate the variability of the estimated hedge ratios from the various techniques, Table VIII also reports the means and standard deviations of the estimated 1-step ahead rolling hedge ratios. The means of the utility-based OHRs are bigger than the means of the OLS hedge ratios for five of the seven spot series hedged (corn, gold, silver, soybean oil and sugar), while they are smaller for the remaining two (cotton and heating oil). Thus, most of the time, switching to a utility-based approach that explicitly incorporates higher moments leads to higher hedge ratios, commensurate with a more precise estimate of the risks associated with systematically leptokurtic return distributions. In all cases, OLS-based hedging yields hedge ratios that have slightly lower variances, indicating more stable hedge ratios and therefore a lower cost of hedging.

In order to examine the relative sizes and stabilities of the estimated hedge ratios, Figures 1 and 2 plot the predictive hedge ratios implied by OLS and various utility functions in the HARA class. The hedge ratios are estimated recursively using all in-sample data, with one observation added at each time step, for cotton and gold, respectively. Figure 1 shows that in the case of cotton, the OLS hedge ratio is higher and less variable than those estimated from HARA utility functions, and in particular, logarithmic utility generates a dynamic OHR that has a lower mean but much higher variance than the others. Similarly, for gold (Figure 2), the OLS hedge ratio is much less volatile than that of the other utility functions (although now the OLS hedge also has a lower average value). This increased variability of the utility-based hedge ratios suggests that more frequent rebalancing of the hedged portfolio would be required, which could have consequences for the cost of implementing the hedges. In order to investigate this issue, we repeat the analysis of Table VIII but now computing the HP based on transactions-cost adjusted returns. Following Locke and Venkatesh (1997), we assume transactions costs
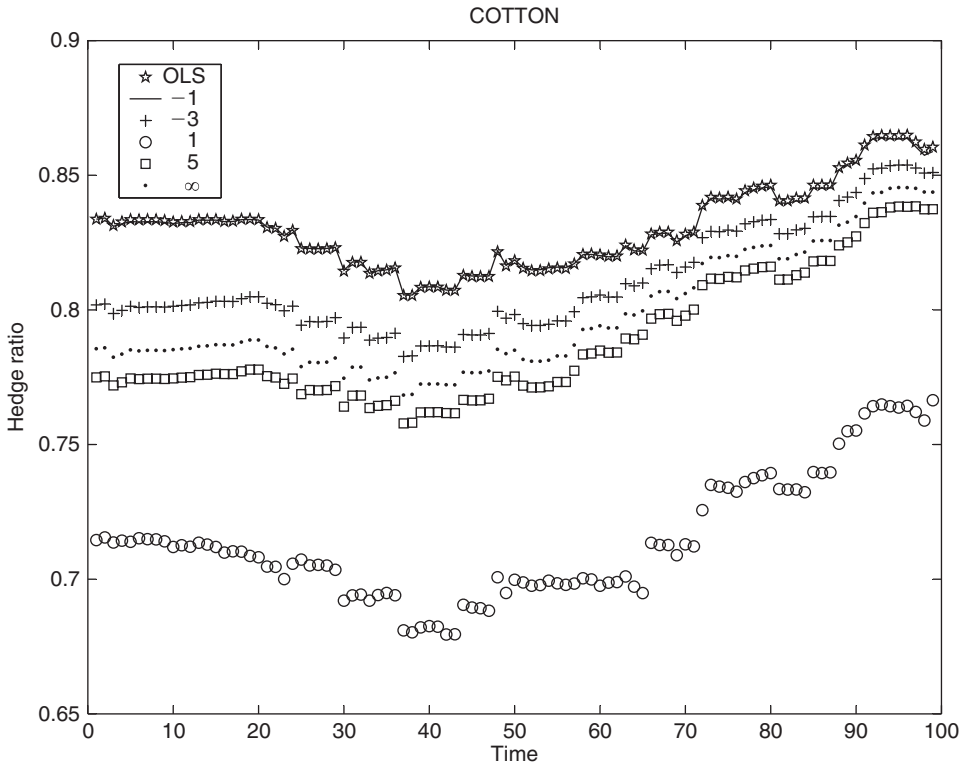
FIGURE 1

Out-of-sample hedge ratios for cotton, $\lambda = 1$, for the OLS and OHRs, $\gamma = -1, -3, 1, 5, \infty$. OHR, optimal hedge ratio.

of 0.033% per round-trip trade. The net of cost results are presented in Table IX.

HP is typically reduced by around 0.06, but the reduction is smaller for series where the hedge ratio is more stable (e.g. silver) and larger, perhaps up to 0.11, for series where it is more volatile (e.g. heating oil). The reduction in HP is, as expected, almost always larger for optimal hedging than OLS hedging. But the difference is very small and for no values of γ and for none of the commodities does the relative merit of one approach over the other qualitatively alter. The reduction in HP averaged across all commodities and all values of γ is 0.048 for OLS hedging and 0.056 for optimal hedging. Thus, we conclude from this analysis that hedging using all moments is certainly feasible from a practical perspective.

Figure 3 shows how the OHR varies with λ for the OLS, three-moment ($\gamma = -2$) and four-moment ($\gamma = -3$) cases for cotton.[22] The figure illustrates

[22]Similar plots for all seven of the commodity series examined in this study are available from the authors upon request but are not presented due to space constraints.
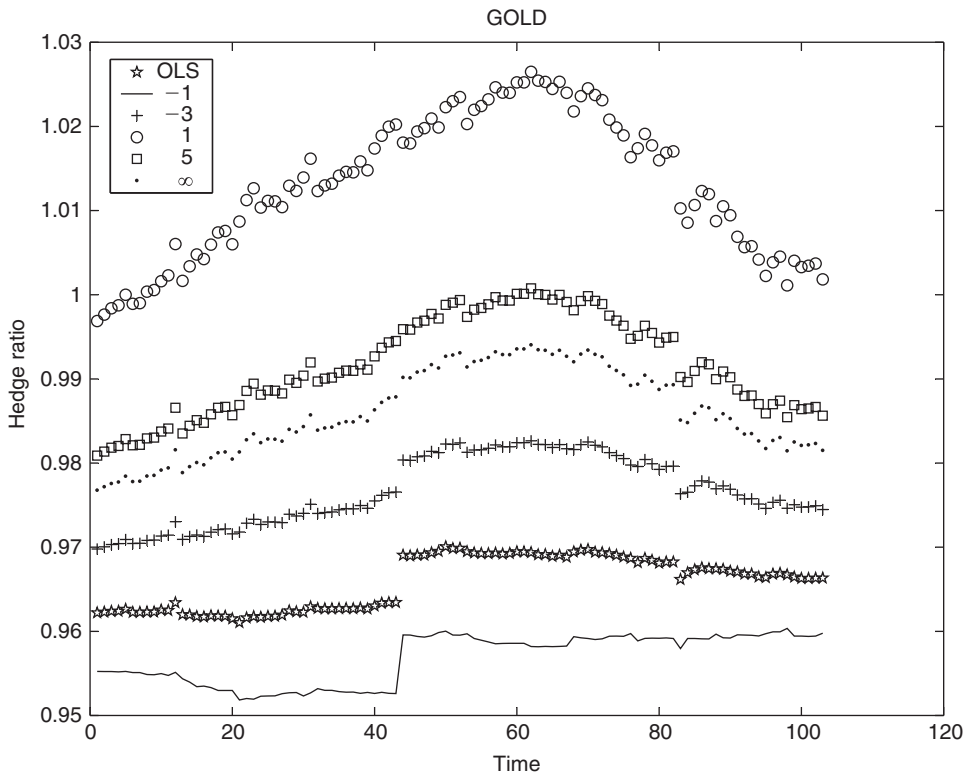
FIGURE 2

Out-of-sample hedge ratios for gold, $\lambda = 1$ for the OLS and OHRs, $\gamma = -1. -3, 1, 5, \infty$. OHR, optimal hedge ratio.

that for small values of λ, there is very little difference between the OLS hedge ratio and the three- and four-moment utility-based hedge ratios. However, as λ rises toward 10 (and therefore the normalised exposure to the spot asset grows), the OHR for the hedges allowing for higher moments diverge from the OLS hedge ratios. The impact of incorporating kurtosis into the mix when moving from the three-moment to the four-moment utility makes very little difference to the OHR, although the OHR is heavily dependent on λ.

Table X shows the first four moments of the hedged portfolio returns out-of-sample. Comparing the results with the no hedge case, both the OLS and the utility function-based hedges successfully reduce the variance of the hedged portfolio returns—sometimes moderately and sometimes spectacularly (e.g. gold). Also in this table, for comparison we report the summary statistics of the hedged portfolio returns when a multivariate GARCH (diagonal VECH) model is employed.

Perhaps precisely because by design OLS will minimize the (in-sample) variance of the hedged portfolio returns, it also results in out-of-sample

TABLE IX

Out-of-Sample Hedging Performance Net of Transactions Costs, $\lambda = 1$

Commodity (in, out)		γ				
		-1	-3	1	5	∞
Corn	OLS HP	17.77	17.81	17.92	17.85	17.83
(204, 103)	OHR HP	17.63	17.83	18.33	18.01	17.94
Cotton	OLS HP	19.78	19.76	19.76	19.76	19.76
(196, 99)	OHR HP	19.78	19.35	17.40	18.90	19.09
Gold	OLS HP	3.62	3.61	3.60	3.61	3.61
(204, 103)	OHR HP	3.60	3.63	3.66	3.64	3.64
Heating oil	OLS HP	-11.71	-11.86	-12.04	-11.95	-11.92
(180, 91)	OHR HP	-11.65	-11.51	-11.50	-11.47	-11.47
Silver	OLS HP	14.09	14.10	14.12	14.10	14.10
(170, 86)	OHR HP	13.89	14.27	14.85	14.53	14.43
Soybean oil	OLS HP	12.11	12.05	12.01	12.02	12.03
(204, 103)	OHR HP	12.07	12.07	12.12	12.08	12.07
Sugar	OLS HP	2.17	2.15	2.12	2.13	2.14
(204, 103)	OHR HP	2.19	2.12	2.02	2.08	2.10

Note. The values in parentheses below the series labels denote the number of in-sample and out-of-sample observations, respectively. The entries for each asset in the remaining columns give first the hedging potentials (welfare gains) of the OLS and of the utility-based hedges, respectively. Transactions costs of 0.033% per round trip have been deducted. Parameter γ determines the shape of HARA utility; see Equation (8) and Appendix B. OHR, optimal hedge ratio; HARA, hyperbolic absolute risk aversion; HP, hedging potential.

portfolio variances that are often lower than those of the utility-based hedges. The multivariate GARCH results also confirm the findings of the previous literature that such models are unable to outperform time-invariant hedges out-of-sample. For only two of the seven series does the MGARCH approach yield a lower out-of-sample return variance than that of the OLS hedge. In some cases, such as corn and cotton, the OLS approach is considerably superior, while for heating oil it is considerably inferior.

Also of interest are the impacts of hedging using the various methods on the higher moments of the hedged portfolio returns out-of-sample. For corn, cotton, heating oil, silver, and especially gold, there is a marked difference between the two. Focusing on the gold case, use of the HARA utility function with $\gamma = 1$ gives a return distribution with higher mean, lower variance, lower kurtosis, but also lower skewness and a larger minimum loss than the corresponding OLS hedge.[23]

To test the robustness of our findings to the sample analyzed, we replicate the analysis performed in Tables VI,VII,VIII, and X with a shorter data set that

[23]Given that the MGARCH model fails to provide better hedges or noticeable differences in the higher moments of the hedged portfolios compared with OLS, in the interests of brevity our comparison proceeds only with the OLS and utility-based hedges.
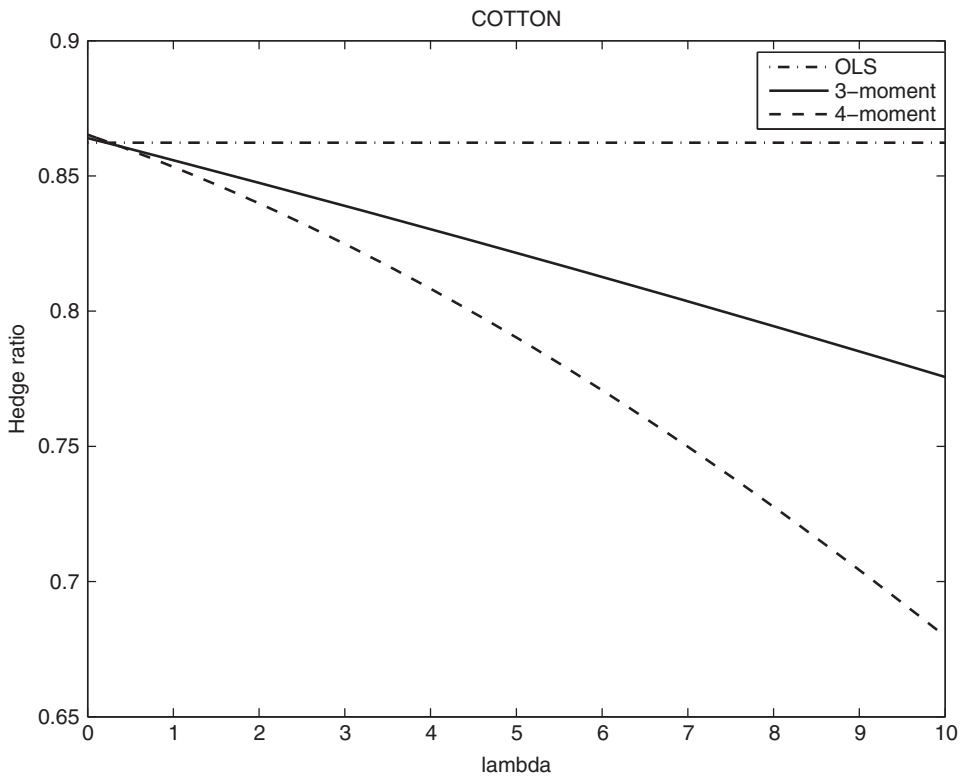
FIGURE 3
OHRs for different values of λ for cotton. The three and four-moment hedge ratio corresponds to HARA utility with $\gamma = -2$ and -3, respectively. OHR, optimal hedge ratio.

excludes the first five years of data. The results are qualitatively the same and available from the authors upon request. Again, with the noticeable exception of cotton, the conclusion from Table VIII of an overall increase in welfare when higher moments are modeled is robust to the sample analyzed. This suggests that the results are not sensitive to any learning that may have taken place in the first five years of our data set.

4.2. Large Commodity Exposure, $\lambda = 10$

Table XI shows the hedging performance for the OLS and utility-based hedge ratios when commodity exposure is large ($\lambda = 10$). Almost irrespective of the utility function considered, there is a clear tendency for the HP of the OHR to exceed the HP of OLS for five of our seven commodities. This is the case for gold, heating oil, silver, soybean oil, and sugar. The differences are very small in terms of HP in most cases. However, when measured in terms of the change

TABLE X

Moments of Hedged Commodity Returns, Out-of-Sample Results, λ = 1.

Commodity	No Hedge	OLS	MGARCH	γ				
				−1	−3	1	5	∞
Corn								
Mean (% p.a.)	−5.11	9.59	17.83	9.46	9.62	10.03	9.76	9.70
Std (% p.a.)	97.37	42.98	123.57	43.16	43.00	42.62	42.86	42.91
Skew	−0.45	−1.54	0.66	−1.52	−1.55	−1.61	−1.57	−1.56
Kurt	0.05	6.26	−0.64	6.11	6.27	6.64	6.40	6.35
Min (% p.a.)	−22.66	−17.36	−11.28	−17.33	−17.38	−17.47	−17.41	−17.40
Cotton								
Mean (% p.a.)	−2.83	13.50	4.56	13.49	13.10	11.30	12.66	12.85
Std (% p.a.)	102.49	43.68	70.86	43.70	43.92	47.23	44.37	44.15
Skew	0.45	0.36	−0.46	0.36	0.38	0.49	0.41	0.40
Kurt	0.04	1.02	2.45	1.02	0.96	0.54	0.88	0.92
Min (% p.a.)	−18.64	−7.51	−23.08	−7.51	−7.40	−7.46	−7.31	−7.35
Gold								
Mean (% p.a.)	1.37	4.11	4.21	4.09	4.13	4.21	4.17	4.15
Std (% p.a.)	47.13	8.13	8.27	8.21	8.07	8.08	8.03	8.03
Skew	1.03	1.09	0.90	1.24	0.90	0.35	0.66	0.75
Kurt	2.99	4.27	3.52	5.07	3.38	1.32	2.34	2.70
Min (% p.a.)	−8.82	−1.21	−1.39	−1.13	−1.32	−1.66	−1.46	−1.40
Heating oil								
Mean (% p.a.)	20.26	3.80	1.24	3.85	4.16	4.32	4.28	4.24
Std (% p.a.)	142.73	94.05	66.44	94.04	94.02	94.03	94.03	94.02
Skew	0.28	−0.20	0.19	−0.20	−0.18	−0.17	−0.17	−0.17
Kurt	0.71	1.99	3.38	1.96	1.86	1.82	1.82	1.83
Min (% p.a.)	−31.18	−24.24	−15.51	−24.12	−23.47	−23.08	−23.20	−23.28
Silver								
Mean (% p.a.)	−0.11	12.26	11.52	12.05	12.45	13.05	12.71	12.62
Std (% p.a.)	69.71	24.97	20.80	24.67	25.30	26.65	25.84	25.63
Skew	0.01	2.04	2.17	2.07	1.98	1.78	1.89	1.92
Kurt	0.85	10.30	14.36	10.36	10.13	9.28	9.80	9.93
Min (% p.a.)	−19.19	−5.26	−4.23	−5.16	−5.37	−5.67	−5.50	−5.45
Soybean oil								
Mean (% p.a.)	4.82	12.61	12.58	12.58	12.62	12.71	12.65	12.64
Std (% p.a.)	109.24	37.04	37.31	37.24	36.96	36.36	36.73	36.82
Skew	0.54	0.96	0.99	0.98	0.95	0.87	0.92	0.93
Kurt	1.46	2.94	3.17	3.00	2.91	2.66	2.82	2.86
Min (% p.a.)	−22.73	−5.60	−5.60	−5.62	−5.60	−5.53	−5.57	−5.58
Sugar								
Mean (% p.a.)	−0.02	−0.18	0.63	−0.16	−0.18	−0.21	−0.20	−0.19
Std (% p.a.)	93.40	54.88	57.24	54.81	55.01	55.49	55.20	55.13
Skew	0.27	−0.42	−0.58	−0.40	−0.46	−0.55	−0.50	−0.48
Kurt	−0.23	2.31	2.93	2.26	2.39	2.56	2.47	2.44
Min (% p.a.)	−16.09	−17.15	−17.96	−17.06	−17.26	−17.55	−17.39	−17.34

Note. The entries for each asset show in rows the first four sample moments of realized returns on (un)hedged portfolios, followed by a row specifying the lowest monthly return of the (un)hedged position. Columns correspond to no hedge, OLS hedge, multivariate GARCH hedge, and utility-based hedge ratios for different utility functions. Parameter γ determines the shape of HARA utility; see Equation (8) and Appendix B. OHR, optimal hedge ratio; HARA, hyperbolic absolute risk aversion; HP, hedging potential.

TABLE XI
Out-of-Sample Hedging Performance, $\lambda = 10$.

Commodity		OLS	γ				
			-1	-3	1	5	∞
Corn	OLS HP		3.9846	4.3069	–	5.6320	4.8545
(204, 103)	OHR HP		3.9660	4.3064	–	5.5811	4.8392
	mean HR	1.0014	0.9900	1.0049	–	0.9901	1.0020
	std HR	0.0084	0.0060	0.0070	–	0.0103	0.0081
Cotton	OLS HP		4.4615	4.4077	–	5.0336	4.6942
(196, 99)	OHR HP		4.4608	3.5488	–	$-\infty$	−5.2704
	mean HR	0.8306	0.8303	0.5598	–	−1.5309	−0.5499
	std HR	0.0156	0.0152	0.0599	–	0.1116	0.1006
Gold	OLS HP		1.1396	1.0717	–	1.0737	1.0693
(204, 103)	OHR HP		1.1369	1.0737	–	1.0773	1.0724
	mean HR	0.9658	0.9568	0.9750	–	0.9836	0.9809
	std HR	0.0031	0.0030	0.0043	–	0.0056	0.0051
Heating oil	OLS HP		2.4318	1.6983	–	4.2407	2.2051
(180, 91)	OHR HP		2.4362	1.9128	–	4.1995	2.5295
	mean HR	0.8285	0.8226	0.7251	–	0.8593	0.7488
	std HR	0.0106	0.0111	0.0120	–	0.0134	0.0108
Silver	OLS HP		2.8108	2.8868	–	3.3040	3.0663
(170, 86)	OHR HP		2.7975	2.9032	–	3.3306	3.0897
	mean HR	0.9401	0.9233	0.9632	–	0.9787	0.9736
	std HR	0.0073	0.0071	0.0119	–	0.0140	0.0133
Soybean oil	OLS HP		4.4840	4.2833	–	5.5855	4.7980
(204, 103)	OHR HP		4.4749	4.2872	–	5.5923	4.8043
	mean HR	0.9940	0.9874	0.9957	–	0.9983	0.9978
	std HR	0.0169	0.0162	0.0145	–	0.0138	0.0141
Sugar	OLS HP		1.9795	1.8440	–	1.9804	1.9138
(204, 103)	OHR HP		1.9835	1.8428	–	2.1106	1.9499
	mean HR	0.7600	0.7505	0.7657	–	0.6632	0.7332
	std HR	0.0076	0.0062	0.0158	–	0.0104	0.0172

Note. The values in parentheses below the series labels denote the number of in-sample and out-of-sample observations, respectively. Below that, means and standard deviations of the OLS hedge ratios are also presented in the first column. The entries for each asset in the remaining columns give first the hedging potentials (welfare gains) of the OLS and of the utility-based hedges, respectively, followed by the mean and standard deviations of the hedge ratios for the OLS and utility-based hedges. Parameter γ determines the shape of HARA utility; see Equation (8) and Appendix B. OHR, optimal hedge ratio; HARA, hyperbolic absolute risk aversion; HP, hedging potential.

in CE wealth, the differences can be economically important (as illustrated in the example of the airline company in Tables III and IV).

The opposite applies to the remaining two commodities, for which explicitly modeling higher moments is detrimental to utility, either marginally (in the case of corn) or substantially (in the case of cotton). Hence, considering a high value of normalized exposure ($\lambda = 10$), it can be seen that paying too much attention to higher moments can be counterproductive. In the case of cotton, for example, the exponential utility hedge ratio is far inferior to the OLS hedge, as reflected by a negative OHR HP.

The example is symptomatic of a more general issue that is pervasive in finance. Suppose that the data-generating process behind the spot and futures returns for cotton is accurately represented by the historical data. Suppose further that the manager of company A ignores extreme negative returns in the past data, arguing that such extreme events are unlikely to occur again in future. Consequently, manager A selects a hedging strategy that very much resembles an OLS hedge. By contrast, the manager of company B heeds the warning issued by the data and selects a hedging strategy leaning toward the minimax hedge ratio. The problem for manager B is that until the extreme scenario captured in the historical data repeats itself, his strategy will be extremely costly. For cotton, this is exemplified in Table XI by HARA utility-based hedging strategies with $\gamma = 5, \infty$ in particular, although the loss in performance is visible for all nonquadratic utility-based strategies.

As we did at the end in Section 4.1, we reproduce the analysis of Table XI over a shorter sample that excludes the first five years of data. This is to test the robustness of our inferences as learning that took place in the early days of trading might have had an impact on the results obtained. Irrespective of the sample analyzed, there are WGs to be earned from adopting a more sophisticated hedging approach for some commodities (gold, heating oil, silver, and sugar). However, the conclusion does not apply throughout as over the shorter period, modeling higher moments can also be damaging (again, especially for cotton).[24]

5. CONCLUSIONS

This study has proposed a utility-based framework for the determination of OHRs that can allow for the impact of higher moments on the hedging decision. The approach is illustrated using the example of an airline hedging its fuel exposure and is then applied to a set of seven commodities that are hedged with futures contracts. The derivation of a general utility-based hedging framework that allows for higher moments is well motivated given the non-normality in most financial asset return series and the empirical observation in the literature that conventional approaches can lead to hedged portfolio returns with less attractive skewness and kurtosis properties than if no hedge had been implemented at all.

[24]Tables VIII and XI present the WGs/losses that can be obtained from modeling higher moments for two levels of commodity exposure, a moderate level $\lambda = 1$ and a high level $\lambda = 10$. We measure for two commodities (corn and cotton) the out-of-sample hedging performance of the different hedge ratios with values of λ ranging from two to eight in increments of 2. The results (available upon request from the authors) are robust to the commodity exposure λ since, as before in Tables VIII and XI, modeling higher moments is found to increase welfare for corn and decrease welfare for cotton, although in the case of corn, the improvement of higher moment hedging over OLS reduces as λ increases.

We find that for both moderate and large spot (commodity) exposures, the performance of out-of-sample hedges constructed allowing for nonzero higher moments is better than the performance of the simpler OLS hedge ratio. Hence, it seems that for most of our cross section, higher moments do matter when it comes to determining OHRs. The picture is, however, not uniform across our seven spot commodities as there is one commodity (cotton) for which the modeling of higher moments decreases welfare out-of-sample relative to the simpler OLS approach. We attribute the lack of performance of the higher moment hedge ratio to the very negative skew of the cotton spot returns (driven by the sharp rise in stocks in the mid 1980s). This case highlights the limitation of the utility-based hedge ratio that works better when departures from normality are not too extreme.

It would be a useful step for future research to determine whether our broad conclusions also hold for other hedging assets, sample periods, and data frequencies. As the framework proposed here is sufficiently general, it seems also of interest to apply it to the type of strategic asset allocation problems that equity investors face.

APPENDIX A: NUMERICAL ALGORITHM

The problem

$$\alpha = \arg \max_{\vartheta \in R^n} E(f(Y + \vartheta X))$$

$$a = f^{-1}(E(f(Y + \alpha X)))$$

can be solved by Newton's iteration method provided that the initial guess ϑ_0 is close to the optimal portfolio α. In practice, the quadratic approximation $\vartheta_0 = -E(XY)/E(X^2)$ works very well. We define $g: \mathbb{R} \rightarrow \mathbb{R} \cup \{-\infty\}$

$$g(\vartheta) = E(f(Y + \vartheta X))$$

and assume that in each iteration $g(\vartheta) > -\infty$. Starting at ϑ_0 we use the iteration

$$\vartheta_{k+1} = \vartheta_k - \frac{g'(\vartheta_k)}{g''(\vartheta_k)}$$

where

$$g'(\vartheta) = E(Xf'(Y + \vartheta X))$$

$$g''(\vartheta) = E(X^2 f''(Y + \vartheta X)).$$

Assuming sufficient smoothness of f the Taylor expansion yields

$$|f^{-1}(g(\vartheta_k)) - f^{-1}(g(\alpha))| = -\frac{1}{2f'(f^{-1}(g(\vartheta_k)))}\frac{(g'(\vartheta_k))^2}{g''(\vartheta_k)} + o((\vartheta_k - \alpha)^2). \quad \text{(A1)}$$

Accordingly, we stop the iteration when

$$-\frac{(g'(\vartheta_k))^2}{f'(f^{-1}(g(\vartheta_k)))g''(\vartheta_k)} < 10^{-12}$$

which in practice guarantees $|f^{-1}(g(\vartheta_k)) - f^{-1}(g(\alpha))| < 10^{-12}$. The last inequality means a very close proximity of the final iterate to the optimal hedging decision in terms of the resulting CE. Equation (A1) hints that $|\vartheta_k - \alpha|$ has half the number of zeros compared with the target function, i.e. in practice the final iterate satisfies $|\vartheta_k - \alpha| \approx 10^{-6}$. Rigorous proof of the quadratic convergence can be found, for example, in Dennis and Schnabel (1996).

APPENDIX B: HARA UTILITY

The HARA family is described in Cass and Stiglitz (1970). We use a slight modification of the parametrization suggested in Ingersoll (1987).

Definition 7.1: The utility function

$$U_\gamma(V; a, b) := \begin{cases} \dfrac{(aV/\gamma + b)^{1-\gamma} - 1}{1/\gamma - 1} & \text{for } \gamma > 0, \gamma \neq 1 \\[2mm] \ln(aV + b) & \text{for } \gamma = 1 \\[2mm] \dfrac{|aV/\gamma + b|^{1-\gamma} - 1}{1/\gamma - 1} & \text{for } \gamma < 0 \\[2mm] 1 - e^{-aV} & \text{for } \gamma = \pm\infty. \end{cases}$$

with $a > 0$ is called the hyperbolic absolute risk-aversion utility. We denote the corresponding effective domain by $\mathcal{D}^{(\gamma)}(a, b)$ and the maximal domain on which U_γ is increasing by $\overline{\mathcal{D}}_\gamma(a, b)$.

The HARA utility is an infinitely differentiable utility in the sense of Definition 2.1. U_γ is strictly increasing and unbounded from above for $\gamma \in (0, 1]$; it is strictly increasing and bounded from above for $\gamma > 1$ and for $\gamma < 0$ it has a bliss point at $-\gamma b/a$. For $\gamma \in \mathbb{R}\setminus\{0\}$, the coefficient of absolute risk aversion at $v \in \overline{\mathcal{D}}_\gamma(a, b)$ reads

$$A_\gamma(v; a, b) = \frac{1}{v/\gamma + b/a}$$

hence the acronym HARA. The HARA class has several advantages over the more frequently used power utility functions. Fixing a positive initial wealth level v one can, with an appropriate choice of a, b, make the HARA utility increasing at v even when $\gamma < 0$. Second, power utility ($\gamma > 0$, $b = 0$) produces unreasonable levels of risk aversion for large values of γ This can be corrected in the HARA class by selecting an appropriate value of $b > 0$.

Proposition 7.2: Fix $\gamma \in \mathbb{R} \cup \{\pm\infty\}$, $\gamma \neq 0$, $a > 0$, $b \in \mathbb{R}$ and $v \in \overline{D}_\gamma(a, b)$. Then f_γ in Equation (8) is a normalized utility to $U_\gamma(\cdot; a, b)$ at v, in the sense of Definition 2.4. Consequently the normalized utility is independent of the specific values of a, b, and v.

APPENDIX C: PROOFS

Proof of Theorem 2.3: See Černý (2003), Theorem 2.

Proof of Theorem 2.5: By a straightforward calculation

$$A(v)CE(v, \eta, \vartheta) = f^{-1}(E(f(\lambda Y + \theta X))) = a(\lambda, \theta) \tag{C1}$$

for *any* normalized utility f, with θ given in Equation (4). From here and (6) follows

$$\varphi(v, \eta) = \alpha(\lambda)/A(v),$$

$$OHR(v, \eta) = \frac{\varphi(v, \eta) - \varphi(v, 0)}{\eta} = \frac{\alpha(\lambda) - \alpha(0)}{\lambda} = \hat{h}(\lambda). \tag{C2}$$

Equation (C1) implies $\varphi(v, 0) = \alpha(0)/A(v)$. This together with (4) and (C1) yields

$$\frac{CE(v, \eta, \varphi(v, 0) - h\eta)}{A(v)\eta^2} = \frac{f^{-1}(E(f(\lambda Y + (\alpha(0) - h\lambda)X)))}{\lambda^2}$$

$$\frac{CE(v, \eta, \varphi(v, 0))}{A(v)\eta^2} = \frac{f^{-1}(E(f(\lambda Y + \alpha(0)X)))}{\lambda^2}$$

whereby we obtain (12) from (3) and (5). The existence and uniqueness of the maximizer in (10) was shown in Theorem 2.3.

APPENDIX D: OPTIMAL HEDGING AND OLS

Assuming sufficient smoothness ($f \in C^2$), the quantity $\alpha(\lambda)$ is differentiable and we can think of the optimal hedge $\hat{h}(\lambda)$ as the *average* value of the marginal hedge ratios $-\alpha'(s)$ with $s \in [0, \lambda]$.

$$\hat{h}(\lambda) = -\frac{\displaystyle\int_0^\lambda \alpha'(s)ds}{\lambda}.$$

By differentiating the first order condition $E(Xf'(\alpha(\lambda)X + \lambda Y)) = 0$ with respect to λ we have

$$E(X^2 f''(\lambda Y + \alpha(\lambda)X))\alpha'(\lambda) = -E(XYf''(\lambda Y + \alpha(\lambda)X))$$

$$\alpha'(\lambda) = -\frac{E(XYf''(\lambda Y + \alpha(\lambda)X))}{E(X^2 f''(\lambda Y + \alpha(\lambda)X))}.$$

In the special case $f'' =$ const, corresponding to quadratic utility, we obtain

$$\alpha'(\lambda) = -\frac{E(XY)}{E(X^2)}$$

which means that $\hat{h}(\lambda)$ is independent of λ. If, in addition, the mean of X is zero (the futures market is unbiased) then the quadratic hedge equals the slope coefficient from the OLS regression of Y onto X and an intercept. For other utility functions, the choice of λ matters to some extent, but our empirical results show that this dependence is extremely weak for $\lambda \in [0, 1]$.

BIBLIOGRAPHY

Agarwal, V., & Naik, N. (2004). Risks and portfolio decisions involving hedge funds. Review of Financial Studies, 17, 63–98.

Arrow, K. J. (1963) Comments on Duesenberry's the portfolio approach to the demand for money and other assets. Review of Economics and Statistics, 45(1), 24–27.

Arrow, K. J. (1971). Essays in the theory of risk-bearing.Chicago: Markham.

Barone-Adesi, G., Gagliardini, P., & Urga, G. (2004), Testing asset pricing models with coskewness. Journal of Business and Economic Statistics, 22(4), 474–485.

Benninga, S., Eldor, R., & Zilcha, I. (1983). The optimal hedging in the futures market under price uncertainty. Economics Letters, 13, 141–145.

Brys, E., Crouhy, M., & Schlesinger, H. (1993). Optimal hedging in a futures market with background noise and basis risk. European Economic Review, 37, 949–960.

Brooks, C., Henry, O. T., & Persand, G. (2002). The effect of asymmetries on optimal hedge ratios. Journal of Business, 75(2), 333–352.

Brooks, C., & Kat, H. (2002) The statistical properties of hedge fund index returns and their implications for investors. Journal of Alternative Investments, 5(2), 26–44.

Brown, G. W., & Toft, K. B. (2002) How firms should hedge. Review of Financial Studies, 15(4), 1283–1324.

Cass, D., & Stiglitz, J. E. (1970) The structure of investor preferences and asset returns, and separability in portfolio allocation: a contribution to the pure theory of mutual funds. Journal of Economic Theory, 2(2), 122–160.

Cechetti, S. G., Cumby, R. E., & Figlewski, S. (1988) Estimation of optimal futures hedge. Review of Economics and Statistics, 70, 623–630.

Černý, A. (2003). Generalized sharpe ratios and asset pricing in incomplete markets. European Finance Review, 7(2), 191–233.

Christie-David, R., & Chaudhry, M. (2001). Coskewness and cokurtosis in futures markets. Journal of Empirical Finance, 8, 55–81.

Chunhachinda, P., Dandapani, K., Hamid, S., & Prakash, A. (1997). Portfolio selection and skewness: evidence from international stock markets. Journal of Banking and Finance, 21, 143–167.

Dennis Jr., J. E., & Schnabel, R. B. (1996). Numerical methods for unconstrained optimization and nonlinear equations. *SIAM*.

Froot, K. A., Scharfstein, S., & Stein, J. (1993). Risk management: coordinating corporate investment and financing policies. Journal of Finance, 68, 1629–1658.

Gilbert, S., Jones, S. K., & Morris, G. H. (2006). The impact of skewness in the hedging decision. Journal of Futures Markets, 26(5), 503–520.

Harris, R. D. F., & Shen, J. (2003). Robust estimation of the optimal hedge ratio. Journal of Futures Markets, 23, 799–816.

Harris, R. D. F., & Shen, J. (2006). Hedging and value at risk. Journal of Futures Markets, 26(4), 369–390.

Harvey, C. (1995). Predictable risk and returns in emerging markets. Review of Financial Studies 8, 773–816.

Harvey, C. R., Liechty, J. C., Liechty, M. W., & Müller, P. (2010). Portfolio selection with higher moments. Quantitative Finance, 10, 469–485.

Harvey, C. R., & Siddique, A. (2000). Conditional skewness in asset pricing tests. Journal of Finance, 55, 1263–1295.

Hong, Y., Tu, J., & Zhou, G. (2007). Asymmetries in stock returns: statistical tests and economic evaluation. Review of Financial Studies, 20(5), 1547–1581.

Ingersoll, J. E. (1987). Theory of financial decision making. Savage, MD: Rowman & Littlefield.

Joanes, D. N., & Gill, C. A. (1998) Comparing measures of sample skewness and kurtosis. The Statistician, 47(1), 183–189.

Jondeau, E., & Rockinger, M. (2006). Optimal portfolio allocation under higher moments. European Financial Management, 12(1), 29–55.

Kallberg, J. G., & Ziemba, W. T. (1983). Comparison of alternative utility functions in portfolio selection problems. Management Science 29(11), 1257–1276.

Kraus, A., & Litzenberger, R. (1976). Skewness preference and the valuation of risky assets. Journal of Finance, 31, 1085–1100.

Kraus, A., & Litzenberger, R. (1983) On the distributional conditions for a consumption-oriented three moment CAPM. Journal of Finance, 38(5), 1381–1391.

Levy, H. (1969). A utility function depending on the first three moments. Journal of Finance, 24, 715–719.

Locke, P., & Venkatesh, P. (1997). Futures market transactions costs. Journal of Futures Markets, 17, 229–245.

Miffre, J. (2004). Conditional OLS minimum variance hedge ratios. Journal of Futures Markets, 24(10), 945–964.

Morrell, P., & Swan, W. (2006). Airline jet fuel hedging: theory and practice. Transport Reviews, 26(6), 713–730.

Moschini, G., & Lapan, H. (1995). The hedging role of options and futures under joint price, basis, and production risk. International Economic Review, 36, 1025–1049.

Post, T., Vliet, P., & Levy, H. (2008). Risk aversion and skewness preference, Journal of Banking and Finance, 32(7), 1178–1187.

Pratt, J. W. (1964). Risk aversion in the small and in the large. Econometrica, 32, 122–136.

Pulley, L. B. (1983). Mean-variance approximations to expected logarithmic utility. Operations Research, 31(4), 685–696.

Samuelson, P. (1970). The fundamental approximation theorem of portfolio analysis in terms of means, variances and higher moments. Review of Economic Studies, 37(4), 537–542.

Scott, R., & Horvath, P. (1980). On the direction of preference for moments of higher order than the variance. Journal of Finance, 35, 915–919.

Simaan, Y. (1993), What is the opportunity cost of mean-variance investment strategies? Management Science, 39(5), 578–587.

Yamada, Y., & Primbs, J. A. (2004). The effect of higher order moments on hedging loss value-at-risk and conditional value at risk. Proceedings of the 2nd IASTED International Conference: Financial Engineering and Applications (pp. 21–26), Cambridge, MA.