Forward Copper Price Models
A Kalman Filter Analysis

Gordon Sick, PhD*
Mark Cassano, PhD†

July 9, 2012

Contents

1 Introduction 2

2 Two-factor Mean Reverting Model 3
2.1 Futures Prices . . . 4

3 Kalman Filter Estimates from Historical Data 4
3.1 Long-run growth rate μ . . . 6
3.2 Strength of Mean Reversion and Half-life of Mean Reversion . . . 6
3.3 Long-run growth risk premium λ_1 . . . 6
3.4 Mean-reversion Risk Premium λ_2 . . . 6
3.5 Volatility of the Long-run Growth Rate σ_1 . . . 7
3.6 Volatility of the Mean-reversion Factor σ_2 . . . 7
3.7 Correlation between the Long-run and Shock Factors ρ . . . 7
3.8 Long Run Risk-neutral Growth Rate $\mu - \lambda_1$ . . . 7
3.9 Systematic Risk β_1, β_2 of the Risk Factors . . . 8

4 Time Series of the Latent Kalman Factors 10

5 Time Series of Actual and Kalman Model Prices 10

6 References 11

*Haskayne School of Business, University of Calgary
†Independent Consultant

1 Introduction

Copper is an important commodity to analyze for several reasons.

- It is used extensively in the construction of capital goods, so its spot pricing is a good indicator of worldwide capital construction and its forward prices are a good indicator of expectations of future capital construction.
- It is dense in terms of value per weight and is easy to transport worldwide with little concern for environmental damage, as happens with crude oil. Thus, it is not subject to regional supply bottlenecks, as we see with intercontinental natural gas or the current Brent-WTI crude oil spread.
- Its consumption is not subject to seasonal fluctuations, as occurs for natural gas or electricity.
- It has a liquid forward market with deliveries well into the future, unlike iron ore or steel.

This report analyzes copper forward markets. This analysis is useful for several reasons.

- The model provides forward prices that are consistent with current and historical forward prices.
  - It is estimated using time series data of historical forward and spot prices.
  - The model uses the cross-sectional term structure of forward prices at each of the historical and current dates.
  - It provides for mean-reverting shocks to supply and demand for copper.
  - It estimates a long-run growth rate for copper prices that is consistent with long-run equilibrium.
  - Estimates are provided for both empirical (true) probability measures, as well as for risk-neutral measures that give certainty equivalents, which are useful in derivatives and real options pricing.
- The model facilitates risk analysis, because it provides volatility estimates of the mean-reverting and long-run mean factor.
- The model is useful for real options models that estimate the value of investment opportunities and provide criteria for starting, delaying, expanding and abandoning projects.

In this report, we are particularly interested in studying the structural break in copper prices that occurred in early 2006. A copper producer or consumer at that time saw significant increases in copper prices and would be wondering about the future pricing before embarking on the construction of large mines or general capital goods that consume copper.

2 Two-factor Mean Reverting Model

We estimate a two-factor model for the spot price of copper S_t in the spirit of Schwartz and Smith (2000); Cortazar and Naranjo (2006). We use the notation of Cortazar and Naranjo (2006).

In order to keep commodity prices positive, we model the logarithm of spot prices $\log S_t$ at time t, which we measure in years. The log is the sum of growth, Brownian motion and mean reverting (latent) factors:

$$\begin{aligned} \log S_t &= \mu t + x_{1,t} + x_{2,t} \\ dx_{1,t} &= \sigma_1 dw_{1,t} \\ dx_{2,t} &= -\kappa_2 x_{2,t} dt + \sigma_2 dw_{2,t} \end{aligned} \tag{1}$$

where

$$\begin{aligned} \mu &= \text{mu} = \text{long-run growth rate of long-term price} \\ w_1, w_2 &= \text{Brownian motion random variables, normally distributed with no drift} \\ \sigma_1 &= \text{sigma1 is the volatility of the long-run factor} \\ \sigma_2 &= \text{sigma2 is the volatility of the mean-reverting shock} \\ \kappa_2 &= \text{strength of mean reversion} \\ \rho_{1,2} &= \text{correlation between the two random factors} \end{aligned}$$

The strength of mean reversion describes how quickly the mean reverting shock dissipates, and it is common to think of it in terms of the half life of the shock, which is $\log(2)/\kappa_2$.

This describes the true or empirical probability process for the spot price. Futures prices are expectations of future spot prices that incorporate a risk-premium. That is, futures prices are “risk-neutral” expectations of future spot prices. The mean growth rates of the true and risk neutral distributions a separated by risk premia λ_1 for the long-run growth and λ_2 for the mean-reverting shock. The risk premia can also be obtained from an asset-pricing model like the capital asset pricing model as[1]

$$\begin{aligned} \lambda_i &= \beta_i \times (\text{CAPM Market Risk Premium}) \\ \beta_i &= \frac{\text{cov}(\sigma_i dw_i, dr_m)}{\text{var}(dr_m})) \\ dr_m &= \text{rate of return on the broad stock market index} \end{aligned}$$

The CAPM market risk premium (MRP) is the difference between the expected rate of return on the stock market index and the risk less rate of return. It is commonly taken to be in the range of 5% to 8% per year.

[1] We can equivalently define

$$\beta_1 = \frac{\text{cov}(dx_1, dr_m)}{\text{var}(dr_m})$$

but we can’t directly estimate β_2 from the dx_2 because the mean reversion in x_2 means that there is not a constant proportionality relationship between dx_2 and dw_2.

2.1 Futures Prices

Futures prices are the risk-neutral expectations of future spot prices. Let $F(x_{1,t}, x_{2,t}, t, T)$ denote the futures price set at time t for a futures contract to deliver at time $T \geq t$ when the two latent factors are $(x_{1,t}, x_{2,t})$. Cortazar and Naranjo (2006) give this (for the two-factor model) as:

$$F(x_{1,t}, x_{2,t}, t, T) = \exp\left(x_{1,t} + e^{-\kappa_2(T-t)}x_{2,t} + \mu t + \left(\mu - \lambda_1 + \sigma_1^2/2\right)(T-t)\right) \times \exp\left(-\frac{(1-e^{-\kappa_2(T-t)})\lambda_2}{\kappa_2}\right) \tag{2}$$

3 Kalman Filter Estimates from Historical Data

Table 1 provides estimates of the Kalman filter model for the full sample of 29 July 1988 to 10 June 2011, as well as for the full sample broken into 4 equal sub periods. Standard errors of the estimates are provided for the full-sample estimates, in the second column. The estimated standard errors are surprisingly small and are similar to those reported for crude oil by Cortazar and Naranjo (2006). This is simply an artifact arising from the very large number of observations of the model (daily data for 5 or 20 years).[2]

In particular, the parameter estimates tended to vary across sub periods, which suggests that the accuracy implied by the standard errors is overstated.

Another robustness check for the model estimation is to re-estimate the model with new initial starting points, since it is a numerical routine to maximize a likelihood function. The final estimate of the long-run mean growth rate μ and its risk-premium λ_1 was quite sensitive to the choice of initial estimates, but the difference between the two, $\mu - \lambda_1$, was quite stable for various initial starting points. This difference is the risk-neutral growth rate $\hat{\mu}$, as used in the pricing of futures contracts

[2] With the very large number of observations, standard errors tend to be very small, since the square root of the sample size is in the denominator of the calculation.

This is something that statisticians have written about and it really comes about because people tend to be abusive in the interpretation of classical statistical inference. That is, when we calculate a standard error or p-value, we assume the model is correct (linear, etc), and use the standard error to assess the accuracy of the estimated value of a parameter. We take a null hypothesis about that parameter value (zero for no effect or 1 for no mean reversion, etc) and test the null hypothesis with the data. For example, we ask where the observed estimate of the parameter would be 95% of the time if the null hypothesis holds to get a 95% confidence interval. But, the crucial thing is that we are assuming that the rest of the model is correct (linearity, residuals uncorrelated with explanatory variables, asymptotically normal disturbances, etc).

To put the problem another way, suppose the correct model for y as a function of x is a quadratic function, but that we observe y with error and that we usually observe y on the upward-sloping part of the function. If we estimated it as a linear function with enough data, we will get arbitrarily small standard errors for the slope estimate when we have arbitrarily large samples. But, we should not use this apparently high accuracy of the estimate of the slope coefficient to conclude that the model is linear. Indeed, we should really be testing a larger model that includes second and third order terms. The appropriate test really is the null hypothesis that these extra explanatory terms are zero.

Table 1: Kalman Filter model estimates for NYMEX Copper futures prices for full sample and for subsamples. The standard errors are for the full sample estimates. The first estimates of $\mu, \lambda_1, \lambda_2$ are from the Kalman filter model. The second set of estimates are from a CAPM risk premium for the respective factors, assuming that the Standard and Poor 500 Index has a market risk premium (MRP) of 5% or 8%, respectively. This leads to alternative estimates of the long-term growth rate μ by adding the CAPM risk-premium λ_1 to the Kalman filter estimate $\hat{\mu}_1$ of the risk-neutral growth rate of the long-term factor. This estimate, denoted μ_1, is the continuously compounded geometric mean growth rate. It is the growth rate of the logarithm of the long-term factor. A more common interpretation of the expected rate of return is the expected one-year return after taking the exponential function, which entails an Itô correction for the non-linearity of the exponential function. This is denoted by $\mu_1 + \sigma_1^2/2$ and is a continuously compounded growth rate for the arithmetic mean growth rate.

	29-Jul-88 10-Jun-11	Standard Error	29-Jul-88 12-Apr-94	13-Apr-94 05-Jan-00	06-Jan-00 29-Sep-05	30-Sep-05 10-Jun-11
μ	3.76%	4.08%	8.23%	-20.16%	5.03%	-12.49%
κ	0.6579	0.0028	1.1896	0.7264	0.5210	0.2766
$\ln(2)/\kappa$	1.05		0.58	0.95	1.33	2.51
λ_1	8.39%	4.08%	9.12%	-23.84%	7.08%	-3.59%
λ_2	2.72%	0.95%	3.31%	5.96%	3.62%	-5.67%
σ_1	22.56%	0.00%	16.19%	20.95%	18.01%	44.19%
σ_2	15.58%	0.00%	21.26%	56.30%	17.46%	64.22%
ρ	-0.0500	0.0001	0.0262	-0.5687	-0.0401	-0.6500
Calculated values						
$\hat{\mu}_1 = \mu - \lambda_1$	-4.63%		-0.90%	3.69%	-2.05%	-8.89%
$\hat{\mu}_1 + \sigma_1^2/2$	-2.09%		0.41%	5.88%	-0.43%	0.87%
Betas on S&P500						
β_1	0.2289	0.0161	0.1083	0.0262	0.1520	0.3999
β_2	0.0540	0.0110	0.0704	0.0014	0.0669	0.0635
MRP SP500	5%					
Calculated values						
$\lambda_1 = \beta_1 \times \text{MRP}$	1.14%		0.54%	0.13%	0.76%	2.00%
$\lambda_2 = \beta_2 \times \text{MRP}$	0.27%		0.35%	0.01%	0.33%	0.32%
$\mu_1 = \hat{\mu}_1 + \lambda_1$	-3.49%		-0.35%	3.82%	-1.29%	-6.89%
$\mu_1 + \sigma_1^2/2$	-0.94%		0.96%	6.01%	0.33%	2.87%
MRP SP500	8%					
Calculated values						
$\lambda_1 = \beta_1 \times \text{MRP}$	1.83%		0.87%	0.21%	1.22%	3.20%
$\lambda_2 = \beta_2 \times \text{MRP}$	0.43%		0.56%	0.01%	0.54%	0.51%
$\mu_1 = \hat{\mu}_1 + \lambda_1$	-2.80%		-0.03%	3.90%	-0.83%	-5.69%
$\mu_1 + \sigma_1^2/2$	-0.25%		1.28%	6.09%	0.79%	4.07%

and options. Thus, the risk-neutral growth rate is well-identified by the Kalman filter model, even though the true mean and risk premium are not well-identified by the Kalman model.

Next, we provide a discussion of each of the parameter estimates.

3.1 Long-run growth rate μ

The long-run growth rate μ is the growth of the logarithm of the long-run mean price. We can see that it is estimated in the Kalman filter model as 3.76% (per annum) for the full sample, but with a standard error of 4.08%, so this is not a really accurate estimate. As noted above, this estimate depends strongly on the choice of initial conditions in the estimation procedure. Furthermore, the estimates vary widely across the subsamples.

We are more confident about the estimate of the risk-neutral growth rate, and we will discuss an alternative estimate of μ from this more accurate base, below.

3.2 Strength of Mean Reversion and Half-life of Mean Reversion

The next rows provide estimates of κ, which is the strength of mean-reversion of the shock, and the half-life of the mean-reverting shock, which is $\log(2)/\kappa$. The standard error of the estimate of κ is quite small, but the estimated value varies from 0.2766 to 1.1896 in the subsamples. The half-life of a shock is 1.05 years for the full sample and varies from 0.58 years to 2.51 years in the subsamples. Note that the half-life is increasing as we proceed to later samples, which suggests that the mean reversion in recent copper markets is not as strong as it was in earlier markets.

3.3 Long-run growth risk premium λ_1

As noted above, the Kalman filter model is quite sensitive to initial conditions when estimating this parameter and it varies strongly across subsamples. Thus, we do not regard the estimate as being reliable, even if the standard error is only half of the magnitude of the estimated parameter. Below, we will discuss the alternative approach to estimating λ_1.

3.4 Mean-reversion Risk Premium λ_2

This parameter is basically a technical parameter that is useful for short-term hedging with futures contracts, but is not terribly important for a firm making long-term capital budgeting decisions. The standard error of the parameter estimate is quite small (one-third of the size of the parameter), but it does vary across sub periods. We provide an alternative estimate and further discussion below.

3.5 Volatility of the Long-run Growth Rate σ_1

This is an important parameter, because the risk in the long-run growth is a major determinant of real option value and affects investment decisions with the real option rule. The overall volatility was estimated as 22.56%, and the standard error is infinitesimally small (which also happens in Cortazar and Naranjo (2006) for crude oil futures). The volatility is slightly higher than that of the overall stock market, but in the same range. It was stable across sub periods, but soared in the last sub period to 44.19%, which is a period where there are a structural breaks in the long-run growth factor x_1. Thus, the high volatility estimate in that sub period may simply reflect the structural change that occurs in that sample when the long-run mean soared, then collapsed and then soared.

3.6 Volatility of the Mean-reversion Factor σ_2

The mean reversion factor is slightly less volatile than the long-run growth factor. And, given the mean-reversion, any shocks from this factor dissipate over time. Thus, this factor does not tend to be as important in long-term capital budgeting decisions as the long-run volatility.

3.7 Correlation between the Long-run and Shock Factors ρ

Overall the correlation between the two risk factors is small and slightly negative at -0.05. But, in some of the subsamples, such as the second and fourth, it assumes a large negative value of less than -0.5. This means that the short-run mean reversion factor provided a natural hedge to the long-run shocks, but the hedge was short-lived because of the mean reversion.

3.8 Long Run Risk-neutral Growth Rate $\mu - \lambda_1$

The risk-neutral logarithmic growth rate for the long run factor was quite stable for various choices of the initial conditions of the estimation, so we place a greater degree of confidence in the accuracy of this estimate than we do of the components separately. Overall, the growth rate is negative at -4.63%. Note that the long-run risk-neutral growth rate became quite negative at -8.89% in the last sub period, which initialy seems to be an unreasonably bearish view on copper, given though copper prices were basically flat over that quarter.

The negative growth rates arise from the combination of our logarithmic transformation (which is concave) and the very high volatility of the long-run growth factor x_1 in the last sub period. By Jensen's inequality (or Itô's lemma), this has the effect of depressing the estimate of the long-run growth rate, since these are geometric mean growth rates.

In fact, it is more common to think of growth rates as arithmetic means, in which case we should add $\sigma_1^2/2$ to the growth rates by Itô's lemma,[3] as shown in the next line labelled $\mu - \lambda_1 + \sigma_1^2/2$. This increases the risk-neutral growth rates substantially to -2.09% overall and 0.87% in the final sub period.

3.9 Systematic Risk β_1, β_2 of the Risk Factors

The Capital Asset Pricing Model (CAPM) provides a measure of risk premium as a product of a risk-measure β and the overall Market Risk Premium for an index against which the beta is measured. We used the Standard and Poor 500 index as the market proxy, and measured the betas of the two factors x_1 and x_2 by regressing daily changes in estimates of the values of these latent factors on the S&P Index returns, just as one would do when estimating the beta of a stock.[4] The betas for this procedure are reported in Table 1 in the lower rows. The beta for the long-run factor is 0.2289 for the full sample and varies from 0.0262 to 0.3999. The last sub-sample has the largest beta, suggesting that the long-run risk in copper prices was becoming an overall macro-economic risk, rather than just a commodity-specific risk as in earlier sub periods.

We assumed a market risk premia of 5% and 8% for the S&P 500 Index, which can be multiplied by the beta to get an alternative estimate of the factor risk premia λ_1 and λ_2. The long-run risk premium is 1.14% in the full sample and varies from 0.13% to 2.00% in the subsamples. This risk-premium is not affected by the difficulties of the initial conditions of the Kalman estimation procedure and can be regarded as more accurate than the risk premia given higher above from the Kalman model, which was 8.39% for the full-sample long-run risk.

Thus, we can add this CAPM-derived risk premium to the risk-neutral growth rate obtained from the Kalman filter model to get a more reliable estimate of the long-run mean growth rate μ, which we denote by μ_1 in the table. We get a separate estimate for the two different MRPs of 5% (8%). For the overall sample, this growth rate μ_1 is -3.49% (-2.80% when the MRP is 8%).

It varies from -6.89% to 3.82% (-5.69% to 3.90%) in the subsamples, with the most negative observations occurring in the last sub period.

But, as noted above, we also need to add the Itô adjustment, which we only do for the long-run growth volatility. This yields the line labelled $\mu_1 + \sigma_1^2/2$. This would be the expected arithmetic growth rate for the sub periods and is a slightly negative -0.94% (-0.25%) in the full sample. In the sub periods it varies from 0.33% to 2.87% (0.79% to 6.09%).

[3] We are only performing the Itô adjustment for the long-term factor, since the Itô adjustment for the mean-reverting factor is proportionally smaller for long horizons, and most capital budgeting is focussed on the long-term factor.

[4] Since x_2 is a mean-reverting factor, this is not a well-specified regression, so we instead scaled the Brownian motion w_2 by its volatility σ_2 and ran the regression with this factor. There is no such problem with the long-run factor and either procedure will the same beta estimate for it.

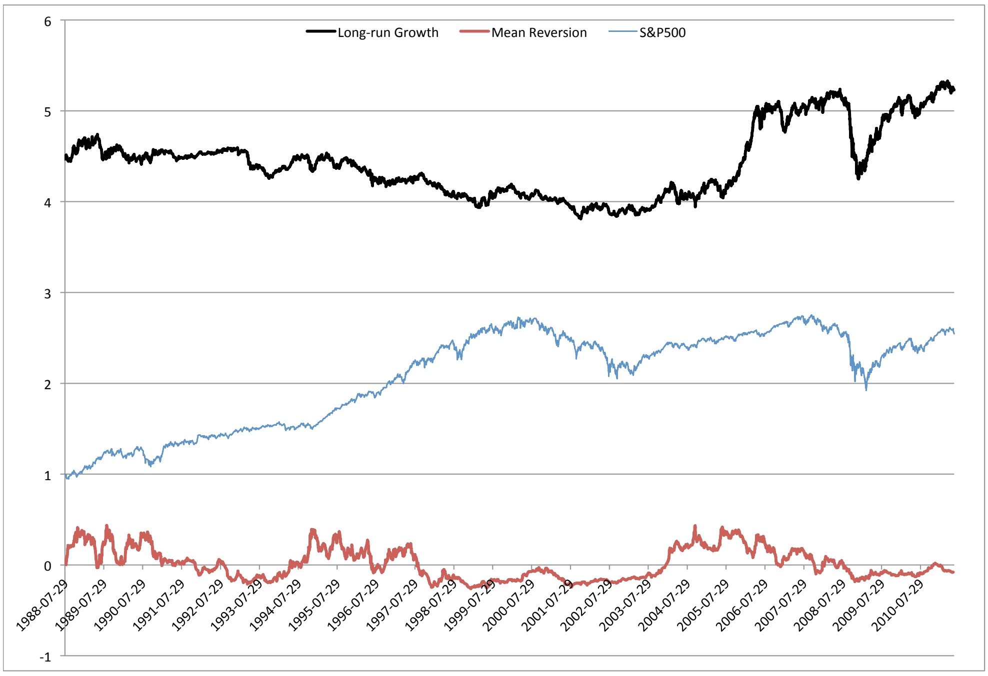

Figure 1: Time series of the latent factors of the Kalman filter model for copper prices. All factors are natural logarithms of the actual value. The long-run growth factor x_1 is the black graph varies between 3.8 and 5.4. The mean-reverting growth factor x_2 is the red graph reverting around the horizontal axis at the bottom. For comparison, the natural logarithm of the cumulative value of the S&P500 index (with dividends reinvested) is reported in blue, starting at an index level of e since $\ln e = 1$.

Once again, the high volatility in the last sub period increases the arithmetic expected rate of return. These are not large expected rates of return, but they are the returns for holding copper, which is an unlevered investment, while mining is levered by operating costs, which can increase the overall expected return on the project.

4 Time Series of the Latent Kalman Factors

Figure 1 shows the time series of the latent factors x_1, x_2 in the Kalman filter model. All factors are natural logarithms of the actual value. The long-run growth factor x_1 is the black graph, which varies between 3.8 and 5.4. The mean-reverting growth factor x_2 is the red graph reverting around the horizontal axis at the bottom. For comparison, the natural logarithm of the cumulative value of the S&P500 index (with dividends reinvested) is reported in blue, starting at an index level of e since $\ln e = 1$.

The long-run growth factor x_1 exhibits a strong upward structural break at the beginning of 2006. There is a drop in this factor through 2008, recovering in 2009, which is likely related to the global liquidity crisis of 2007–2009. This drop and recovery resemble mean-reversion, so we can see how there is less certainty about the true growth rate in that period. This mean reversion appears to have a longer half life than the mean reversion in the earlier periods. Indeed, there seems to still be the high-frequency mean reversion from the earlier periods in the last sub-period, but the large cyclical effect of the last sub period seems to take over the estimation. It may well be that the Kalman filter model would identify two mean-reverting factors for the last period. But, the extra mean reversion with the long half life in the last sub period may be spurious, because a similar cyclicality occurs for the S&P500 in the last sub period, and few financial economists would be prepared to say that a stock index like the S&P500 would have mean reversion.

Also, the upward shock and recessionary shock coincided with world-wide events, which explains the positive beta for those periods.

5 Time Series of Actual and Kalman Model Prices

We have prepared a Quicktime movie CopperStateVarMovie.mov, which shows the NYMEX futures prices as circles on a bi-weekly basis, along with the Kalman filter model futures curve as a solid curve. At the top of the figure is the date of the prices. The horizontal axis represents the numbers of years to delivery of the futures contract and the vertical axis is the copper price in cents per pound.

Once again, the structural break at the beginning of 2006 is very apparent. Also, in 2007, we can see that the Kalman model has difficulty tracking all the short-term shocks in copper prices. By July of 2007, the Kalman model deviates from the empirical futures prices at both the short end (mean-reverting shocks) and the long end of the futures. The deviation on the long end of the futures curve arises because the futures market became very bearish on copper at that point and endured a sharper correction than had previously been seen with the Kalman filter model. As a result, the Kalman model forecast higher long-run prices than the futures markets actually showed. This corrected itself by early 2008 as the futures data flattened out and approached the long-term curve that the Kalman model had shown earlier.

Also, note that in the 2006âĿ“-2008 period when spot prices were high, the NYMEX futures prices were much lower. This backwardation helped contribute to the negative risk-neutral growth rate that was estimated for the last quarter of the data period, which we have discussed earlier. Using the methodology of Subsection 3.9, this negative risk-neutral long-run growth rate contributes to a negative long-run growth rate, even when we add back the beta-determined risk premium. This partly explains why the last quarter of data shows a negative long-run growth rate, even though spot prices appear to be flat, on average, for the last quarter in Figure 1. The other explanation for the negative growth rate, as discussed above, is the high volatility in the last period, which pushes down the estimates of the logarithmic growth by Jensen's inequality.

Finally, it is interesting to note the relative performance of the long run copper price factor to the S&P stock index factor. From 1988 to 2000, the S&P index rose strongly, while the copper price fell somewhat. After 2000, copper started to outperform the stock index, leading to the sharp rise at the beginning 2006. Thereafter, the two tended to move closely in sync, since copper seemed to be proxying for the strong growth in emerging market economies.

This dichotomy between futures and spot prices highlights the value of using the Kalman filter approach to integrating futures and spot prices to get a holistic view of market estimates of future copper prices.

6 References

Cortazar, G. and Naranjo, L. (2006). An n-factor gaussian model of oil futures prices. *Journal of Futures Markets*, 26(3):243–268.

Schwartz, E. S. and Smith, J. E. (2000). Short-term variations and long-term dynamics in commodity prices. *Management Science*, 46(7):893–911.