

Markov modeling of Gas Futures

Leif Andersen

Banc of America Securities



February 2008

Agenda

- This talk is based on a working paper with two parts: i) theory for Markov representations of term structures of futures curves (w. mean reversion, stochastic volatility, jumps, and regime-switch); ii) an application for seasonal natural gas modeling.
- Here, we'll focus on ii), thereby providing a simple, specific example of i). Objective is to develop a practical trading model that represents the evolution of gas futures prices over time well.
- In particular, we want to model *seasonality* in: 1) futures levels, 2) implied volatilities, 3) correlations, and 4) implied volatility skews (!)
- ...and we want the model to have call option pricing formulas of the same complexity as Black-Scholes, with perfect fit to all ATM options
- ...and we want the model to have three or less Markov state-variables describing the entire futures evolution.
- ...and we want the model to imply nice stationary (in a seasonality-adjusted sense) dynamics of futures levels and volatilities.

Natural Gas Option Snapshot - I

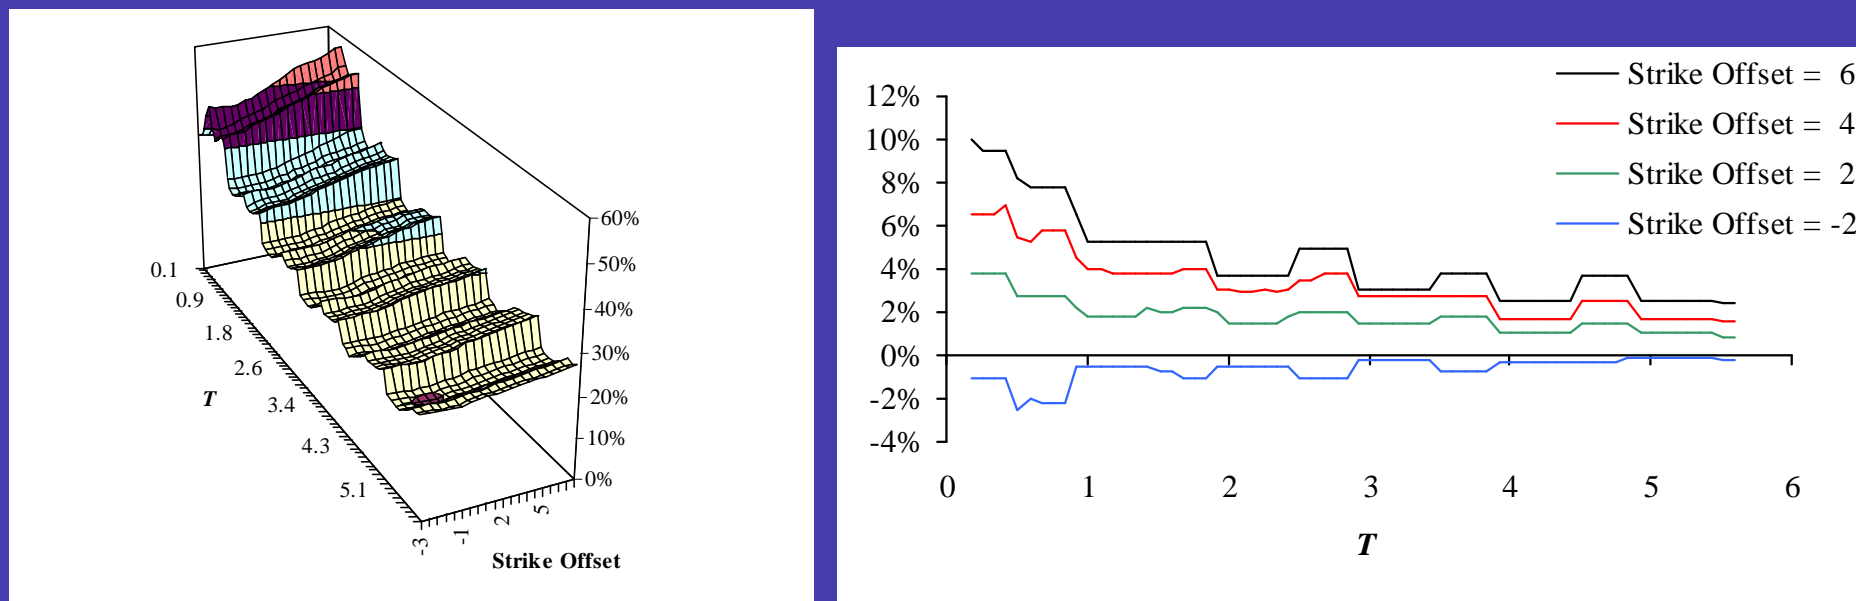
Figure 1: Futures Prices and ATM Volatilities, USD Gas Market



- **Notes:** The left panel shows the futures curve $F(0, T)$; the right shows implied ATM volatilities for European T -maturity call options. April 2007.
 - Gas futures prices strongly seasonal: high in winter, low in summer
 - Gas ATM implied volatilities strongly decaying in option maturity, with “jagged” seasonal overlay. Volatility approaches a constant plateau (with seasonality) as maturity gets large.


Natural Gas Option Snapshot - II

Figure 2: Volatility Surface



- **Notes:** Left panel: volatility smile, as a function of option maturity (T) and strike. Right panel: “skew” (volatility minus ATM volatility). Both panels: strike = $F(0, T) + \text{offset}$.
- A strong “reverse” volatility skew, which dies out as maturity is increased.
- Seasonality component, where skew is higher for winter delivery than for summer delivery.

Figure 3: Selected USD Gas Futures



- **Notes:** $F(t, t + \Delta)$, for $\Delta = \{1 \text{ month}, 13 \text{ months}, 25 \text{ months}, 61 \text{ months}\}$.
 - $F(t, T)$: time t gas price for delivery at time T .
 - Short-dated futures more volatile than long-dated (as expected).
 - The futures curve will occasionally (in winter) go into strong *backwardation*, after a rapid “spiky” increase in the short-term futures prices.

Correlation Analysis - I

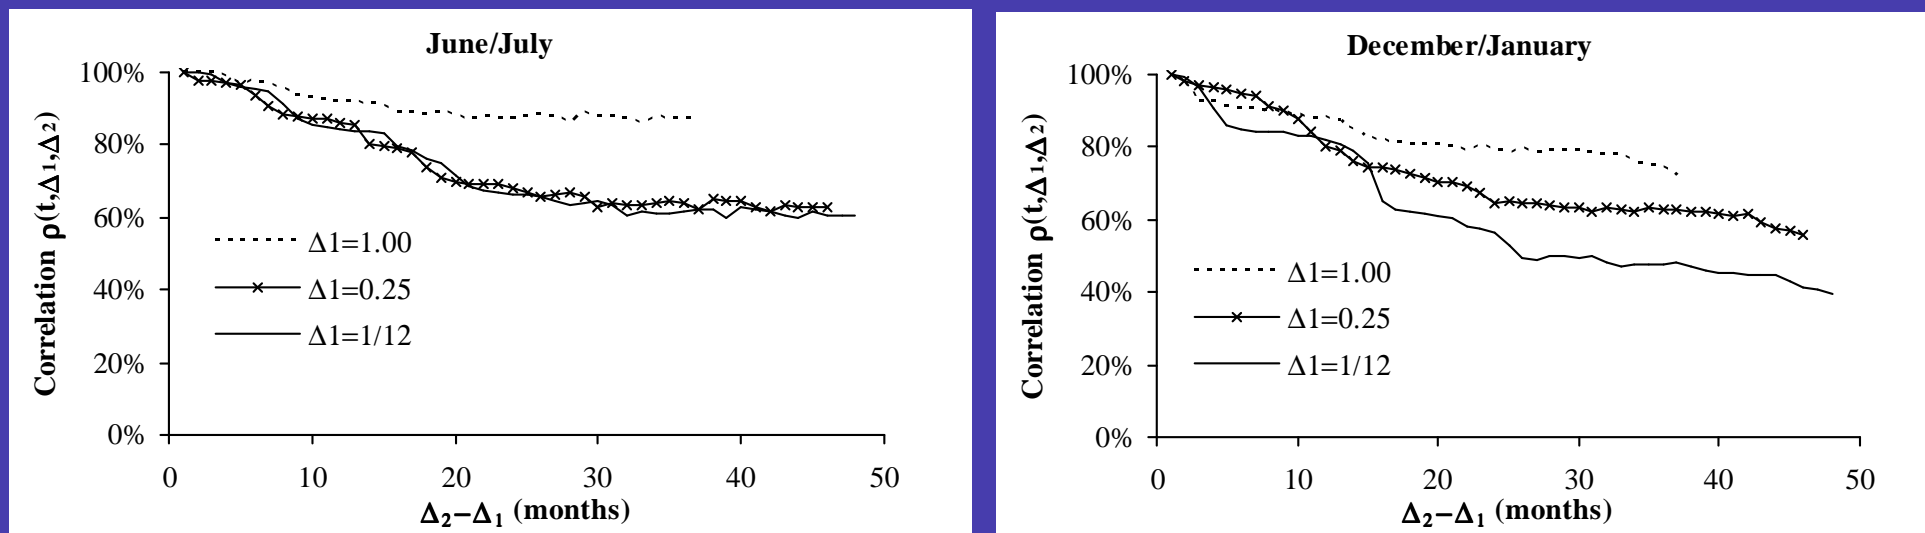
- Due to seasonality effects, care must be taken when attempting to measure futures correlation across maturities (the term structure of correlation)
- In particular, correlation structure depends strongly on the season of the observation intervals
- To get going, set $X(t, T) = \ln F(t, T)$ and define a correlation function

$$\rho(t, \Delta_1, \Delta_2) = \text{corr} (dX(t, t + \Delta_1), dX(t, t + \Delta_2)) .$$

- Seasonality effects will cause $\rho(t, \Delta_1, \Delta_2)$ to depend on t , for fixed time-to-maturity arguments Δ_1 and Δ_2 .
- We are also interested in the asymptote $f_\infty(t) = \lim_{\Delta \rightarrow \infty} \rho(t, t, t + \Delta)$.

Correlation Analysis - II


Figure 4: Empirical Correlation Structure for USD Gas Futures



- **Notes:** Daily time-series data covering January 2000 to June 2007.
 - Correlations between futures observed in the winter are generally lower than when observed in the summer.
 - Broadly speaking, the correlations $\rho(t, \Delta_1, \Delta_2)$ also tend to decline in $|\Delta_2 - \Delta_1|$ and, for fixed $|\Delta_2 - \Delta_1|$, increase in $\min(\Delta_1, \Delta_2)$.
 - But seasonal “undulations”.

Correlation Analysis - III

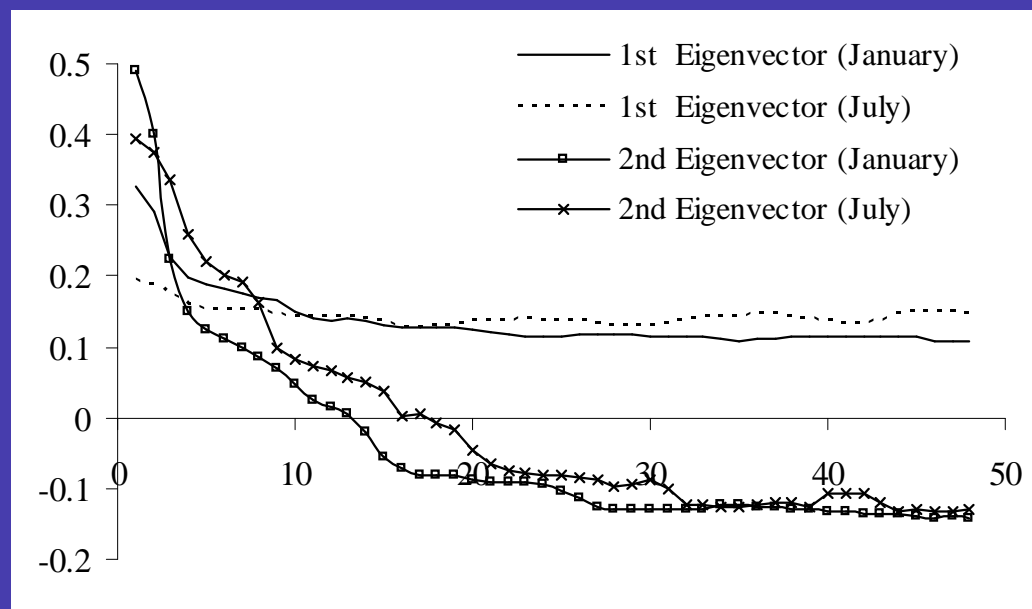
Figure 5: Empirical Correlation Asymptotes for USD Gas Futures



- **Notes:** $f_{\infty}(t)$, as a function of the calendar month to which t belongs. Daily time-series data covering January 2000 to June 2007.
 - The asymptotic correlation function $f_{\infty}(t) = \lim_{\Delta \rightarrow \infty} \rho(t, t, t + \Delta)$ is higher in summer than in winter.
 - And also “undulates” in the typical seasonal fashion.

Principal Components Analysis

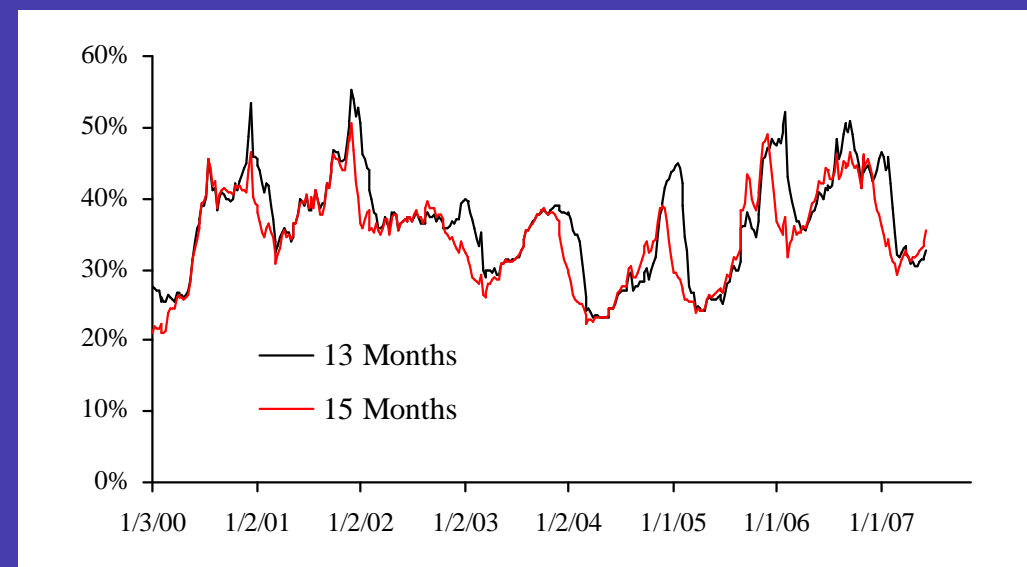

Figure 6: Principal Components Analysis for USD Gas Futures



- **Notes:** PCA of 48 futures price daily log-increments (January 2000 to June 2007).
 - Due to seasonality, we perform PC analysis on a month-by-month basis
 - Generally, first PC explains about 80% of futures curve variation; first two PCs explain about 95% of variation. First PC explains more variance in summer than in winter.

Implied Volatility Time Series

Figure 7: Selected USD Gas Implied Volatilities



■ **Notes:** Implied at-the-money volatilities for options on the spot gas price.

- As expected, short-term volatilities are consistently higher than long-term volatilities
- Options maturing a few months apart (not a multiple of one year) show a “lag” and some “dampening” in their implied volatility dynamics

Continuous-Time Model - I

- PCA: suffices to have two (indep.) Brownian motions W_1 and W_2 .
- Futures are martingales in risk-neutral measure, i.e.

$$dF(t, T)/F(t, T) = \sigma_1(t, T)dW_1(t) + \sigma_2(t, T)dW_2(t), \quad T > t.$$

- For now assume that σ_1 and σ_2 are deterministic such that $F(T, T)$ is log-normal, and put/call options can be priced by Black-Scholes.
- To ensure that the evolution of the entire futures curve can be represented through a few (two, in fact) Markov state-variables, let us specialize to

$$\sigma_1(t, T) = e^{b(T)} h_1 e^{-\kappa(T-t)} + e^{a(T)} h_\infty; \quad \sigma_2(t, T) = e^{b(T)} h_2 e^{-\kappa(T-t)}.$$

- Here κ is a mean-reversion speed; a and b are seasonality functions oscillating around zero at an annual frequency; h_∞ is the level of volatility for large maturities; and h_1 and h_2 are constants that together determine the short-term volatility.

Continuous-Time Model - II

- Defining $d(T) = a(T) - b(T)$, we have

$$dF(t, T)/F(t, T) = e^{a(T)} \begin{pmatrix} h_1 e^{d(T)} e^{-\kappa(T-t)} + h_\infty \\ h_2 e^{d(T)} e^{-\kappa(T-t)} \end{pmatrix}^\top d \begin{pmatrix} W_1(t) \\ W_2(t) \end{pmatrix}. \quad (1)$$

- So $a(T)$ is a persistent seasonality function, $d(T)$ is a transitory one.
- We note that the SDE (1) is of the *separable* type, in the sense that

$$dF(t, T)/F(t, T) = \beta(T)^\top \alpha(t)^\top d \begin{pmatrix} W_1(t) \\ W_2(t) \end{pmatrix},$$

with

$$\beta(T) = \begin{pmatrix} e^{a(T)+d(T)} e^{-\kappa T} \\ e^{a(T)} \end{pmatrix}, \quad \alpha(t) = \begin{pmatrix} h_1 e^{\kappa t} & h_\infty \\ h_2 e^{\kappa t} & 0 \end{pmatrix}.$$

- As shown in paper (unsurprisingly), this allows for a representation of all futures prices as functions of two Markov variables. We return to this later.

Basic Option Pricing

- Consider a T_1 -maturity option on a gas futures contract that matures at time $T \geq T_1$. The implied Black-Scholes term volatility is, at time $t < T_1$:

$$\sigma_{term}^2(t, T_1; T) = e^{2a(T)} \left\{ (h_1^2 + h_2^2) e^{2d(T)} \frac{e^{-2\kappa(T-T_1)} - e^{-2\kappa(T-t)}}{2\kappa(T_1 - t)} + 2h_\infty h_1 e^{d(T)} \frac{e^{-\kappa(T-T_1)} - e^{-\kappa(T-t)}}{\kappa(T_1 - t)} + h_\infty^2 \right\}. \quad (2)$$

- Normally, we calibrate the model to spot options, where $T = T_1$.
- An aside: swaptions can be easily priced, too, through the Markov representation we shall show shortly.
- Note that rolling futures price $F(t, t + \Delta)$ has volatility

$$\sigma_F(t, t + \Delta) \approx e^{a(t+\Delta)+d(t+\Delta)} \sqrt{(h_1 e^{-\kappa\Delta} + h_\infty)^2 + h_2^2 e^{-2\kappa\Delta}},$$

so as Δ is increased volatilities should – as seen earlier – be dampened due to mean-reversion and be “time-shifted” by $e^{a(t+\Delta)+d(t+\Delta)}$.

Correlation Structure

- Recall $\rho(t, \Delta_1, \Delta_2) = \text{corr}(d \ln F(t, t + \Delta_1), d \ln F(t, t + \Delta_2))$. In our model:

$$\rho(t, \Delta_1, \Delta_2) = \frac{e^{d(T_1)} e^{d(T_2)} e^{-\kappa(\Delta_1 + \Delta_2)} + q (e^{d(T_1)} e^{-\kappa \Delta_1} + e^{d(T_2)} e^{-\kappa \Delta_2}) + w}{\sqrt{e^{2d(T_1)} e^{-2\kappa \Delta_1} + 2q e^{d(T_1)} e^{-\kappa \Delta_1} + w} \sqrt{e^{2d(T_2)} e^{-2\kappa \Delta_2} + 2q e^{d(T_2)} e^{-\kappa \Delta_2} + w}},$$

where $T_1 = t + \Delta_1$ and $T_2 = t + \Delta_2$ and

$$q = \frac{h_1 h_\infty}{h_1^2 + h_2^2}, \quad w = \frac{h_\infty^2}{h_1^2 + h_2^2} = q \frac{h_\infty}{h_1}.$$

- For the case $\Delta_1 = 0$ and $\Delta_2 = \infty$, we get

$$\rho(t, 0, \infty) = f_\infty(t) = \frac{q e^{d(t)} + w}{\sqrt{e^{2d(t)} + 2q e^{d(t)} + w} \sqrt{w}}, \quad (3)$$

which depends on time through $d(t)$ only (not $a(t)$). If we know $f_\infty(t)$, we can use this to back out $d(t)$ analytically.

Parameter Reformulation

- The model so far has been parameterized through constants $h_1, h_2, h_\infty, \kappa$ and two seasonality functions. In a trading setting, it is sometimes easier to work through more “intuitive” parameters:

$$\begin{aligned}\sigma_0 &\equiv \sigma_F(t, t) = \sqrt{(h_1 + h_\infty)^2 + h_2^2}, \\ \sigma_\infty &\equiv \sigma_F(t, \infty) = h_\infty, \\ \rho_\infty &\equiv \frac{h_1 + h_\infty}{\sigma_0}.\end{aligned}$$

- ρ_∞ is the correlation function $f_\infty(t)$ when $d(T) = 0$. So: the seasonality-free correlation between the short and long futures prices.
- There is a one-to-one mapping between h_1, h_2, h_∞ and $\sigma_0, \sigma_\infty, \rho_\infty$. We use both representations interchangeably going forward.

Markov Representation of Futures Curve

- Define $X(t, T) = \ln F(t, T)$ and set

$$dx_1(t) = e^{\kappa t} (h_1 dW_1(t) + h_2 dW_2(t)), \quad dx_2(t) = h_\infty dW_1(t).$$

with $x_1(0) = x_2(0) = 0$. Then $F(t, T) = e^{X(t, T)}$, where

$$X(t, T) = \ln F(0, T) + e^{a(T)} \left(x_1(t) e^{-\kappa T + d(T)} + x_2(t) \right) - \frac{1}{2} e^{2a(T)} \frac{e^{2d(T) - 2\kappa T} (e^{\kappa t} - 1) (h_1^2 + h_2^2) + 4h_1 h_\infty e^{d(T) - \kappa T} (e^{\kappa t} - 1) + 2h_\infty^2 t \kappa}{2\kappa}.$$

- The state variables x_1 and x_2 are martingales, with the former having exponential variance. A better choice of state-variables:

$$z_1(t) = e^{-\kappa t} x_1(t), \quad z_2(t) = x_2(t),$$

$$dz_1(t) = -\kappa z_1(t) dt + h_1 dW_1(t) + h_2 dW_2(t), \quad dz_2(t) = h_\infty dW_1(t).$$

- SDE easy to implement by Monte Carlo or by 2-D finite difference methods. Entire futures curve can be *reconstituted* at time t from $z_1(t)$ and $z_2(t)$.

Calibration Algorithm

- By working in a futures curve setting (rather than with the spot gas price) our model is auto-calibrated to the seasonal futures curve. It remains to calibrate the model to ATM spot option volatilities and to correlation structure. (We deal with skew later).
- Here is an algorithm:
 1. Pick a value of ρ_∞ , based on empirical data.
 2. Set $a(T) = d(T) = 0$ (temporarily), and best-fit σ_0, σ_∞ and κ to the decaying ATM implied volatility term structure. This should give good stationarity properties.
 3. Decide if we want to model correlation seasonality. If no, set $d(T) = 0$; if yes, use (3) to set d from empirical observations. Some smoothing may be useful.
 4. Find the function $a(T)$ by matching perfectly ATM option volatilities, using (2). This can be done algebraically involving no root-search.
- Note: we may specify $\sigma_0, \sigma_\infty, ..$ to be functions of time.

Calibration Example (to Snapshot Data)


- We select $\rho_\infty = 0.5$ (Step 1); a subsequent fitting procedure (Step 2) then gives $\kappa = 1.35$, $\sigma_0 = 0.50$, $\sigma_\infty = 0.17$. We can compute that then $h_1 = 0.08$, $h_2 = 0.43$, and $h_\infty = 0.17$.
- To demonstrate the effects from the selection of $d(T)$, in Step 3 of our calibration algorithm we use two choices for $f_\infty(t)$: i) $f_\infty(t) = \rho_\infty = 0.5$, and ii) $f_\infty(t) = q(t)$, where $q(t)$ is a sine-function loosely fitted to the empirical data in Figure 5:

$$q(t) = 0.5 + 0.1 \sin(2\pi(t - 0.4)). \quad (4)$$

- Note: convention is that $t = 0$ is April 2007, such that $q(t)$ has peaks and valleys in summers and winters, respectively.
- (Note: if we want to ignore correlation completely, we can work with a single-factor model by setting $\rho_\infty = 1$)

Calibration Example - II


Figure 8: $a(T)$ for USD Gas Options, April 2007



- **Notes:** The graph shows two cases: i) $d(T) = 0$ and ii) $d(T)$ set to match the correlation seasonality implied by the function $q(t)$ above.
 - Seasonality function is close to stationary – a good sign.

Calibration Example - III

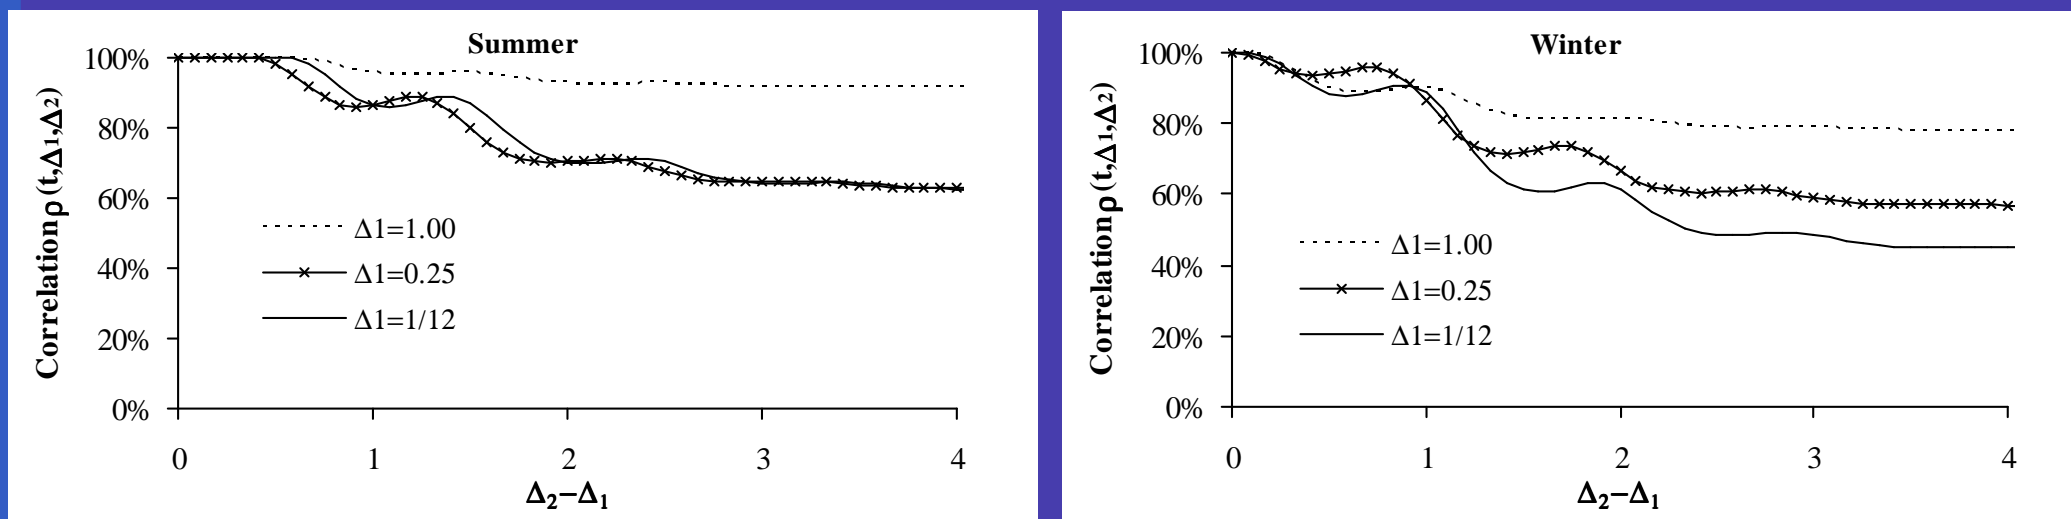
Figure 9: ATM Term Volatilities



- **Notes:** At-the-money term volatility $\sigma_{term}(t, T)$ as function of $T - t$ at two different points in time. For reference, the term volatilities corresponding to $a(T) = 0$ are also shown.
 - ATM volatility structure also close to stationary – after considering seasonality effects
 - Note: figure shows case where $d(T) = 0$. Even more stationary if $d(T)$ allowed to model seasonality in correlations (where $f_{\infty}(t) = q(t)$)

Calibration Example - IV

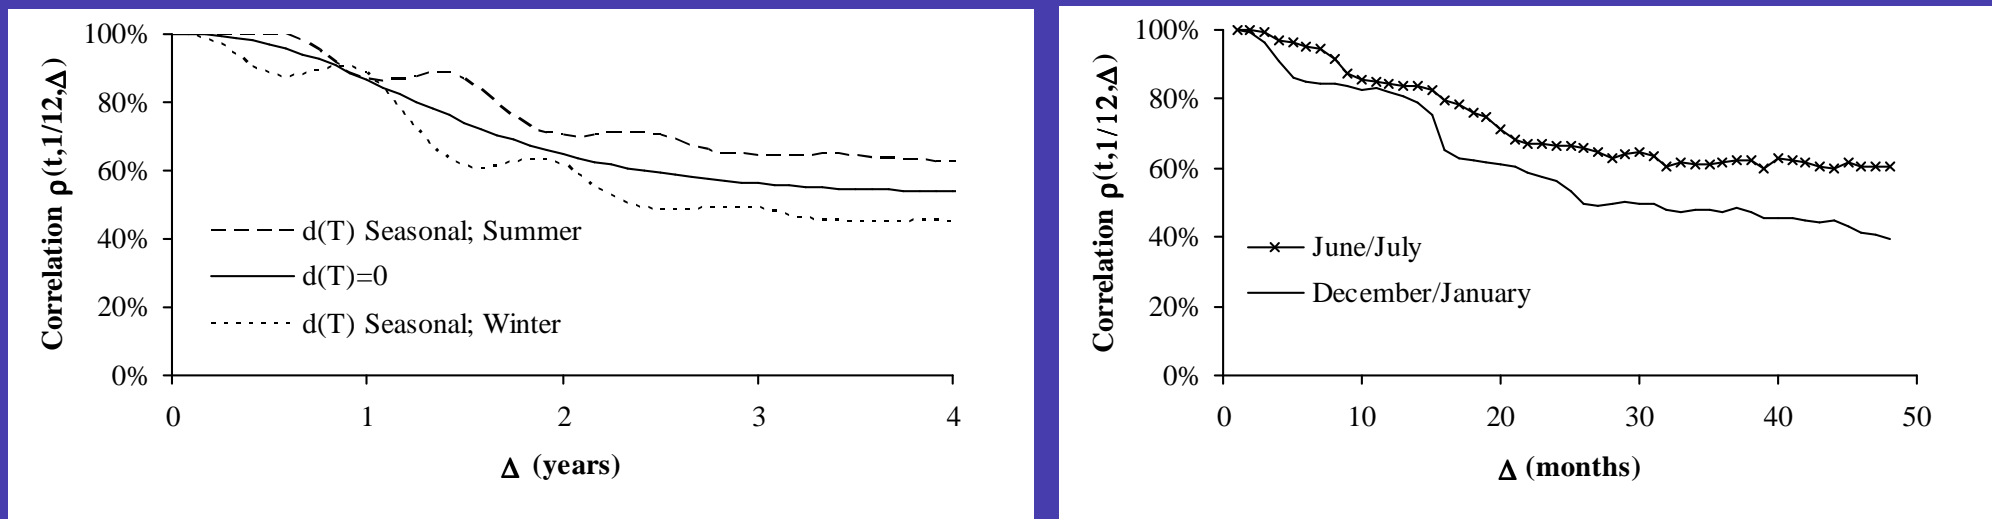
Figure 10: Model Correlation Structure



- **Notes:** Model-predicted correlation $\rho(t, \Delta_1, \Delta_2)$ for various values of Δ_1 and Δ_2 .
 - The correlation structure is qualitatively quite similar to empirical data in Figure 4.
 - The seasonality effects are somewhat more pronounced in model than in data; a closer fit can (if needed) be obtained by a more elaborate choice of d -function.

Calibration Example - V

Figure 11: Model vs. Empirical Correlation Structure



- **Notes:** Left panel: model-predicted values of $\rho(t, 1/12, \Delta)$ for $t = 0.1$ (summer) and $t = 0.6$ (winter). Right panel: empirical data.

Skew Modeling - Basics

- As discussed in paper, there are (at least) three approaches to skew modeling, all of which can be mixed-and-matched: **stochastic volatility** (positive correlation between spot and volatility); **mean-reverting jump-diffusion**; or **regime-switch** modeling.
- Stochastic volatility has some empirical support, but comes with complications: i) an increase in the number of state-variables (from 2 to 6, see paper); ii) the need for Fourier transformations when pricing puts/calls. i) is the main concern as it rules out finite difference grids. See paper.
- Strongly mean-reverting jump-diffusion can be used to model the upward winter-spikes in the time-series data, and adds (in its simplest form) only a single state-variable. However, put/call option pricing involves Fourier transforms which are quite time-consuming to compute unless mean reversion is zero (which precludes spike-behavior). See paper for details.
- Regime-switching is a natural alternative to mean-reverting jump-diffusion that, as we shall see, can allow for much simpler call/put pricing.

Two-State Regime Switch Model - I

- Consider a simple two-state model (paper has more complicated models)

$$F(t, T) = F_c(t, T)F_J(t, T), \quad (5)$$

where F_c is generated by the diffusion model from earlier, and F_J is a jump-martingale, i.d. of F_c , driven by a regime switch mechanism.

- The regime switch model has two states: $c = 0$ (“low”) and $c = 1$ (“high”). The “up” move (from $c = 0$ to $c = 1$) is characterized by an intensity $h_1(t)$; the “down” move is characterized by an intensity $h_0(t)$.
- Transition probabilities can be found by Markov chain methods (see paper). Define $\tilde{h}_i = \tilde{h}_i(t, T) = \int_t^T h_i(u)du$, $i = 0, 1$. Then

$$p_0^0(t, T) \equiv \Pr(c(T) = 0 | c(t) = 0) = (\tilde{h}_0 + \tilde{h}_1 e^{-\tilde{h}_0 - \tilde{h}_1}) (\tilde{h}_0 + \tilde{h}_1)^{-1},$$

$$p_0^1(t, T) \equiv \Pr(c(T) = 1 | c(t) = 0) = (\tilde{h}_1 - \tilde{h}_1 e^{-\tilde{h}_0 - \tilde{h}_1}) (\tilde{h}_0 + \tilde{h}_1)^{-1},$$

$$p_1^0(t, T) \equiv \Pr(c(T) = 0 | c(t) = 1) = (\tilde{h}_0 - \tilde{h}_0 e^{-\tilde{h}_0 - \tilde{h}_1}) (\tilde{h}_0 + \tilde{h}_1)^{-1},$$

$$p_1^1(t, T) \equiv \Pr(c(T) = 1 | c(t) = 1) = (\tilde{h}_1 + \tilde{h}_0 e^{-\tilde{h}_0 - \tilde{h}_1}) (\tilde{h}_0 + \tilde{h}_1)^{-1}.$$

Two-State Regime Switch Model - II

- To generate futures price dynamics from the Markov chain c , introduce a jump process J . When $c(t) = 0$, $J(t)$ will also be zero; however, if $c(t)$ jumps to 1 then $J(t)$ will simultaneously jump to a random value drawn from a Gaussian distribution $\mathcal{N}(\mu_J, \gamma_J)$.
- We assume that $\Pr(J(t) = 0 | c(t) = 1) = 0$, which requires that either $\gamma_J > 0$ or $\mu_J \neq 0$.
- For some freely specifiable deterministic function $s(T)$, we then finally set

$$F_J(t, T) = \frac{\mathbf{E}_t \left(e^{s(T)J(T)} \right)}{\mathbf{E} \left(e^{s(T)J(T)} \right)} \equiv \mathbf{E}_t \left(e^{s(T)J(T) - G(T)} \right), \quad (6)$$

where

$$G(T) = \ln \mathbf{E} \left(e^{s(T)J(T)} \right).$$

- Notice that $F_J(t, T)$ by construction is a martingale in the risk-neutral measure and that our scaling with $\exp(-G(T))$ ensures that $F_J(0, T) = 1$ for all T .

Two-State Regime Switch Model - III

- Since $c(t) = 1_{J(t) \neq 0}$, the regime-switch model above adds only a single Markov state variable ($J(t)$) to our setup.
- We already know how to reconstitute $F_c(t, T)$ from our two continuous state variables; we need to do the same for $F_J(t, T)$ given $J(t)$. The required result is below:

$$\begin{aligned} \mathbb{E}_t \left(e^{s(T)J(T)} \right) &= \mathbb{E}_t \left(e^{s(T)J(T)} | J(t) \right) \\ &= \begin{cases} \left(p_0^0(t, T) + p_0^1(t, T) e^{\mu_J s(T) + s(T)^2 \gamma_J^2 / 2} \right), & J(t) = 0, \\ e^{-G(T)} \left(p_1^0(t, T) + \left(p_1^1(t, T) - e^{-\tilde{h}_0} \right) e^{\mu_J s(T) + s(T)^2 \gamma_J^2 / 2} + e^{-\tilde{h}_0} e^{s(T)J(t)} \right), & J(t) \neq 0. \end{cases} \end{aligned}$$

- The same result can be used to establish $G(T)$, such that $F_J(t, T) = \mathbb{E}_t \left(e^{s(T)J(T) - G(T)} \right)$ can be computed in closed form.

Call Option Pricing in Regime Switch Model - I

- Call option pricing in the regime switch model is simple. To demonstrate, set the spot price

$$S(T) = F(T, T) = e^{-G(T)} F_c(T, T) e^{s(T)J(T)},$$

where $F_c(T, T)$ is known to be log-normal with mean $F(0, T)$ and term volatility $\sigma_{term}(0, T)$.

- Conditional on $c(T) = 0$, $S(T)$ is log-normal with mean and non-central 2nd moment

$$m_0 = e^{-G(T)} F(0, T), \quad s_0 = e^{-2G(T)} F(0, T)^2 e^{\sigma_{term}(0, T)^2 T}.$$

- Further, from known properties of log-normal distributions, conditional on $c(T) = 1$ we know that $S(T)$ is log-normal with moments

$$m_1 = e^{-G(T)} F(0, T) e^{s(T)\mu_J + s(T)^2 \gamma_J^2 / 2}$$

$$s_1 = e^{-2G(T)} F(0, T)^2 e^{\sigma_{term}(0, T)^2 T} e^{2s(T)\mu_J + 2s(T)^2 \gamma_J^2}.$$

Call Option Pricing in Regime Switch Model - II

- Therefore (ignoring discounting), for a call we get, assuming $c(0) = 0$,

$$\begin{aligned} C(0, T) &= \mathbb{E}((S(T) - K)^+ | c(T) = 0) p_0^0(0, T) + \mathbb{E}((S(T) - K)^+ | c(T) = 1) p_0^1(0, T) \\ &= p_0^0(0, T) \left(m_0 \Phi(d_+^{(0)}) - K \Phi(d_-^{(0)}) \right) + p_0^1(0, T) \left(m_1 \Phi(d_+^{(1)}) - K \Phi(d_-^{(1)}) \right) \end{aligned}$$

where

$$d_{\pm}^{(i)} = \frac{\ln\left(\frac{m_i}{K}\right) \pm \frac{1}{2} \ln\left(\frac{s_i}{m_i^2}\right)}{\sqrt{\ln(s_i/m_i^2)}}, \quad i = 0, 1.$$

- In the expression for $d_{\pm}^{(i)}$, we can use

$$\ln \frac{s_0}{m_0^2} = \sigma_{term}(0, T)^2 T, \quad \ln \frac{s_1}{m_1^2} = \sigma_{term}(0, T)^2 T + s(T)^2 \gamma_J^2.$$


- The expression for call options if $c(0) = 1$ is simple, too – left to the audience (or see paper)

Calibration Example - I

- The combination of regime-switching with continuous dynamics give us a range of parameters (h_1 , h_0 , μ_J and σ_J) that we can use to calibrate the volatility smile surface.
- On top of this, we now have *three* sources of seasonality in the model: the functions $a(T)$, $d(T)$, and $s(T)$. [We can also make h_1 and h_0 dependent on calendar time, but this does not give much seasonality effect]. In practice, we would normally parameterize two of these functions directly (most likely d and s), and solve for the last one to perfectly match at-the-money volatilities.
- The seasonality in the volatility skew originates with the function $s(T)$.

Calibration Example - II

Figure 12: Skew Fit for USD Gas Options



- **Notes:** Volatility skew (= difference between the implied volatility minus ATM volatility) as a function of option maturity T . The function $s(T)$ was a simple periodic function; $a(T)$ was set to provide a perfect fit to the ATM volatilities in Figure 1.
- Fit is decent, particularly given the somewhat rough market data.
- A better fit will require a more complicated s -function, and/or stochastic volatility.

Conclusion

- We have shown a simple, practical model for the evolution of seasonal futures prices. Model has general applicability (oil, gas, electricity,..); we focused on its fit to gas
- With *one* factor, the model can handle seasonality in futures prices and in ATM volatilities; with *two* factors, the model can also handle seasonality in correlations; with *three* factors, the model can also handle seasonality in the volatility skews
- More factors can be added for additional realism, at the expense of tractability.
- For this, and for numerical methods, see paper.