

The Stochastic Behavior of Commodity Prices with Heteroscedasticity in the Convenience Yield

Peng Liu¹

Cornell University

Ke Tang^{2*}

Renmin University of China

July 27, 2010

¹ Cornell University, 465 Statler Hall, Ithaca, NY 14850, United States, pl333@cornell.edu

² Hanqing Advanced Institute of Economics and Finance and School of Finance, Renmin University of China, Beijing, 100872, China, ketang@ruc.edu.cn

* *Corresponding author*

Author Note

We thank Han Hong, Tong Li, Jun Liu, Geert Rouwenhorst, Wei Xiong and Lu Zhang, seminar participants at 2010 Financial Management Association Annual Conference, the University of California at Berkeley and Renmin University of China for helpful comments.

We are also grateful to two anonymous referees whose suggestions greatly improved the paper. Liu acknowledges the research fund from Cornell University.

Tang acknowledges financial support from the Program for New Century Excellent Talents in University, and the Fundamental Research Funds for the Central Universities, and the Research Funds of Renmin University of China.

We are also grateful for helpful discussion with Andrew Harvey regarding the implementation of the Kalman filter.

Abstract

We document a new stylized fact regarding the dynamics of the commodity convenience yield: the volatility of the convenience yield is heteroskedastic for industrial commodities; specifically, the volatility (variance) of the convenience yield depends on the convenience yield level. To explore the economic and statistical significance of the improved specification of the convenience yield process, we propose an affine model with three state variables (log spot price, interest rate, and the convenience yield). Our model captures three important features of commodity futures—the heteroskedasticity of the convenience yield, the positive relationship between spot-price volatility and the convenience yield and the dependence of futures risk premium on the convenience yield. Moreover our model predicts an upward sloping implied volatility smile, commonly observed in commodity option market.

Keywords: commodity, convenience yield, heteroscedasticity, affine model, volatility smile

JEL Classification: G12, G13

The Stochastic Behavior of Commodity Prices with Heteroscedasticity in the Convenience Yield

The convenience yield, defined as the flow of benefit of immediate ownership of a physical commodity, is regarded as a distinguishing feature of commodities as an asset class. It is also considered as the key to modeling commodity futures prices. By assuming a stochastic latent convenience yield process, reduced-form contingent-claims models are often used for risk management and for real option valuation. However, since the convenience yield is not directly observable, distinct convenience yield models will result in varying results in contingent claim evaluation. Allowing the convenience yield to depend on the (log) spot price and the risk-free rate, Casassus and Collin-Dufresne (CCD, 2005) document that mis-specification of the convenience yield can have a significant impact on option valuation and risk management. Therefore, correctly modeling the convenience yield process is a significant step towards modeling commodity-related contingent claims.

In this paper, we empirically document a new feature of the convenience yield—heteroskedasticity. The volatility (variance) of the convenience yield depends on the convenience yield level. This is also consistent with equilibrium models in the theory of storage. For example, Deaton and Laroque (1992, 1996) and Routledge et al. (2000) indicate that the volatility of commodity prices is driven not only by the volatility of inventories but also by the non-linear response of prices to demand–supply shocks due to the non-negativity constraint on inventories. This tends to cause a non-linear relationship between the convenience yield and the inventory level, which is consistent with Fama and French (1987, 1988). As shown in Section 3 of this paper, this non-linear relationship results in the heteroskedasticity of the convenience yield.

The existence of heteroskedasticity in the convenience yield is quite important because understanding the dynamics of and relationships between the first and second moments of the convenience yield is crucial in modeling commodity futures prices and pricing commodity derivatives. In order to incorporate the heteroskedasticity discovered in our empirical work, we specify a three-factor ‘maximal’ affine framework to model the commodity price. In addition to the heteroskedasticity, our model also reflects the positive correlation between the convenience yield and spot price volatility, as empirically documented in Ng and Pirrong (1994).

Furthermore, Gorton et al. (2007) show that there is a strong positive relationship between the convenience yield and the futures risk premium. This is because a high convenience yield will result in high spot-price volatility, which in turn causes a high risk premium for futures. Similarly Dincerler et al. (2005) show a positive relationship between futures returns and inventory withdrawals, which also indicates a positive relationship between the convenience yield and futures risk premia from the theory of storage. Up to now, however, these features have not been incorporated into the reduced-form commodity pricing models.

Many scholars have modeled the stochastic behavior of the convenience yield. For example, Gibson and Schwartz (1990) use two factors—the log-spot price and the convenience yield—to model the movements of futures prices. The log-spot price is modeled as a Brownian motion and the convenience yield as an Ornstein–Uhlenbeck (OU) process. Schwartz (1997) introduces a third factor—stochastic interest rates, and shows that interest rate dynamics is important in valuing long-term commodity claims. Nielsen and Schwartz (2004) allow the volatility of the spot price depending on the convenience yield.¹

¹ In Nielsen and Schwartz (2004) the conditional correlation of spot volatility and convenience yield volatility is restricted to be one; thus this model is not general enough.

Beyond the empirical documentation of the heteroskedasticity of the convenience yield, we also build a “maximal” affine model which copes with the heteroskedasticity effect of the convenience yield. We then utilize oil and copper data to calibrate our model. The results show that both the heteroskedasticity of the convenience yield and the dependence of spot volatility on the convenience yield exist for copper and oil. Furthermore, using a likelihood ratio test, we show that our model is better than models that fail to specify convenience yield heteroskedasticity or the dependence of spot volatility on the convenience yield.

The rest of the paper is organized as follows. Section 2 examines the heteroskedasticity of the convenience yield empirically. Section 3 proposes an economic model to explain the stylized fact. Section 4 proposes a reduced-form model to model the futures dynamics based on the theory of storage and our empirical findings in Section 2. Section 5 discusses the model calibration results using oil and copper data. Section 6 proposes the implementation of option pricing as an application of our model. Section 7 concludes.

Empirical Tests of the Heteroscedasticity

In this section, we first test the heteroskedasticity of the convenience yield. We perform three heteroskedasticity tests using daily data on WTI crude oil and high-grade copper from the New York Mercantile Exchange (NYMEX). The data contain 1544 observations for each commodity from January 2000 to February 2006. The two statistical tests reject the homoscedasticity hypothesis at a 1% significance level. The results are robust across different industrial commodities. The theory of storage (such as Kaldor, 1939, Working, 1949, Brennan, 1958, and Telser, 1958) implies that the convenience yield depends mainly on commodity inventories. Hence, it is natural to examine if the heteroskedasticity also exists in commodity inventories.

Heteroskedasticity test of the convenience yield

Since convenience yield is not directly observable, we infer the implied convenience yield, $\delta(t, T_1, T_2)$ from commodity futures prices and the interest rates using the following equation,

$$\delta(t, T_1, T_2) = r_t - \frac{\ln F(t, T_2) - \ln S_t}{T_2} \approx r_t - \frac{\ln(F(t, T_2)) - \ln(F(t, T_1))}{T_2 - T_1} \quad (1)$$

where T_1 and T_2 are futures maturities, and r_t is the risk-free rate (here we use the three-month Treasury bill). Since good spot-price data are not available for most commodities, following Bessembinder et al. (1995) and others, we use futures prices on maturing contracts to represent spot prices.

To test the high convenience yield–high convenience yield volatility relationship (the heteroskedasticity of the convenience yield),² we thus run the following regressions³:

$$\delta(t, T_1, T_2) - \delta(t - 1, T_1, T_2) = \iota + \kappa \times \delta(t - 1, T_1, T_2) + \mu_t \quad (2)$$

$$\text{Breusch-Pagan Test : } \mu_t^2 = a + b \times \delta(t - 1, T_1, T_2) + \zeta_t \quad (3)$$

$$\text{Glesjer Test : } |\mu_t| = c + d \times \delta(t - 1, T_1, T_2) + \pi_t \quad (4)$$

$$\text{GARCH 11 Test : } \mu_t = \sqrt{h_t} \epsilon_t,$$

$$h_t = g_0 + g_1 h_{t-1} + g_2 \mu_{t-1}^2 \quad (5)$$

$$h_t = e + f \times \delta(t - 1, T_1, T_2) + \omega_t$$

where $\mu_t, \zeta_t, \pi_t, \epsilon_t$ and ω_t represent error terms at time t . Regressions (3) and (4) examine whether the convenience yield volatility at t , which is approximated by μ_t^2 or $|\mu_t|$, can be explained by the convenience yield level at $t-1$. However, using μ_t or $|\mu_t|$ to represent the

² Note that in this paper we are interested in the “cross-sectional” relationship between convenience yield volatility and the convenience yield level but not the autoregressive behavior of the convenience yield volatility.

³ See also Wooldridge (2003) and Goldfeld and Quandt (1972) for details of Breusch–Pagan and Glejser tests.

volatility of the convenience yield might be rough, we thus use a GARCH(1,1) model to estimate the latent “real” volatility and then analyze its relationship to the level of the convenience yield.⁴

We name this test as the GARCH11 test. Note that from the Durbin–Watson statistics, we usually observe weak autocorrelations in the regressions, we hence use the Newey–West (1987) method to correct them.

Using oil and copper futures prices we run regressions (2) to (5). From Table 1, the regression coefficients b , d and f are positive and highly significant from the Breusch–Pagan, Glejser and GARCH11 tests.⁵ This suggests that the volatility of the convenience yield is heteroskedastic. In particular, a high convenience yield level is associated with a high convenience yield volatility. To the best of our knowledge, we are the first to document this phenomenon explicitly. To check the robustness of our result, we also test three groups of other commodities: the oil products, heating oil and gasoline; the industrial metals, zinc and nickel; and the precious metals, gold and silver. We find that production commodities (such as oil products and industrial metals) exhibit heteroskedastic behavior similar to that of copper and oil. However, non-production commodities like silver and gold do not show heteroskedasticity. (See Appendix A, Table 5 for detailed results.)

Heteroskedasticity test of the inventory

Many studies have documented the importance of the inventory in explaining the behavior of the commodity convenience yield, so it is natural to test whether the volatility of inventory dynamics is heteroskedastic. To explore the heteroskedasticity of the inventory, we run the Breusch–Pagan, Glejser and GARCH11 tests (Eqs. (2) to (5) using inventory data on oil and

⁴ We thank the referee for suggesting us doing this test.

⁵ The only case when it is not significant in a 95% level (but in a 90% level) is when the oil convenience yield is calculated using 3 month futures.

copper.⁶ In Appendix A, we perform the same test using zinc, nickel, and gold data as a check for robustness. The insignificant b, d and f values in Table 1 show that the inventory variance (or volatility) of both copper and oil does not depend on the inventory level, which suggests that the heteroskedasticity does not exist for the inventory process. Table 6 in Appendix A shows the same results for inventories of zinc, nickel, and gold.

Therefore, the convenience yield process exhibits a high convenience yield–high convenience yield volatility relationship, however, inventory does not display this relationship. This rules out the possibility that the convenience yield heteroskedasticity is driven by heteroskedastic inventories. Section 3 shows that the nonlinear relationship between the convenience yield and inventory from the theory of storage will result in the heteroskedastic convenience yield and the homoskedastic inventory for industrial commodities. Since gold and silver are not mainly used for production and hence have a large amount of inventory, their convenience yields are very likely to be inelastic with inventory changes. Hence, the convenience yield heteroskedasticity does not exist in gold and silver.

Economic Explanation of the Heteroskedastic Convenience Yield and Homoscedastic Inventory

We build a simple model to demonstrate that the stylized fact explored in Section 2 is a direct implication of the theory of storage, developed by Brennan (1958), Telser (1958) and Fama and French (1987, 1988).

⁶ Weekly crude oil inventory data (excluding strategic petroleum reserves) are obtained from the U.S. Energy Information Administration for a period running from January 1995 to February 2006. The weekly copper stocks in warehouses are obtained from COMEX for a period running from February 1997 to February 2006. The units for oil and copper inventories are, respectively, 1 million bbl and 1 million tons.

The convenience yields link to two interconnected markets: the financial futures market and the storage market. In the futures markets, the fluctuation of the convenience yield is reflected by changes in the futures slope (i.e. contango and backwardation). This is shown in the left part of Fig. 1. In the storage market, the theory of storage considers the inventory level at t , I_t as the quantity of storage and the convenience yield as the price of storage. According to classical microeconomics theory, the convenience yield is a marginal utility gained from holding one more unit of inventory, and thus it should be a function of the inventory level, i.e. $\delta_t = f(I_t)$, Brennan (1958), Fama and French (1987) and Pindyck (2001) show the convenience yield δ_t as a decreasing convex function of inventory level I_t . That is, δ_t decreases as I_t increases, but at a decreasing rate, or $\frac{d\delta_t}{dI_t} < 0$ and $\frac{d^2\delta_t}{dI_t^2} > 0$. Deaton and Laroque (1992, 1996) and Routledge et al. (2000) indicate that the volatility of commodity prices is driven not only by the volatility of inventories but also by the non-linear response of prices to demand–supply shocks due to the non-negativity constraint on inventories. This indicates a non-linear relationship between the convenience yield and the inventory level. The right part of Fig. 1 shows the relationship between inventory and the convenience yield. This convex shape of the function in the storage market implies that, at high inventory levels, the convenience yield function is almost flat. There can be a large inventory response to a demand shock without a large change in the convenience yield. In contrast, at low inventory levels, the convenience yield rises much faster when inventory is used to meet an increase in demand. Using inventory data from 33 commodities, Gorton et al. (2007) confirmed the convex structural relationship between the convenience yield and the commodity inventory empirically.

We further specify the reduced-form stochastic process of inventory as follows:

$$dI_t = \zeta(I_t, \Xi_t)dt + \gamma(\Xi_t)dW_t, \quad (6)$$

where I_t denotes inventory at time t , Ξ_t denotes a set of exogenous variables, $\zeta(I_t, \Xi_t)$ is the drift part of the inventory, which depends on inventory and exogenous variables, $\gamma(\Xi_t) > 0$ is the volatility of the inventory, which depends only on exogenous variables, and W_t is a Wiener process. Note that the results of Section 2 suggest that the volatility of inventory does not depend on inventory itself and hence is independent of the convenience yields as well. The independence of the convenience yield volatility and inventory level is consistent with Routledge et al. (2000), in which the change of inventory (dI_t) conditional on the current inventory level (I_t) is related only to exogenous net demand shocks (modeled as a normally distributed variable).

Applying Ito's lemma, the process for the convenience yield therefore follows:

$$d\delta_t = \left[f'(I_t)\zeta(I_t, \Xi_t) + \frac{1}{2}f''(I_t)\gamma(\Xi_t)^2 \right] dt + \gamma(\Xi_t)f'(I_t)dw_t, \quad (7)$$

where $f'(I_t)$ and $f''(I_t)$, respectively, denote the first and second derivatives of δ_t on I_t . We now investigate the relationship between the volatility of the convenience yield $|v_t := |\gamma(\Xi_t)f'(I_t)| = -\gamma(\Xi_t)f'(I_t)$ and that of the convenience yield δ_t . As we know from the assumptions,

$$\frac{dv_t}{d\delta_t} = -\gamma(\Xi_t) \frac{f''(I_t)}{f'(I_t)} > 0. \quad (8)$$

Thus, holding the exogenous variables unchanged, the volatility of convenience yield increases monotonically with the convenience yield level and thus confirms the high convenience yield–high convenience yield volatility relationship. Also, Eq. (8) indicates that a stronger heteroskedasticity is likely to be caused by a more convex inventory–convenience yield curve (i.e. a larger $f''(I_t)$).

Therefore, the heteroskedasticity detected in major industrial commodities is structural; it is not caused by the heteroskedasticity in inventory, but is rather caused by the relationship

between the convenience yield and inventory. Since gold and silver are not used primarily for production and have a large amount of inventory, hence the convex relationship between the convenience yield and inventory will hardly show. This is consistent with the empirical finding that the convenience yield for gold and silver does not show significant convenience yields (see Schwartz, 1997 and CCD, 2005).

Since heteroskedasticity commonly exists in industrial commodities, it is important to incorporate this feature in futures dynamics and contingent claim pricing. In the following, we specify a “maximal” affine model for futures prices with the feature of heteroskedasticity of the convenience yield.

Dynamics of the Convenience Yield

In continuous time we assume the variance of the convenience yield V_t^δ following an affine structure on the convenience yield, $V_t^\delta = h(\delta_t + w)$ where h ($h > 0$) and w are constants. This affine specification makes our empirical model and calibration much more trackable. We define $\hat{\delta}_t := \delta_t + w$. w is an important value that relates to measuring the degree of heteroskedasticity for a certain commodity. The larger w is, the less heteroskedasticity the convenience yield process embodies; see also Eq. (11).

The risk neutral process

Assuming the risk-free interest rate follows an autonomous one-factor CIR process, we propose our model in the risk neutral measure as,

$$\begin{aligned}
 dr_t &= \kappa_r (\theta_r - r_t)dt + \sigma_r \sqrt{r_t} W_{1,t}^Q \\
 d\hat{\delta}_t &= [\kappa_\delta (\hat{\theta}_\delta - \hat{\delta}_t) + \kappa_{\delta r} (\theta_r - r_t)]dt + \sigma_\delta \sqrt{\hat{\delta}_t} dW_{2,t}^Q \\
 dx_t &= \left[r_t + w - \hat{\delta}_t - \frac{1}{2} (\sigma_{x\delta}^2 \hat{\delta}_t + \sigma_{xr}^2 r_t + \hat{v}_0 + v_{x\delta} \hat{\delta}_t + v_{xt} r_t) \right]
 \end{aligned} \tag{9}$$

$$+\sigma_{x\delta}\sqrt{\widehat{\delta}_t}dW_{1,t}^Q + \sigma_{xr}\sqrt{r_t}dW_{2,t}^Q + \sqrt{\widehat{v}_0 + v_{x\delta}\widehat{\delta}_t + v_{xr}r_t}dW_{3,t}^Q,$$

where $x_t := \ln(S_t)$ is the log-spot price, and r_t is the risk-free interest rate. W denotes the Wiener process with superscript Q representing the risk-neutral measure.

We also want the model to be maximal in a sense that the maximum number of identifiable parameters exists. That is, our model has the largest number of free parameters in a certain affine model class.⁷ Following the procedure of Dai and Singleton (2000), it is easy to show that our model belongs to a “maximal” affine model of $\mathbb{A}_1(3)$ in terms of Dai and Singleton (2000). For brevity, we omit the derivation, the complete details can be provided on request.

Defining $\widehat{\theta}_\delta := \theta_\delta + w$ and $v_0 := \widehat{v}_0 + v_{x\delta}w$, Eq. (9) can be rewritten in terms of convenience yield δ_t , interest rate r_t and log-spot price x_t ,

$$\begin{aligned} dr_t &= \kappa_r(\theta_r - r_t)dt + \sigma_r\sqrt{r_t}dW_{1,t}^Q \\ d\delta_t &= [\kappa_\delta(\theta_\delta - \delta_t) + \kappa_{\delta r}(\theta_r - r_t)]dt + \sigma_\delta\sqrt{\delta_t}dW_{2,t}^Q \\ dx_t &= \left[r_t - \delta_t - \frac{1}{2}(\sigma_{x\delta}^2(\delta_t + w) + \sigma_{xr}^2r_t + v_0 + v_{x\delta}\delta_t + v_{xt}r_t) \right] dt + \sigma_{x\delta}\sqrt{w + \delta_t}dW_{1,t}^Q \\ &\quad + \sigma_{xr}\sqrt{r_t}dW_{2,t}^Q + \sqrt{v_0 + v_{x\delta}\delta_t + v_{xr}r_t}dW_{3,t}^Q, \end{aligned} \tag{10}$$

Eq. (10) reflects a high convenience yield–high convenience yield volatility relationship discovered from the above empirical tests.

The heteroskedasticity measure and other steady-state statistics

The instantaneous variance of the convenience yield $V_t^\delta = \sigma_\delta^2(w + \delta_t)$ is composed of the two parts—a constant part σ_δ^2w and a stochastic part $\sigma_\delta^2\delta_t$. In the steady state (the state

⁷ Note that maximality is a theoretical concept. Dai and Singleton (2000) determine maximality by considering a series of ‘invariant rotations’ of the fundamental PDE (satisfied by path-independent European contingent claims) that leave all security prices unchanged.

when $t \rightarrow \infty$, $V_{t \rightarrow \infty}^\delta$ becomes $V_{t \rightarrow \infty}^\delta = \sigma_\delta^2(w + \theta_\delta)$, hence the fraction of the stochastic part in the convenience yield variance, z_δ , is written as,

$$z_\delta := \frac{\theta_\delta}{w + \theta_\delta} \quad (11)$$

$z_\delta (0 \leq z_\delta \leq \frac{\theta_\delta}{s_c + \theta_\delta})$ can be considered as a measure of the degree of heteroskedasticity for different commodities. Large z_δ corresponds to strong heteroskedasticity. w plays an important role in calculating z_δ . For example if w is very large compared with δ_t for a certain commodity, which means that the variance of convenience yield V_t^δ almost does not change with δ_t , the convenience yield for this commodity is almost homoskedastic, in this case z_δ is correspondingly very small.

The variance of spot prices V_t^s is

$$V_t^s := (\sigma_{x\delta}^2 + \nu_{x\delta})\delta_t + (\sigma_{xr}^2 + \nu_{xr})r_t + (u_0 + \sigma_{x\delta}^2 w). \quad (12)$$

Thus, V_t^s depends on the state of δ_T ; we thus define f_δ as the fraction that the total variance of *spot* prices explained by δ_t in the steady state,

$$f_\delta := \frac{(\sigma_{x\delta}^2 + \nu_{x\delta})\theta_\delta}{(\sigma_{x\delta}^2 + \nu_{x\delta})\theta_\delta + (\sigma_{xr}^2 + \nu_{xr})\theta_r + u_0 + \sigma_{x\delta}^2 w}. \quad (13)$$

Eq. (12) implies a high convenience yield–high spot volatility relationship. The correlation between the convenience yield and the log-spot price is also stochastic and dependent on the state of the convenience yield. We thus define the steady-state correlation $\rho_{x\delta}$ as

$$\rho_{x\delta} := \frac{\sigma_{x\delta} \sqrt{\theta_\delta + w}}{\sqrt{\sigma_{x\delta}^2(\theta_\delta + w) + \sigma_{xr}^2 \theta_r + u_0 + \nu_{x\delta} \theta_\delta + \nu_{xr} \theta_r}} \quad (14)$$

Specification of risk premia

To explain the historical time-series dynamics of prices, we need to specify the risk premia $\lambda(Y_t)$ in the relation between the risk-neutral (Q) and physical (P) measures,

$$dW_t^Q = dW_t^P + \lambda(Y_t)dt, \quad (15)$$

where W_t^P is a 3×1 vector of independent Brownian Motions and $\lambda(Y_t)$ is a 3×1 vector. To ensure that the process for Y_t under P also has the affine form, similar with Dai and Singleton (2000) we assume risk premia $\lambda(Y_t)$ to be a square root of the state variables,

$$\lambda(Y_t) = \sqrt{Y}J, \quad (16)$$

where $J = \left[\frac{\eta_r}{\sigma_r}, \frac{\eta_\delta}{\sigma_\delta}, \eta_x \right]^T$ is a 3×1 vector. This specification of risk premia allows the dependence of the futures risk premia on the convenience yield (see also Eq. (22)). Our formulation can be easily generalized to a non-affine diffusion for Y_t under P , but non-affine risk premia specification usually brings difficulties in empirical estimations.

Specifically, in the P measure, the dynamics of the three factors are:

$$\begin{aligned} dr_t &= (\kappa_r(\theta_r - r_t) + \eta_r r_t)dt + \sigma_r \sqrt{r_t} dW_{1,t}^P \\ d\hat{\delta}_t &= (\kappa_\delta(\hat{\theta}_\delta - \hat{\delta}_t) + \kappa_{\delta r}(\theta_r - r_t) + \eta_\delta \hat{\delta}_t)dt + \sigma_\delta \sqrt{\hat{\delta}_t} dW_{2,t}^P \\ dx_t &= \left[r_t + w - \hat{\delta}_t - \frac{1}{2}(\sigma_{x\delta}^2 \hat{\delta}_t + \sigma_{xr}^2 r_t + \hat{v}_0 + v_{x\delta} \hat{\delta}_t + v_{xr} r_t) + \frac{\eta_\delta \sigma_{x\delta}}{\sigma_\delta} \hat{\delta}_t + \frac{\eta_r \sigma_{xr}}{\sigma_r} r_t \right. \\ &\quad \left. + \eta_x(\hat{v}_0 + v_{x\delta} \hat{\delta}_t + v_{xr} r_t) \right] dt + \sigma_{xr} \sqrt{r_t} dW_{1,t}^P + \sigma_{x\delta} \sqrt{\hat{\delta}_t} dW_{2,t}^P \\ &\quad + \sqrt{\hat{v}_0 + v_{x\delta} \hat{\delta}_t + v_{xr} r_t} dW_{3,t}^P. \end{aligned} \quad (17)$$

In the futures pricing and model calibration, we assume $\kappa_{\delta r}$ to be zero, i.e. interest rates do not influence the movement of the convenience yields. This is mainly because there is no

analytical solution if $\kappa_{\delta r} \neq 0$.⁸ Also, since in this paper we are not mainly interested in the influence of interest rates on the convenience yield,⁹ we hence make this assumption to make our model parsimonious.

Pricing of commodity futures

It is well known (e.g., Cox et al., 1981) that the futures price $F(t, T)$ with maturity T at time t follows

$$F(t, T) = E_t^Q(S_T) = E_t^Q[e^{x_T}]. \quad (18)$$

Thus the futures price must satisfy the following Feynman-Kac equation with boundary condition $F(T, T) = S_T$.

$$\begin{aligned} \frac{\partial F}{\partial \tau} = & \frac{\partial F}{\partial x_t} \left[r_t + w - \widehat{\delta}_t - \frac{1}{2} (\sigma_{x\delta}^2 \widehat{\delta}_t + \sigma_{xr}^2 r_t + \widehat{v}_0 + v_{x\delta} \widehat{\delta}_t + v_{xr} r_t) \right] + \frac{\partial F}{\partial \widehat{\delta}_t} [\kappa_\delta (\widehat{\theta}_\delta - \widehat{\delta}_t)] \\ & + \frac{\partial F}{\partial r_t} [\kappa_r (\theta_r - r_t)] + \frac{1}{2} \frac{\partial^2 F}{\partial x_t^2} (\sigma_{x\delta}^2 \widehat{\delta}_t + \sigma_{xr}^2 r_t + \widehat{v}_0 + v_{x\delta} \widehat{\delta}_t + v_{xr} r_t) + \frac{1}{2} \frac{\partial^2 F}{\partial \widehat{\delta}_t^2} \sigma_\delta^2 \widehat{\delta}_t \\ & + \frac{1}{2} \frac{\partial^2 F}{\partial r_t^2} \sigma_r^2 r_t + \frac{\partial F}{\partial x_t} \frac{\partial F}{\partial r_t} \sigma_r \sigma_{xr} r_t + \frac{\partial F}{\partial x_t} \frac{\partial F}{\partial \widehat{\delta}_t} \sigma_\delta \sigma_{x\delta} \widehat{\delta}_t, \end{aligned} \quad (19)$$

where $\tau = T - t$. To solve Eq. (19) we first guess

$$\ln F(x_t, \widehat{\delta}_t, r_t, \tau) = x_t + A(\tau) \widehat{\delta}_t + C(\tau) r_t + B(\tau), \quad (20)$$

and then solve $A(\tau)$, $B(\tau)$, and $C(\tau)$. Appendix B shows the solution of $A(\tau)$, $B(\tau)$, and $C(\tau)$.

In the Q measure, the future process follows

$$\begin{aligned} \frac{dF(t, T)}{F(t, T)} = & (\sigma_{x\delta} + A(\tau) \sigma_\delta) \sqrt{\widehat{\delta}_t} dW_{1,t}^Q + (\sigma_{xr} + C(\tau) \sigma_r) \sqrt{r_t} dW_{2,r}^Q \quad (21) \\ & + \sqrt{\widehat{v}_0 + v_{x\delta} \widehat{\delta}_t + v_{xr} r_t} dW_{3,t}^Q. \end{aligned}$$

⁸ Numerical integration has to be employed.

⁹ Which is studied in a $A_0(3)$ framework by CCD (2005).

In the P measure, the futures process is

$$\begin{aligned} \frac{dF(t, T)}{F(t, T)} = & \left(\frac{\eta_1 \sigma_{x\delta}}{\sigma_\delta} \hat{\delta}_t + \frac{\eta_2 \sigma_{xr}}{\sigma_r} r_t + \eta_3 (\hat{v}_0 + v_{x\sigma} \hat{\delta}_t + v_{xr} r_t) + A(\tau) \eta_1 \hat{\delta}_t + C(\tau) \eta_2 r_t \right) dt \\ & + (\sigma_{x\delta} + A(\tau) \sigma_\delta) \sqrt{\hat{\delta}_t} dW_{1,t}^P + (\sigma_{xr} + C(\tau) \sigma_r) \sqrt{r_t} dW_{2,t}^P. \quad (22) \end{aligned}$$

From Eq. (22) we see that the futures risk premium depends on the convenience yield δ_t (note that $\hat{\delta}_t := \delta_t + w$), which is consistent with Gorton et al. (2007).

In general, our model captures three important features of commodity futures—the heteroskedasticity of the convenience yield, the positive relationship between spot volatility and the convenience yield, and the dependence of the futures risk premium on the convenience yield.

Since we want to calibrate our model by jointly using futures prices and bond prices, we need to specify the bond pricing formula given the model setup. Bond prices driven by a CIR process has been studied for long time (e.g., Brown and Dybvig, 1986).

Empirical Implementation

The data

Our dataset consists of futures contracts on crude oil, copper, and zero-coupon bond prices. For all commodities, we use daily data from January 2000 to February 2006 (1544 observations for each commodity). The futures prices of *WTI* crude oil (*CL*) and high-grade copper (*HG*) are from the New York Mercantile Exchange.¹⁰ Table 2 contains the summary statistics for commodity prices and returns. The time to maturity ranges from 1 month to 17 months for both oil and copper contracts. We denote F_n as the n^{th} contract closest to maturity; e.g., F_1 is the future contract that is closest to maturity. Since the maturities are in consecutive

¹⁰ Note that the oil and copper data are from the NYMEX and COMEX divisions, respectively.

calendar months, n also roughly denotes the time to maturity (in monthly units). In this paper, we use seven time series for oil and copper— $F_1, F_3, F_6, F_9, F_{12}, F_{15}$, and F_{17} contracts. The interest rate data from January 2000 to February 2006 are from the Federal Reserve Board. We use daily Constant Maturity Treasury yields (CMT) in our calibration.

Estimation

One of the difficulties in the calibration of our model is that the three factors are not directly observable. Several calibration methodologies have been proposed to solve this problem, such as efficient method of moments (Gallant and Tauchen, 1996), maximum likelihood estimation (i.e., the Chen and Scott, 1993 method), and the Kalman filter method. Duffee and Stanton (2004) compare these methods and conclude that the Kalman filter is the best method among those three, especially when the model is complicated. We specify our model in a state-space form and use the Kalman filter to calibrate the model.

The state-space form normally consists of a transition equation and a measurement equation. The transition equation shows the stochastic process of the data-generating process. Thus, the transition equation in the model should be the discrete version of Eq. (17). The measurement equation relates the time series of multivariate observable variables (futures and bonds prices for different maturities in our case) to an unobservable vector of state variables (factors $\hat{\delta}, r, x$). The measurement equation is obtained using Eq. (20) with uncorrected noises taking account of the pricing errors. These errors may be caused by bid–ask spreads, non-simultaneity of the observations, etc.

To describe the transition and measurement equations in greater detail, suppose the data are sampled in an equal interval: $t_n, n = 1, \dots, T$. Let $\Delta = t_{n+1} - t_n$ be the interval between two observations and $X = (X_1, \dots, X_T)$ and $Z = (Z_1, \dots, Z_T)$ be the total latent state variables and

observations (from 1 to T). $X_n = [\hat{\delta}_n, r_n, x_n]$ represents the vector for state variables at the time t_n , and $Z_n = [F_1, \dots, F_{17}, R_6, R_{60}]^T$ represents the nine observations of each time n with the first seven as the futures prices ($F_1, F_3, F_6, F_9, F_{12}, F_{15}$, and F_{17}) and the last two as bond yields (6 and 60 months). Note that $X_{n+1}|X_n$ in the transition equation is not normally distributed in continuous time. However, since we use daily data in our calibration, $X_{n+1}|X_n$ should be very close to the one obtained simply by Euler discretization of Eq. (17) which does show a normal distribution for $X_{n+1}|X_n$, i.e.

$$X_{n+1}|X_n \sim \mathbb{N}(E(X_{n+1}|X_n), \text{var}(X_{n+1}|X_n)), \quad (23)$$

where \mathbb{N} denotes normal distribution. This is also suggested by many previous studies of calibrating and simulating stochastic volatility models (see Johannes and Polson, 2002; Kahl and Jackel, 2006, and Andersen, 2007). Since our model is very close to the Heston (1993) stochastic volatility model, we adopt those methodologies. Therefore, the transition equation is specified as,

$$X_{n+1} = HX_n + L + \omega_n, \quad (24)$$

where

$$H = \begin{bmatrix} 1 + (\eta_1 - \kappa_\delta)\Delta & 0 & 0 \\ 0 & 1 + (\eta_2 - k_r)\Delta & 0 \\ \left(-1 - \frac{1}{2}(\sigma_\delta^2 + v_{x\delta})\right) \Delta & \left(-1 - \frac{1}{2}(\sigma_r^2 + v_{xr})\right) \Delta & 1 \\ + \frac{\eta_1 \sigma_{x\delta}}{\sigma_\delta} + \eta_3 v_{x\delta} & + \frac{\eta_2 \sigma_{xr}}{\sigma_r} + \eta_3 v_{xr} & \end{bmatrix} \quad (25)$$

$$L = \left[\kappa_\delta \hat{\theta}_\delta \Delta \quad \kappa_r \theta_r \Delta \quad \left(w - \frac{1}{2} \hat{v}_0 + \eta_3 \hat{v}_0 \right) \Delta \right]^T. \quad (26)$$

From the futures pricing formula, we can get the measurement equation.

$$Z_n = V_n X_n + U_n + \varepsilon_n, \quad (27)$$

where

$$V_n = \begin{bmatrix} A(\tau_{F_1}) & C(\tau_{F_1}) & 1 \\ \vdots & \vdots & \vdots \\ A(\tau_{F_{17}}) & C(\tau_{F_{17}}) & 1 \\ 0 & D(\tau_{R_6}) & 0 \\ 0 & D(\tau_{R_{60}}) & 0 \end{bmatrix} \quad (28)$$

$$U_n = [B(\tau_{F_1}) \cdots B(\tau_{F_{17}}) G(\tau_{R_6}) G(\tau_{R_{60}})]^T, \quad (29)$$

where τ_i denotes time to maturity of contract i , ω_n and ε_n are zero-mean random Gaussian noise vectors at time t_n respectively with variance–covariance matrices Φ_n and Σ . Note that, for simplicity, we assume that the variances in the futures pricing errors of all maturities are the same, and so are the bond yields. This assumption is also based on the belief that our model should price futures across different maturities equally well. We also assume that the futures and bonds pricing errors are independent across different maturities. The variances of pricing errors for futures and bonds are denoted by ξ_F^2 and ξ_R^2 respectively.

$$\Phi_n = \Delta \begin{bmatrix} \sigma_\delta^2 \hat{\delta}_n & 0 & \sigma_\delta \sigma_{x\delta} \hat{\delta}_n \\ 0 & \sigma_r^2 r_n & \sigma_r \sigma_{xr} r_n \\ \sigma_\delta \sigma_{x\delta} \hat{\delta}_n & \sigma_r \sigma_{xr} r_n & \left(\begin{array}{c} \sigma_{x\delta}^2 \hat{\delta}_n + \sigma_{xr}^2 r_n \\ + \hat{v}_0 + v_{x\sigma} \hat{\delta}_n + v_{xr} r_n \end{array} \right) \end{bmatrix} \quad (30)$$

$$\Sigma = \begin{bmatrix} \xi_F^2 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & \xi_R^2 \end{bmatrix}. \quad (31)$$

In Eq. (30) Φ_n should depend on the ‘true’ state variable $\hat{\delta}_n$; however, since the true $\hat{\delta}_n$ is not known, we thus use its estimates in the Kalman filter iteration. Thus, our Kalman filter algorithm is regarded as an ‘extended’ Kalman filter. Although the extended Kalman filter in our case may be biased, as suggested by Duan and Simonato (1999) and De Jong (2000), this Kalman filter recursion still performs reasonably well. We thus use the extended Kalman filter recursion to calibrate our model, as also suggested by Duffee and Stanton (2004) and Chen and Scott (2003). We maximize the likelihood function from the Kalman filter to obtain the parameter

estimates. Note that the Kalman filter algorithm above cannot guarantee the non-negativity of the convenience yield $\hat{\delta}_n$ and interest rate r time series. In case of negative convenience yields or interest rates, we replace them with zeros.¹¹ This is suggested by Chen and Scott (2003) when calibrating interest rates using multi-CIR-type factors in the Kalman filter framework.

Results

Table 3 presents the estimates of our model for oil and copper. For each commodity, we present the risk-neutral parameters, the risk premia parameters, volatility, and the correlation parameters. Table 3 shows that all risk-neutral parameters are significant. This implies that our model setup is indeed necessary to explain the dynamics of the two commodities. For both commodities, we can see that z_δ is significant, i.e., the convenience yield δ_t plays an important role in determining the variance of the convenience yield V_t^δ . This is consistent with the heteroskedasticity discovered in our empirical studies in Section 2. Moreover, comparing z_δ for both oil and copper, we find that the heteroskedasticity for copper is stronger than that for oil. As shown empirically by Gorton et al. (2007), copper has a much more convex inventory–convenience yield relationship than crude oil and hence possesses a stronger heteroskedasticity. $\sigma_{x\delta}$ and $\nu_{x\delta}$ are positive and significant for both oil and copper, which is consistent with the theory of storage and indicates the high convenience yield–high spot price volatility relationship. Also, f_δ is significant, thus the convenience yield factor plays an important role in determining the spot price volatility. The steady-state correlation between the spot price and the convenience yield is large for both oil and copper, which is consistent with the theory of storage and the findings in Schwartz (1997) and CCD (2005).

¹¹ Note that it is very rare that the convenience yields or interest rates are below zero in our Kalman filter iteration for oil and copper futures.

To do formal tests on the heteroskedasticity and high convenience yield–high spot price volatility relationship, we run the likelihood ratio tests between our model and: 1) a model without convenience yield heteroskedasticity, 2) a model with spot volatility independent of the convenience yield. We run the first test by estimating a model with an arbitrarily large w value, i.e., $w = 100$.¹² We run the second test by estimating a model with $\sigma_{x\delta} = v_{x\delta} = 0$. Note that we run the two tests on both oil and copper data. Table 4 lists the results. Table 4 suggests that both the heteroskedasticity and the high convenience yield–high spot volatility relationship are two significant features for both oil and copper futures prices.

For oil, the significantly positive η_δ confirms a high convenience yield–high futures risk premium relationship; although η_δ for copper is not significant the positive sign is consistent with this relationship. The convenience yield process of oil is estimated to be persistent under both the risk-neutral measure and the physical measure with a long-run mean of about 5% and a half-life of about 0.6 years. The insignificant σ_{xr} shows that the interest rate is not significantly correlated with the spot price, which is also consistent with the findings of Schwartz (1997) and CCD (2005). z_δ is about 9%, which shows that in general about 9% of the convenience yield variance can be explained by the heteroskedasticity. We also see that the risk premia of x and the $\hat{\delta}$ factor are quite significant, while that for r is not.

The long-run mean and the half-life of convenience yields for copper are about 10% and 1.9 years, respectively. For copper, the contribution of the stochastic part of the convenience yield variance is quite high, i.e., z_δ is about 62%. This shows that the convenience yield is substantial in explaining the convenience yield variance. For the volatility of the copper spot

¹² Given that the convenience yield for copper and oil fluctuate between -0.3 and 0.3 , $w=100$ is large enough to make the heteroskedastic effect vanish.

price, v_{xr} and σ_{xr} are zeros through the calibration, which shows that the interest rate actually does not influence the volatility of the copper spot prices.

Option Pricing Applications

We price options on commodity futures following the framework of Duffie et al. (2000) and Chacko and Das (2002). In our model, an increase in the convenience yield implies an increase in the futures volatility (indicated by the high convenience yield–high spot volatility relationship), also an increase in the convenience yield tends to associate with an increase in the futures prices because of their significant positive correlation (indicated by the positive correlation between spot return and the convenience yield, see also Eq. (21)). Thus, there exists a positive correlation between the futures volatility and returns, which results in a positive return skewness and in turn causes an upward sloping implied volatility smile pattern. This positive skewness of commodity futures returns is also consistent with Deaton and Laroque (1992) where commodity prices occasionally show large positive spikes. Fig. 2 shows the patterns of implied volatility for a three-month option on a three-month futures calculated from our model. Note that the moneyness is defined as the ratio of futures prices to the option strike prices and parameters in the option pricing are those calibrated from the oil futures in this paper. From the figure, we do see the upward sloping smile pattern and the implied volatility is high given a high initial convenience yield level.

In the oil option market, the option implied volatility often displays the upward sloping pattern. Note that, using the monthly implied volatility smile dataset (from Jan. 2000 to Feb. 2006) of oil futures, we find that about 70% of the data show an upward sloping pattern. Less than 10% of the data show a downward sloping pattern (see Fig. 3). Some researchers find a

positive correlation between the underlying returns and volatility, such as Richter and Sorensen (2004), but do not offer a reliable explanation for this. Existing models such as CCD (2005) and Schwartz (1997) cannot explain this phenomenon either. From our model, we see that the positive correlation between the convenience yields and spot prices and the high convenience yield-high spot price volatility relationship together serve as a good interpretation of upward sloping implied volatility smiles.

Conclusion

This paper primarily investigates the heteroskedasticity of the convenience yield. We empirically discover that industrial commodities such as oil and copper all show heteroskedasticity, i.e., the high convenience yield–high convenience yield volatility relationship. The heteroskedasticity of the convenience yield is caused by the non-linear (convex) relationship between the convenience yield and inventory, which is consistent with the theory of storage. A stronger heteroskedasticity is likely to be caused by the more convex inventory–convenience yield equilibrium curve.

We propose an affine three-factor (log-spot prices, interest rates and the convenience yield) model, where the convenience yield variance is specified as an affine structure of the instantaneous convenience yield, and the interest rate is modeled by a CIR process. Our model copes with three features of commodity futures: the heteroskedasticity of the convenience yields, the high convenience yield–high spot volatility relationship and the high convenience yield–high futures risk premium relationship. The likelihood ratio tests show that our model is better than models that fail to specify either the convenience yield heteroskedasticity or the influence of the convenience yield on spot-price volatility. We also devise a measure of heteroskedasticity—the Z_δ score. Our tests show that the Z_δ score of copper is seven times more than that of crude oil,

which is consistent with the empirical result by Gorton et al. (2007) that copper has a more convex inventory–convenience yield curve than oil. As an application we use our model to price European vanilla options. The implied volatility in our model shows an upward sloping pattern that matches market observations.

Appendix A

Robust checks on convenience yield heteroskedacity

In this appendix, we perform several robustness tests to double check the existence of the heteroskedasticity of the convenience yield and inventories.

First, we run regressions (2) to (5) using three groups of commodities: 1) oil products, i.e., heating oil and unleaded gasoline; 2) industrial metals, zinc and nickel; 3) precious metal, silver and gold. For heating oil and gasoline the data consist of daily futures prices traded on NYMEX for the period from January 2000 to February 2006 (1544 observations). The time to maturity ranges from 1 month to 9 months. For nickel and zinc, the data consist of daily futures prices traded on the London Mercantile Exchange (LME) for the period of January 2000 to December 2005 (1509 observations) with 4 contracts: the spot contract and 3, 15, and 27 months futures. For gold and silver, the data consist of daily futures prices traded on COMEX for the period of January 2000 to February 2006 (1544 observations). Due to data availability constraints, we have only two contracts available for gold and silver, i.e., the one- and three-month futures prices. Table 5 shows the regression results of Eqs. (2) to (5).

The table shows that the oil products and industrial metals also preserve a high convenience yield–high convenience yield volatility relationship from the significant and positive t test statistics. However, for precious metals, this relationship no longer holds. Note that although silver and gold show significant negative d and f values, their b values are not significant.

We then do a robustness check of whether the inventories of zinc, nickel, and gold show heteroskedacity by running regressions (2) to (5). The weekly gold stock in warehouses is

obtained from COMEX through Datastream from February 1997 to February 2006. The weekly zinc and nickel stocks in warehouses are obtained from LME through Datastream from January 1995 to February 2006. We cannot find inventory data for oil products and silver, and therefore we did not test the inventory behavior of them. Table 6 shows the results.

The regression results show that none of zinc nickel and gold shows heteroskedacity.

Note that although the value d is significant for zinc b and f are not.

Appendix B

Pricing commodity futures

Plug $F(x_t, \hat{\delta}_t, r_t, \tau) = \exp(x_t + \hat{\delta}_t A(\tau) + r_t C(\tau) + B(\tau))$ into Eq. (19) yields

$$\begin{aligned} r_t + w - \hat{\delta}_t + A\kappa_\delta(\hat{\theta}_t - \hat{\delta}_t) + C\kappa_r(\theta_r - r) + \frac{1}{2}A^2\sigma_\delta^2\hat{\delta}_t + \frac{1}{2}C^2\sigma_r^2r_t + A\sigma_{x\delta}\sigma_\delta\hat{\delta}_t \\ + C\sigma_{xr}\sigma_r r_t = A'\hat{\delta}_t + C'r_t + B'. \end{aligned} \quad (32)$$

To make the coefficient of $\hat{\delta}_t$ and r_t zero, we thus have the following Ordinary Differential Equations (ODE).

$$\frac{1}{2}\sigma_\delta^2A^2 + (\sigma_{x\delta}\sigma_\delta - \kappa_\delta)A - 1 = A' \quad (33)$$

$$1 - (\kappa_r - \sigma_r\sigma_{xr})C + \frac{1}{2}\sigma_r^2C^2 = C' \quad (34)$$

$$w + A\kappa_\delta\hat{\theta}_\delta + C\kappa_r\theta_r = B' \quad (35)$$

with boundary condition $A(0) = 0, B(0) = 0, C(0) = 0$. A' denotes the first derivative of $A(\tau)$ on τ , so are C' and B' . We have

$$A(\tau) = A_0 + \frac{1}{ge^{-\lambda\tau} - \frac{a}{\lambda}} \quad (36)$$

where $a = \frac{1}{2}\sigma_\delta^2, b = \sigma_{x\delta}\sigma_\delta - \kappa_\delta, \lambda = b + 2a, A_0 = \sqrt{b^2 + 4a}, g = \frac{a}{\lambda} - \frac{1}{A_0}$. Eq. (34) can be

solved using exactly the same procedure used in solving Eq. (32). However, since the sign of

$q = (\kappa_r - \sigma_r\sigma_{xr})^2 - 2\sigma_r^2$ is unknown, we need to discuss all scenarios for $q > 0, q < 0$, and

$q = 0$.

Case I. $q > 0$, define $\eta = \sqrt{q}$, the solution is

$$C(\tau) = C_0 + \frac{1}{ne^{-\eta\tau} - \frac{\sigma_r^2}{2\eta}} \quad (37)$$

where $C_0 = \frac{\kappa_r - \sigma_r \sigma_{xr} + \eta}{\sigma_r^2}$, and $n = \frac{\sigma_r^2}{2\eta} - \frac{1}{C_0}$.

Case II. $q = 0$, the solution is,

$$C(\tau) = \frac{\kappa_r - \sigma_r \sigma_{xr}}{\frac{1}{2} \sigma_r^2 (\kappa_r - \sigma_r \sigma_{xr}) \tau + \sigma_r^2} + \frac{\kappa_r - \sigma_r \sigma_{xr}}{\sigma_r^2} \quad (38)$$

Case III. $q < 0$, define $\gamma = \sqrt{-q}$, the solution is,

$$C(\tau) = \frac{\gamma}{\sigma_r^2} \tan \left[\frac{\gamma}{2} \tau + \tan^{-1} \left(-\frac{\kappa_r - \sigma_r \sigma_{xr}}{\gamma} \right) \right] + \frac{\kappa_r - \sigma_r \sigma_{xr}}{\sigma_r^2} \quad (39)$$

With the solution of (33) and (34), $B(\tau)$ can be expressed as:

Solution for Case I:

$$\begin{aligned} B(\tau) = & \left(w + \kappa_\delta \hat{\theta}_\delta A_0 + \kappa_r \theta_r \frac{\kappa_r - \sigma_r \sigma_{xr}}{\sigma_r^2} - \frac{\kappa_\delta \hat{\theta}_\delta \lambda}{a} - \frac{2\kappa_r \theta_r \eta}{\sigma_r^2} \right) \tau + \frac{\kappa_\delta \hat{\theta}_\delta}{a} \left[\ln \left(\frac{g\lambda - a}{g\lambda e^{-\lambda\tau} - a} \right) \right] \\ & + \frac{2\kappa_r \hat{\theta}_\delta}{\sigma_r^2} \left[\ln \left(\frac{2n\eta - \sigma_r^2}{2n\eta e^{-\eta\tau} - \sigma_r^2} \right) \right]. \end{aligned} \quad (40)$$

Solution for Case II:

$$\begin{aligned} B(\tau) = & \left(w + \kappa_\delta \hat{\theta}_\delta A_0 + \kappa_r \theta_r \frac{\kappa_r - \sigma_r \sigma_{xr}}{\sigma_r^2} - \frac{\kappa_\delta \hat{\theta}_\delta \lambda}{a} \right) \tau + \frac{\kappa_\delta \hat{\theta}_\delta}{a} \left[\ln \left(\frac{g\lambda - a}{g\lambda e^{-\lambda\tau} - a} \right) \right] \\ & + \frac{2\kappa_r \hat{\theta}_\delta}{\sigma_r^2} \left[\ln \frac{1}{2} (\kappa_r - \sigma_r \sigma_{xr}) \tau + 1 \right]. \end{aligned} \quad (41)$$

Solution for Case III:

$$\begin{aligned}
 B(\tau) = & \left(w + \kappa_\delta \hat{\theta}_\delta A_0 + \kappa_r \theta_r \frac{\kappa_r - \sigma_r \sigma_{xr}}{\sigma_r^2} - \frac{\kappa_\delta \hat{\theta}_\delta \lambda}{a} \right) \tau + \frac{\kappa_\delta \hat{\theta}_\delta}{a} \left[\ln \left(\frac{g\lambda - a}{g\lambda e^{-\lambda\tau} - a} \right) \right] \\
 & + \frac{2\kappa_r \hat{\theta}_\delta}{\sigma_r^2} \left[\ln \left(\left| \frac{\cos \left(\tan^{-1} \left(-\frac{\kappa_r - \sigma_r \sigma_{xr}}{\gamma} \right) \right)}{\cos \left(\frac{\gamma}{2} \tau + \tan^{-1} \left(-\frac{\kappa_r - \sigma_r \sigma_{xr}}{\gamma} \right) \right)} \right| \right) \right]. \quad (42)
 \end{aligned}$$

References

- Andersen, L. *Efficient simulation of the Heston stochastic volatility model*, working paper, Bank of America Securities, 2007.
- Bessembinder, H., Coughenour, J., Seguin, P., Smoller, M., 1995. Mean reversion in equilibrium asset prices: evidence from the futures term structure. *J. Finance* 50, 361–375.
- Brennan, M., 1958. The supply of storage. *Am. Econ. Rev.* 48, 50–72.
- Brown, S., Dybvig, P., 1986. The empirical implications of the Cox, Ingersoll, Ross theory of the term structure of interest rates. *J. Finance* 41, 617–630.
- Casassus, J., Collin-Dufresne, P., 2005. Stochastic convenience yield implied from commodity futures and interest rates. *J. Finance* 60, 2283–2331.
- Chacko, G., Das, S., 2002. Pricing interest rate derivatives: a general approach. *Rev. Financ. Stud.* 15, 195–241.
- Chen, R., Scott, L., 1993. Maximum likelihood estimation for a multifactor equilibrium model of the term structure of interest rates. *J. Fixed Income* 3, 14–31.
- Chen, R., Scott, L., 2003. Multi-factor Cox–Ingersoll–Ross models of the term structure: estimates and tests from a Kalman filter model. *J Real Estate Finance Econ.* 27, 143–172.
- Cox, J., Ingersoll, J., Ross, S., 1981. The relation between forward prices and futures prices. *J. Financ. Econ.* 9, 321–346.
- Dai, Q., Singleton, K., 2000. Specification analysis of affine term structure models. *J. Finance* 55, 1943–1978.
- Deaton, A., Laroque, G., 1992. On the behaviour of commodity prices. *Rev. Econ. Stud.* 59, 1–23.

- Deaton, A., Laroque, G., 1996. Competitive storage and commodity price dynamics. *J Polit. Econ.* 104, 896–923.
- De Jong, F., 2000. Time series and cross-section information in affine term structure models. *J Bus. Econ. Stat.* 18, 300–314.
- Duan, J., Simonato, J., 1999. Estimating exponential-affine term structure models by the Kalman filter. *Rev. Quant. Finance Acc.* 13, 111–135.
- Duffee, G., Stanton, R. *Estimation of dynamic term structure models*, working paper, Haas School of Business, Univ. of California, Berkeley, 2004.
- Duffie, D., Pan, J., Singleton, K., 2000. Transform analysis and asset pricing for affine jump-diffusion. *Econometrica* 68, 1343–1376.
- Fama, E., French, K., 1987. Commodity futures prices: some evidence on forecast power. *J. Bus.* 60, 55–74.
- Fama, E., French, K., 1988. Business cycles and the behavior of metals prices. *J. Finance* 43, 1075–1093.
- Gallant, R., Tauchen, G., 1996. Which moments to match? *Econometric Theory* 12, 657–681.
- Gibson, R., Schwartz, E., 1990. Stochastic convenience yield and the pricing of oil contingent claims. *J. Finance* 45, 959–976.
- Dincerler, C., Khokher, Z., Simin, T. *An empirical analysis of commodity convenience yields*, working paper, available on SSRN, 2005.
- Goldfeld, S., Quandt, R., 1972. *Nonlinear Methods in Econometrics* North-Holland. Amsterdam Press.
- Gorton, G., Hayashi, F., Rouwenhorst, G. *The fundamentals of commodity futures returns*, Yale ICF Working Paper, No. 07-08, 2007.

- Heston, S., 1993. A closed form solution for options with stochastic volatility. *Rev. Financ. Stud.* 6, 327–343.
- Kahl, C., Jackel, P. *Fast strong approximation Monte-Carlo schemes for stochastic volatility models*, working paper, ABN AMRO, London, 2006.
- Kaldor, N., 1939. Speculation and economic stability. *Rev. Econ. Stud.* 7, 1–27.
- Johannes, M., Polson, N. *MCMC methods for financial econometrics*, working paper, University of Chicago, 2002.
- Newey, W., West, K., 1987. A simple positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55, 703–708.
- Nielsen, M., Schwartz, E., 2004. Theory of storage and the pricing of commodity claims. *Rev. Deriv. Res.* 61, 341–378.
- Ng, V., Pirrong, S., 1994. Fundamentals and volatility: storage, spreads and the dynamics of metal prices. *J. Bus.* 67, 203–230.
- Pindyck, R., 2001. The dynamics of commodity spot and futures markets: a primer. *Energy J.* 22, 1–29.
- Richter, M., Sorensen, C. *Stochastic volatility and seasonality in commodity futures and options: the case of soybeans*, working paper, Copenhagen Business School, 2004.
- Routledge, B., Seppi, D., Spatt, C., 2000. Equilibrium forward curves for commodities. *J. Finance* 55, 1297–1338.
- Schwartz, E., 1997. The stochastic behaviour of commodity prices: implications for valuation and hedging. *J. Finance* 52, 923–973.
- Telser, L., 1958. Futures trading and the storage of cotton and wheat. *J. Polit. Econ.* 66, 233–255.

Wooldridge, J., 2003. *Introductory Econometrics: A Modern Approach*. Thomson, South Western press, USA.

Working, H., 1949. The theory of the price of storage. *American Economic Review* 39, 1254–1262.

Table 1. Heteroskedacity tests of the convenience yield and inventory. Panel A, Breusch–Pagan, Glejser and GARCH11 heteroskedasticity tests for convenience yield (regressions (2) to (5)); Panel B, Breusch–Pagan, Glejser and GARCH11 heteroskedasticity tests for inventory. Note that T_1 and T_2 are in months; numbers in brackets are t -statistics. Both b , d and f are significantly positive in Panel A, but not in Panel B.

	b (B–P) Test	d (Glejser test)	f (GARCH11)
<i>Panel A</i>			
Oil			
Convenience yield ($T_1 = 1, T_2 = 17$)	0.00034(4.44)	0.0141 (5.83)	0.00019(7.74)
Convenience yield ($T_1 = 1, T_2 = 9$)	0.00063(4.30)	0.0167(5.98)	0.00035(7.30)
Convenience yield ($T_1 = 1, T_2 = 3$)	0.0021(1.92)	0.0217(4.10)	0.0012(5.62)
Copper			
Convenience yield ($T_1 = 1, T_2 = 17$)	0.00031(5.39)	0.0203(11.3)	0.00029(18.93)
Convenience yield ($T_1 = 1, T_2 = 9$)	0.00058(4.59)	0.026(11.1)	0.00054(18.79)
Convenience yield ($T_1 = 1, T_2 = 3$)	0.0079(3.28)	0.078(7.81)	0.0058(16.83)
<i>Panel B</i>			
Oil			
Inventory	-1.59(-0.32)	-0.0001(-0.02)	0.000052(0.49)
Copper			
Inventory	-0.0007(-0.47)	0.0022(1.01)	-0.000012(-0.10)

Table 2. Annualized mean and standard deviation of returns from January 2000 to February 2006. The data is obtained from New York Mercantile Exchange with a daily frequency. The mean and standard deviation of futures returns are annualized.

	F_1	F_3	F_6	F_9	F_{12}	F_{15}	F_{17}
<i>Oil</i>							
Mean (return)	0.1429	0.1604	0.1792	0.1905	0.1989	0.205	0.2079
Std (return)	0.3654	0.3152	0.2777	0.2559	0.2421	0.2334	0.2291
<i>Copper</i>							
Mean (return)	0.1538	0.1512	0.1462	0.141	0.1356	0.1301	0.1266
Std Return	0.2271	0.2222	0.212	0.204	0.2005	0.1992	0.1999

Table 3. Model parameters for oil and copper contracts. Numbers in brackets are standard deviations

Variables	Oil		Copper	
κ_δ	1.2556	(0.0100)	0.3676	(0.0113)
κ_r	0.3028	(0.0058)	0.2757	(0.0051)
σ_δ	0.4055	(0.0134)	0.2066	(0.0097)
σ_r	0.0473	(0.0027)	0.0475	(0.0028)
$\hat{\theta}_\delta$	0.5399	(0.0526)	0.1565	(0.0049)
θ_r	0.0570	(0.0004)	0.0587	(0.0004)
$\sigma_{x\delta}$	0.4282	(0.0142)	0.3621	(0.0374)
σ_{xr}	0.0333	(0.0354)	0.0002	(0.0000)
\hat{v}_0	0.0023	(0.0003)	0.0247	(0.0017)
$v_{x\delta}$	0.0073	(0.0031)	0.1013	(0.0159)
v_{xr}	1.1209	(0.1092)	0.0000	(0.0000)
η_δ	0.0085	(0.0000)	0.0549	(0.253)
η_r	0.0052	(2.5016)	0.0000	(0.1135)
η_x	9.4385	(2.1991)	5.5628	(2.1481)
w	0.4921	(0.0537)	0.0588	(0.0024)
ξ_F	0.0069	(0.0009)	0.0048	(0.0006)
ξ_R	0.0031	(0.0005)	0.0030	(0.0005)
Log likelihood	48,135		52,307	
z_δ	8.89%		62.37%	
$p_{x\delta}$	0.7768		0.6102	
f_δ	5.44%		37.19%	

Table 4. Likelihood ratio test. This table presents the likelihood ratio statistics. 1) A model without the convenience yield heteroskedacity, 2) a model with spot volatility independent of the convenience yield. the 1% significant levels are 6.635 for model 1) and 9.210 for model 2).

	Variables	Oil	Copper
Model 1)	$w = 100$	1340	5808
Model 2)	$\sigma_{x\delta} = v_{x\delta} = 0$	1330	276

Table 5. Convenience yield heteroskedasticity test on three groups of commodities. This table contains results from regressions (2) to (5) of convenience yield for 6 commodities. Both b and d are significantly positive for oil products and industrial metals, but not for precious metals. Numbers in brackets denote the t-statistics.

	b (B-P test)	d (Glejser test)	f (GARCH11)
<i>Oil products</i>			
Heating oil			
$T_1 = 1, T_2 = 9$	0.002(7.14)	0.024(11.49)	0.0018(6.01)
$T_2 = 1, T_2 = 6$	0.002(9.10)	0.026(15.72)	0.0045(6.68)
$T_2 = 1, T_2 = 3$	0.0015(12.22)	0.017(20.40)	0.025(5.11)
Gasoline			
$T_1 = 1, T_2 = 9$	0.00035(2.04)	0.008(3.66)	0.00048(3.70)
$T_2 = 1, T_2 = 6$	0.00038(2.23)	0.008(4.72)	0.0014(2.81)
$T_2 = 1, T_2 = 3$	0.00016(2.71)	0.005(6.68)	0.0065(3.30)
<i>Industrial metals</i>			
Zinc			
$T_1 = 0, T_2 = 27$	0.0024(5.10)	0.026(13.02)	0.00027(12.60)
$T_1 = 0, T_2 = 15$	0.0043(8.60)	0.031(13.85)	0.00035(14.30)
$T_1 = 0, T_2 = 3$	0.018(12.76)	0.13(15.93)	0.0071(7.87)
Nickel			
$T_1 = 0, T_2 = 27$	0.0061(10.41)	0.036(16.32)	0.00057(18.25)
$T_1 = 0, T_2 = 15$	0.0087(10.2)	0.040(14.69)	0.00076(12.83)
$T_1 = 0, T_2 = 3$	0.008(10.36)	0.092(14.50)	0.0044(6.64)
<i>Precious metals (COMEX)</i>			
Silver			
$T_1 = 1, T_2 = 3$	-0.000021(-0.63)	-0.0074(-3.96)	-0.00038(-4.84)
Gold			
$T_1 = 1, T_2 = 3$	-0.000051(-0.73)	-0.0128(-4.63)	-0.000011(-2.69)

Table 6. Inventory heteroskedasticity test on zinc, nickel and gold. This table contains results from regressions (2) to (5) of inventory for 3 commodities. Both b , d and f are insignificant for all commodities except the d value for zinc. Numbers in brackets denote the t -statistics.

	b (B-P test)	d (Glejser test)	f (GARCH11)
Zinc	0.00006(0.86)	0.004167(3.14)	0.000063(1.90)
Nickel	-0.000001(-0.13)	0.00123(0.904)	0.0000011(0.4381)
Gold	-0.00029(-0.209)	0.00299(1.50)	-0.000028(-0.0165)

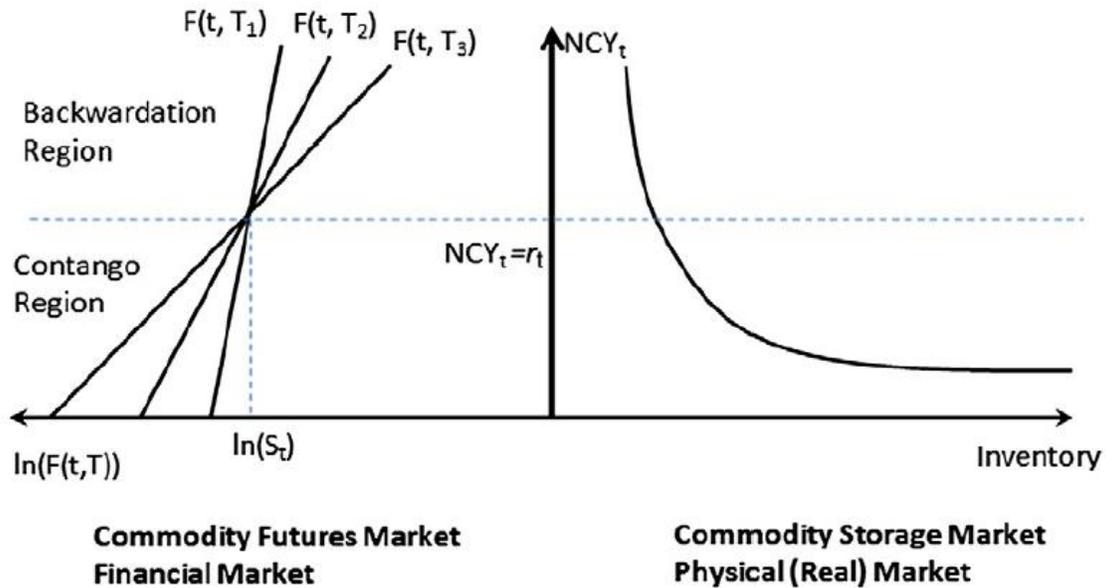


Figure 1. The convenience yield links the commodity storage market with the commodity futures market. The convenience yield can be regarded as implicit value of holding one marginal unit of commodity. Due to the convex structure of convenience yield–inventory equilibrium curve, the same amount of quantity changes in inventory level will result in bigger changes in convenience yield when inventory level is lower than when the level is high. The left panel of the graph shows a negative relationship between the convenience yield and commodity futures prices where $\ln(F(t, T))$ and $\ln(S_t)$ denote log futures prices ($T_3 > T_2 > T_1$) and log spot price. When the convenience yield is higher than the interest rate r_t , the futures price is lower than the spot price, which is called backwardation. On the other hand when the convenience yield is very low, the futures price is higher than spot price, which is called contango.

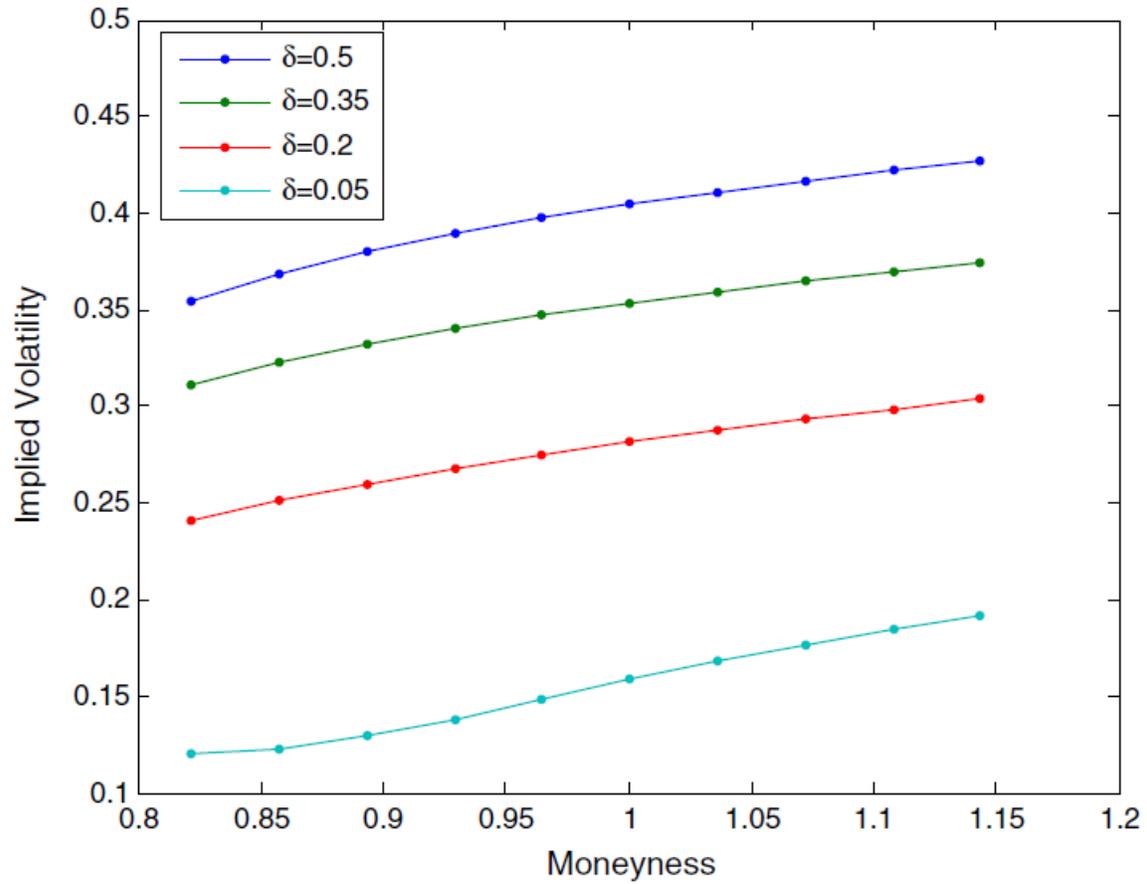


Figure 2. Implied volatility smiles inferred from our model. This figure shows an upward sloping smile pattern. The moneyness is defined as the ratio of futures prices to the option strike prices.

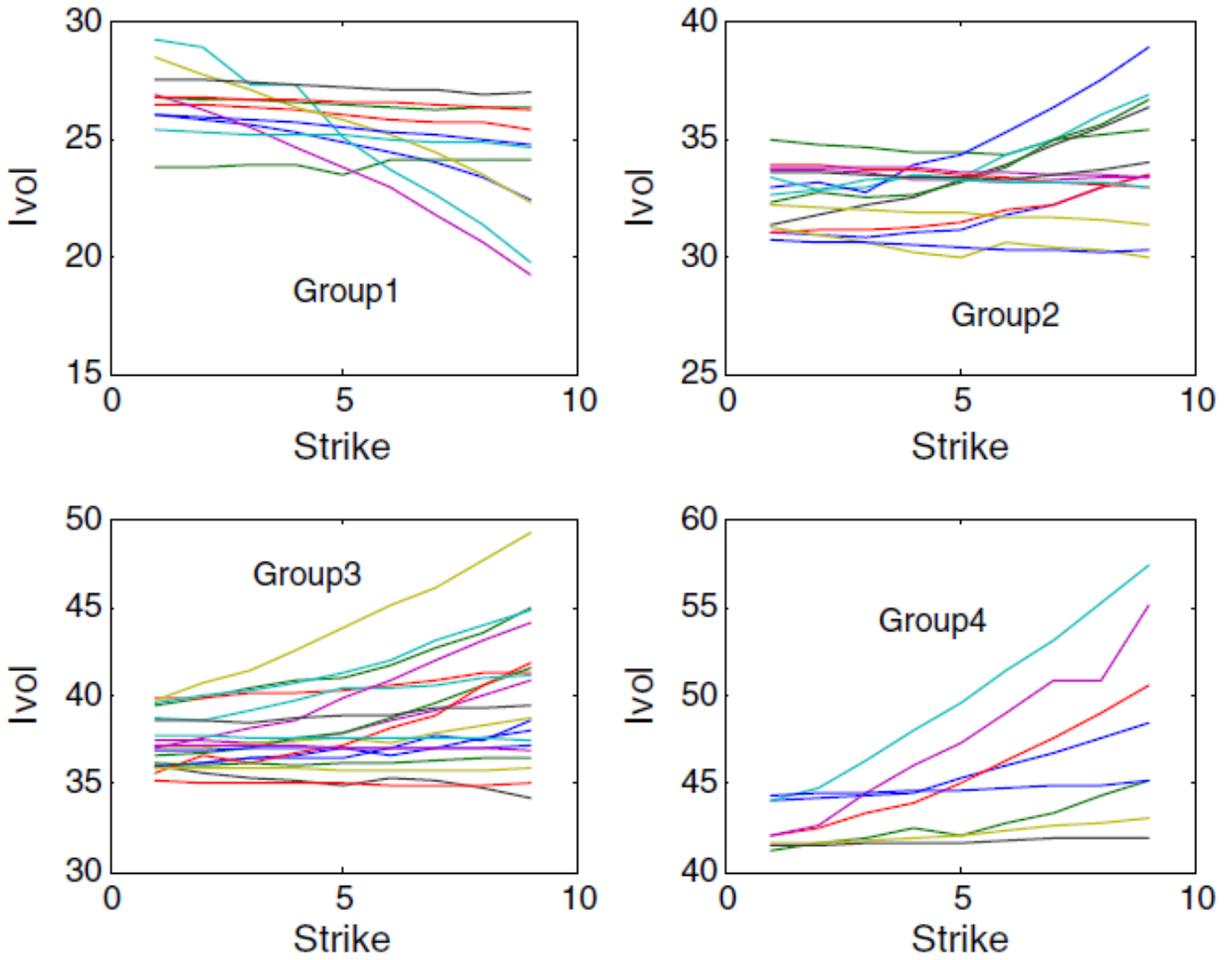


Figure 3. The market observed implied volatility for oil futures. The implied volatilities are grouped by their values. The majority of smiles show an upward sloping pattern.