KIODEX

UNDER THE HOOD

The Models and Analytics that Power The Kiodex Risk Workbench^{SM}

SEPTEMBER 24, 2001

3 NEW YORK PLAZA ▪ 15TH FLOOR ▪ NEW YORK, NY 10004 ▪ TOLL FREE: 877.6.KIODEX ▪ FAX: 646.437.3910 ▪ www.kiodex.com

TABLE OF CONTENTS

1. ABSTRACT

This document gives a basic introduction to the classes of financial derivatives that are currently priced and risk-managed by the Kiodex Risk Workbench^{SM}. We describe the market dynamics that need to be and are captured in the Kiodex Multifactor Model, the features of the model itself, and how it accommodates volatility skew. Also included is a discussion of how we calculate risk (including Front Month Equivalent Delta), Profit and Loss (P&L), and the breakdown of the P&L due to moves in the different market factors. We close with a description of our testing procedures and with a view of what the future holds.

2. PRICING METHODOLOGY

2.1 Properties of Energy Underlyings

Kiodex currently maintains data on natural gas, crude oil and refined petroleum products. Different geographical locations and delivery points correspond to different assets. This means that Kiodex supports about 800 different underlyings in these markets.

2.1.1 Futures vs. Publication Data

The two types of underlying data currently supported by Kiodex are futures prices (WTI, Heating Oil, etc.) and publication prices (Gas Daily, IFERC, etc.). Futures prices are used directly in the pricing of derivatives, whereas publication prices, which are applicable over periods of different durations, must be converted into forward curves prior to being used in a pricing model. We discuss the forward building methodology in section 2.3.

2.1.2 Seasonality

Seasonality effects are incorporated in the forward curves by assuming that a multiplicative deterministic seasonality function is embedded in the forward curve. The dynamics of the model described in section 2.4 govern the non-seasonality component of the forward curve. The seasonality component is then superimposed (multiplicatively) in order to obtain the correct forwards so they can be passed to our pricing engines.

2.1.3 Many-Factoredness of the Forward Curve

Futures curves are built out of a large number of futures contracts. For instance, the crude oil NYMEX WTI futures curve is made up of 36 futures, indicating that its dynamics are described by 36 factors or, equivalently, 36 sources of uncertainty. In commodities markets such as natural gas, crude oil and refined petroleum products, the front end of the curve is much more volatile than the back end of the curve. This leads to the futures curve alternating between states of *backwardation* (where the futures prices in the near contracts exceed the futures prices of those contracts well into the future), *contango* (the opposite of backwardation) and/or a mixture of the two. The same is true for publication curves.

Note that the futures in a single curve are typically highly correlated in commodities. This opens up the possibility of capturing the dynamics of such a curve with much less than 36 factors. In fact, the Kiodex Multifactor Model uses 2 factors, as discussed in Section 2.4.1.

2.2 Pricing Forwards, Futures and Swaps

2.2.1 Forwards

The price of a forward contract depends on the payment date for the forward contract, the level of interest rates for the payment date, the price of the underlying futures contract or appropriately derived publication price, and the strike.

2.2.2 Futures

The price of a futures contract is given by the market. A synthetic futures price can be calculated from the forward curve.

2.2.3 Swaps

Swaps are priced using futures or publication prices. The fixing on a swap typically takes place over a period of time in order to construct an average. Various averaging conventions are available, such as, but not limited to, *business daily, calendar daily, last day, penultimate,* etc.

2.3 Building Forward Curves

Forward curves must be built up from information from futures contracts or from publication data. Forward curves built from futures contracts are the easier of the two to build. Since it does not cost anything to enter into a futures contract, the drift of the futures price of the contract is zero [1]. This means that the forward curve *for that specific contract* is flat from today till the expiry of that contract, and the constant value of that forward curve is equal to the currently observed futures price of that contract. There is one such forward curve for each futures contract thereby giving rise to a forward *surface*. A point on the surface is given by specifying both the point in time on which an observation is made and what contract is being observed. When constructing averages in order to price swap related instruments, the proper cross section of that surface is taken depending on the deal attributes to form the resulting forward *curve*.

Constructing forward curves based on publication data is more involved since the structure of the data received from market participants, brokers and data vendors is typically more varied. The swap quotes may be monthly for the first few months, then quarterly, followed by yearly, or some subset thereof. Kiodex uses a piecewise-polynomial interpolation procedure for converting this information into a forward curve consisting of implied monthly quotes that reprice the input data exactly. The interpolation procedure has been designed so that a change to one of the building blocks of the

---

[1] The Kiodex Multifactor Model introduced in section 2.4 is internally consistent with this fact.

forward curve only affects the latter locally. This feature ensures the absence of spurious numerical spikes when calculating risk as described in section 3.1.

The forward curves described above are inputs to our pricing models.

2.4 Pricing Options

The Kiodex Risk Workbench currently prices European and American options on any single futures contract or publication, and Asian options whose underlyings are averages composed of settings on either a rolling set of futures contracts, prompt or not, or publications. Options on swaps will be supported in a later release.

As mentioned in section 2.1.3, futures curves for gas, crude oil and refined petroleum products tend to fluctuate between backwardation and contango, exhibiting a rich dynamical behavior. The Kiodex Multifactor Model (KMM), discussed in this section, successfully captures the dynamics of the futures curve while reducing the number of independent factors to two. It also accounts for volatility skew (see section 2.4.3). This model extends an approach developed by Jacques Gabillon in “The Term Structure of Oil Futures Prices”, Oxford Institute for Energy Studies, 1991.

The following sections describe the parameterization and calibration of the KMM.

2.4.1 Parameterization of the Futures Curve

The observation that the front end of the futures curve is typically more volatile than the back end indicates that the futures curve may be decomposed into two components responsible for its main modes of variation. The KMM makes the assumption that the dynamics of the futures curves are driven by two factors, namely the spot price S and long-term price L of the underlying commodity. S and L correspond effectively to each extremity of the curve. This model is rich enough to recreate the behavior of the term structure of the underlying.

The factor S obeys a locally mean-reverting process while L is lognormal. Their joint distribution is given by the short volatility σ_S, the long volatility σ_L, and the correlation ρ. We call σ_S, σ_L, and ρ the KMM parameters. For instance, setting the correlation ρ to one ensures that both the front- and the back-end of the curve always move in the same direction. Or, setting the long volatility σ_L to zero causes the back-end of the curve to be fixed. The KMM parameters are not frequently updated, as they are representative of the semi-stationary properties of the underlying futures curve. However, they need to be changed from time to time, thus impacting option valuation. The risk associated with the changes in values of those factors is reported by the Kiodex Risk Workbench.

Information about the futures curve as a whole is implied by parameterizing the futures curve as a function of S and L. Following Gabillon, the price of a futures contract with maturity T observed at time t, is parameterized as follows:

$$F_{t,T}(S,L)=A_{t,T}S^{B_{t,T}}L^{1-B_{t,T}}$$

where $A_{t,T}$ and $B_{t,T}$ are parameterization functions obtained by calibration to futures and options prices in the market. This parameterization allows one to determine from S and L the dynamics that is followed by the futures $F_{t,T}$. The probability distribution for $F_{t,T}$ is lognormal. Options on futures can then be priced using the standard Black-Scholes framework, but only after the appropriate volatility has been extracted from the KMM. The volatility that is calculated is therefore appropriate for the type of underlying applicable to the trade: it is consistent with the dynamics of the futures curve, it is an early-expiry volatility for an option that matures prior to the final fixing of its underlying, and it is consistent with the skewness of the underlying distribution (see section 2.4.3).

2.4.2. Calibration

The calibration procedure ensures that the KMM prices market instruments accurately. The instruments to which we calibrated are futures contracts and liquidly traded options on futures. Depending on the commodity, these options are either European (only exercisable at maturity) or American (exercisable at all times during the life of the option). Given implied volatilities for a set of liquidly traded options, the KMM solves for the parameterization function $B_{t,T}$ which uniquely determines the volatility surface $\sigma(t,T)$ for all t and T, with $t \leq T$ where $\sigma(t,T)$ is the out-turn volatility for a European option expiring at t on a futures contract maturing at T. Note that the volatilities are affected by the KMM parameters σ_S, σ_L, and ρ.

2.4.3 Skew

Skew is a measure of the non-lognormality of the underlying. This phenomenon is illustrated by the so-called volatility smile describing the commonly observed dependency of volatility on the strike of the option, all other parameters kept equal. The KMM accepts market-implied volatilities for out-of-the-money options of all maturities. Currently, the out-of-the-moneyness is expressed by the delta of the put or call option, which is then translated into a strike. The out-of-the-money vols on either side of the at-the-money-forward volatility are referred to as volatility wings.

To accommodate skewed underlyings, the KMM generates a strike-adjusted volatility and treats the underlying as lognormal. This approximation yields good results for plain-vanilla options, but should be used carefully for highly path-dependent options such as barrier options. We will address these issues in a future release when we support barrier options.

2.4.4 European Option Pricing

European options are priced using the Black-Scholes formula with forward and volatility information obtained from the calibrated market info.

2.4.5 American Option Pricing

American options are priced using the Whaley analytic approximation with forward and volatility information obtained from the calibrated market info. This approach is computationally efficient at the cost of a small bias in valuation. However, more sophisticated yet slightly slower numerical pricing techniques will be available in a forthcoming release.

2.4.6 Asian Option Pricing

In Asian options, the payoff is a function of the underlying price averaged over a period of time. The average of lognormally distributed prices is not itself lognormally distributed. The KMM approximates the average's distribution with a lognormal distribution by matching its first two moments, which are mean and variance. Once the effective forward and volatility of the average have been calculated, the Asian option calculation proceeds by inserting those quantities into a European option calculator.

3. RISK CALCULATIONS

3.1 General Methodology

We make available several types of risk measures all of which are discussed in this section: the "Greeks", Front Month Equivalent Delta, and P&L Breakdown.

3.1.1 Standard Greeks

The "Greeks" are a set of risk measures that indicate the sensitivity of the value of a trade (or portfolio of trades) with respect to movements in market variables. More simply, they indicate how quickly a portfolio value will move for a change in, for example, interest rates, vols, futures prices or any relevant underlying prices. The Greeks are a natural outgrowth of the concept of dynamic hedging.

The Greeks currently available through the Kiodex Risk Workbench are:

- *Delta* – The rate of change of an instrument's value with respect to the underlying futures or publication price.
- *Gamma* – The rate of change of an instrument's delta with respect to the underlying futures or publication price. Kiodex adopts the convention to scale the mathematical gamma by $S/100$, where S is the price of the underlying, so that the Kiodex gamma can be interpreted as the amount by which delta changes for a *relative* 1% move in the underlying's price.
- *Vega* – The rate of change of an instrument's value with respect to the underlying futures or publication volatility. Kiodex adopts the convention to scale the mathematical vega by 1/100, so that the Kiodex vega can be interpreted as the amount by which the instrument value changes for an *absolute* 1% move in the underlying's volatility.
- *Theta* – The rate of change of an instrument's value with respect to time.

There are other Greeks; *rho* is an example, which is the change in the value of a financial instrument resulting from a change in interest rates. Although rho is an important risk measure in general, its role in commodities markets is less pronounced than for other markets such as fixed income, foreign exchange and equity.

Given a portfolio P whose value depends on the underlying price S, the delta is obtained by first varying S by an amount δS and then calculating the change in value $\delta P = P(S+\delta S)-P(S)$ of the portfolio. The resulting delta is given by

$$\text{delta}=\delta P/\delta S$$

and will depend on how big δS is chosen. As δS gets smaller, the ratio $\delta P/\delta S$ converges to the instantaneous sensitivity of P with respect to S. For numerical computations, δS is typically chosen to be of the order of $S/10{,}000$. While it is sometimes desired to calculate δP when δS is large (several percent of S or more) for the purposes of scenario analysis, the resulting ratio $\delta P/\delta S$ may be very different from the instantaneous rate of change and cannot be used for dynamic hedging.

In order to calculate the four Greeks above, we take the simple and unified approach of *bumping*. For delta and gamma, we superimpose on the relevant underlying prices a small change and calculate δP. For example, for a strip of futures we bump each futures price one at a time and calculate the resulting delta and Kiodex gamma of the portfolio for each separate bump. This allows the risk to be viewed with respect to individual futures. Note that bumping of the underlying may have consequences on calibrated quantities, such as volatility skew, which contribute to the change in value of the portfolio. For instance, when quoting volatility skew as a function of an option delta, a change in the underlying price affects the volatility for a given strike and a given maturity. Thus, the bumping approach ensures that skew effects due to calibration are captured when calculating delta and Kiodex gamma.

Similarly for vega, volatilities are bumped across all strikes and all option maturities in the vol surface so that vega can be broken down in maturity buckets. For theta, the bump is equivalent to a change in time equal to 1 day, although weekends, but not holidays, are excluded.

3.1.2 Front Month Equivalent Delta (FME Delta)

The Front Month Equivalent denotes a hedging strategy whereby a direct risk exposure to a commodities futures contract is hedged by taking a position in the prompt contract rather than in the contract to which the risk is directly exposed. The FME hedging technique is a useful tool when contracts further into the future are less liquid than the prompt contract, often prompting risk managers to decline taking positions in the less liquid contracts and instead taking some “equivalent” position in the more liquid prompt contract.

There are two ingredients that determine the size of the equivalent position to take in the prompt contract: the relative difference between the vols of the prompt and any illiquid contract, and the correlation between those contracts. This can be understood as follows: assume for the moment that the contracts are perfectly correlated. Then if the prompt contract volatility is for instance double that of the volatility of the illiquid contract, it follows that the size of the equivalent position of the prompt contract need only be half as big as the size one would have needed for the position in the illiquid contract. However, this is only true if the correlation is perfect. A less than perfect correlation would determine a different size for the equivalent hedge. In fact, if the two contracts are not correlated at all, then regardless of the volatilities, it makes no sense to interchange positions between the contracts. So the FME technique only works if the illiquid contracts and the prompt contract are very highly correlated.

Please note that however close the correlation is to being perfect, if it is not indeed perfect, then the FME hedge is also not perfect.

3.2 P&L and the P&L Breakdown

Given “before” and “after” snapshots of a portfolio, Kiodex breaks down the P&L between these two-time snapshots in order to identify what the contribution to the P&L is for each of the market variables that affect the portfolio’s price. A portfolio consisting of trades on WTI and Brent will, in general, experience P&L arising from movements in WTI, Brent, interest rates and volatilities and the P&L breakdown will display how much of that P&L is due to movements in each.

If there were no market moves whatsoever, and only the passage of time affected the P&L, then it should be true that the time component of the P&L breakdown must match the theta of the “before” snapshot. The P&L algorithms at Kiodex obey this consistency condition. We calculate the time component of P&L by superimposing the market configuration of the “before” snapshot onto the “after” snapshot and then we compare the price difference of the resulting “after” snapshot with the “before” snapshot. Conversely, to extract the contribution of a single market parameter, like interest rate movements to the P&L, we superimpose the interest rate market data of the “after” snapshot onto the “before” snapshot and measure the affect of this change on the “before” snapshot.

A mathematical analysis of this approach to P&L breakdown shows that these individual contributions will sum to the total P&L if the nonlinear effects between all the movements in market parameters is small. Otherwise, those nonlinear effects must also be quantified by measuring the gamma and cross-gamma risk. The Kiodex Risk Workbench currently computes all these “higher order” gamma and cross-gamma effects as one lump sum referred to as cross-risk, but does not break them down any further in Kiodex’s initial release of the Risk Workbench (v1.1). The higher order breakdown will be available in a forthcoming release.

4. TESTING

4.1 Accuracy

Every valuation tool in the Risk Workbench is tested on trades that are new, in-the-life, and expired. This ensures that users can rerun reports on days in the past, and if necessary obtain cash flows from trades that have generated payments in the past. In order to validate the correctness of the calculators and risk reporting engines, the values produced by every single operation that can be performed by the calculation engines of the Risk Workbench are recomputed independently over a range of parameters applicable to whatever operation is being tested. Answers must match.

We begin our testing at the most basic level: date manipulation routines and numerical routines (root finders, distributional related calculations, etc.) that form the basis for the calculators that will be tested later on. Next in the hierarchy are the libraries that deal with the construction of interest rate curves and calculation of discount factors, and the construction of forward curves built from exchange-based data and publication-based data. These objects are used to build our asset and market objects on which our calculators depend. In our tests, the asset and market objects are applied to all the commodities we support, and again, all answers are recalculated by independent means and required to match those results produced by the asset and market objects. We continue this process on through the calculators and result aggregation. Finally, we move on to risk reports and test all those reports we support, recalculating independently the risk that should be produced for portfolios that contain forwards, futures, swaps, and options. We require that all results reconcile.

Although the testing scheme described above validates the mathematical accuracy of a large number of trades in all the markets we support, it is not a regression test. A regression test is useful for checking, for instance, that the values of a large number of trades priced on all days in a given year do not change their values in unexpected ways as development continues on extending the functionality of the Kiodex Risk Workbench. We run our regression tests every day to monitor any such value changes. If a change is detected and the change was not anticipated, we find out what happened and address the issue.

All these tests are carried out in Kiodex's modeling environment. Accuracy tests and regression tests are then carried through a second round on the entire Kiodex Risk Workbench as a whole each night to ensure no spurious changes arise when the system is run as an integrated unit.

4.2 Speed

As a general indication, we can produce a P&L report on 250 different trades with 6-month maturities with daily averaging in less than 10 seconds. This exercises our forward curve building, vol grid calibration, pricing calculators and P&L calculation and reporting engines. Risk will take longer because the individual market parameters need to be bumped individually and for each such bump the portfolio must be revalued (although we do take into account the dependencies – or lack thereof – of instruments on the bumped parameters in order to streamline the computation of risk).

5. GOING FORWARD

Plans for a chain of module releases that will appear starting in 2002 will include FAS 133, Value at Risk and Cash flow at Risk, and FX enhancements for the commodities trades we support. Subsequent to that, work is in progress for the release of modules that support power, foreign exchange and fixed income.