

Spot Convenience Yield Models for the Energy Markets

René Carmona and Michael Ludkovski

ABSTRACT. We review that part of the literature on energy spot price models which involves convenience yield as a factor, our goal being to document the shortcomings of the most commonly used models. From a mathematical point of view, the introduction of the convenience yield is usually justified by the desire to reconcile dynamical models for the time evolution of commodity prices with standard arbitrage theory. Since the convenience yield appears as a factor which cannot be observed directly, stochastic filtering has been proposed as a strategy of choice for its estimation from observed market prices. We implement these ideas on the models we review, and on some natural extensions. We illustrate the inconsistencies of the spot models on readily available data, paving the way for the empirical analysis of models of the term structure of convenience yield recently proposed as a viable alternative.

1. Introduction

As the energy markets continue to evolve, valuation of energy-linked financial instruments has been one of the focal topics of recent mathematical finance research. One of the most popular choices for describing asset movements is a class of so-called "convenience yield models". Such models introduce a new unobserved quantity related to physical ownership of the asset. In turn, convenience yield models can be broadly split into "spot" models and "term structure factor" models. To the first group belong the classical Gibson and Schwartz model [GS90], as well as later models of Schwartz [Sch97], Hilliard and Reis [HR98], and Casassus and Collin-Dufresne [CCD03]. The "term structure" models, which have been discussed among others by Miltersen and Schwartz [MS98], and Bjork and Landen [BL01] are similar to Heath-Jarrow-Morton constructions originally introduced in the analysis of fixed income markets.

In this paper we review the literature of spot convenience yield models, and we analyze in detail two new extensions. First, we discuss a variant of the Gibson-Schwartz model with time-dependent parameters. This was first suggested by Miltersen [Mil03], but we provide the first full implementation of the model using empirical data. Second, we describe a new three-factor affine model with stochastic convenience yield and stochastic market price of risk. The existence of a third factor allows us to achieve a good fit to the cross-section of futures prices. The idea of time-dependent risk premia in the context

1991 *Mathematics Subject Classification.* Primary 62P05, 62M20; Secondary 60H10.

Key words and phrases. Spot price, convenience yield, energy markets.

of spot convenience yield models has been recently proposed by Casassus and Collin-Dufresne in [CCD03]. However, they only consider the deterministic case. We believe that a stochastic version is much more natural from a theoretical standpoint, especially with a filtering application in view. Nevertheless, our results indicate that the standard assumption of Ornstein-Uhlenbeck process for convenience yield may be mis-specified. Overall, we fail to find a fully satisfactory model that is both consistent with the spot and the forward curve. In our opinion, this supports the view that the term-structure paradigm is more appropriate for energy commodities.

Because a commodity can be consumed, its price is a combination of future asset and current consumption values. However, unlike financial derivatives, storage of energy products is costly and sometimes practically impossible like in the case of electricity. Consequently, physical ownership of the commodity carries an associated flow of services. On the one hand, the agent has the option of flexibility with regards to consumption (no risk of commodity shortage). On the other hand, the decision to postpone consumption implies storage expenses. The net flow of these services per unit of time is termed the convenience yield δ . From now on we assume that δ is quoted on a continuously compounded basis. Intuitively, the convenience yield corresponds to dividend yield for stocks.

$$\delta = \text{Benefit of direct access} - \text{Cost of carry.}$$

Commodity pricing models are obtained via various assumptions on the behavior of δ_t . The implicit assumption is that S_t the spot price process of the commodity in fact exists. This is not true for some commodities, such as electricity. Even for mature markets like crude oil where spot prices are quoted daily, the exact meaning of the spot is difficult to pin down. Nevertheless, we will maintain the industry-standard assumption of traded spot asset.

By a basic no-arbitrage argument it follows that the price of a forward contract $F(t, T)$ which has payoff S_T at future time T must equal

$$(1.1) \quad F(t, T) = S_t \mathbb{E}_Q \left[e^{\int_t^T (r_s - \delta_s) ds} \right].$$

Indeed, we can replicate the payoff by either entering into the forward contract, or by borrowing S_t today and holding the commodity itself from now until T . As usual, \mathbb{Q} is an equivalent martingale measure used for risk neutral pricing. Reformulating, (1.1) implies that the risk-neutral drift of the spot commodity must equal $(r_t - \delta_t)S_t$. From a modelling point of view, it remains to specify the stochastic processes for S_t , r_t and δ_t . However, the difficulty lies in the fact that the convenience yield is *unobserved*. Notice that δ_t is defined indirectly as a "correction" to the drift of the spot process. Thus, for any model we specify we must address the issue of estimating or filtering the convenience yield process δ_t given model observables.

In this paper we will present several stochastic convenience yield models and discuss the resulting filters. We implement the filtering problems and compare the features of the resulting estimations using crude oil data. Section 2 reviews the original Gibson-Schwartz model and its limitations. In Section 3 we investigate the deterministic shift extension and its relation to work done in fixed income. Our empirical study of this model is the first of its kind. Section 4 is the main thrust of the paper. We analyze the extended model with stochastic risk premia which is a new method of consistently estimating convenience

yield given the entire forward curve. The filtering results from this model are summarized and compared in Section 5. Finally, Section 6 concludes.

2. The Gibson-Schwartz Model

The basic spot model for convenience yield was introduced by Gibson and Schwartz in 1990 [GS90]. This is a 2-factor model for which the risk neutral dynamics of the commodity spot price S_t is given by a geometric Brownian motion whose rate of growth is corrected by a stochastic mean-reverting convenience yield δ_t . To be more specific, we assume that $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ is a filtered probability space, and we consider a bivariate state process comprising the spot commodity asset S_t and the spot instantaneous convenience yield δ_t . According to Gibson and Schwartz, the dynamics of the state are given under the risk-neutral measure \mathbb{Q} by a system of Ito stochastic differential equations of the form:

$$(2.1a) \quad dS_t = (r_t - \delta_t)S_t dt + \sigma S_t dW_t^1,$$

$$(2.1b) \quad d\delta_t = \kappa(\theta - \delta_t)dt + \gamma dW_t^2,$$

where W^1 , and W^2 are 1-dimensional Wiener processes satisfying $d\langle W^1, W^2 \rangle_t = \rho dt$. We assume throughout that $\{\mathcal{F}_t\}_t$ is the filtration generated by this bivariate Wiener process. Depending on market conditions, convenience yields can be either positive or negative, and so unlike with interest rates, the choice of an Ornstein-Uhlenbeck process for δ_t in (2.1b) makes sense. From now on we shall also assume that

ASSUMPTION 1. The interest rate r_t is deterministic.

In practice, one observes that the volatility of the convenience yield is an order of magnitude higher than the volatility of interest rates. Consequently, letting r_t be stochastic as in Schwartz [Sch97] does not give much qualitative improvement to the model.

The implications of (A1) are crucial. Indeed, non-stochastic interest rates imply that futures and forward prices are the same, and both equal the risk-neutral expectation of the future spot price.

If $\rho > 0$, then the stochastic convenience yield induces weak mean-reversion in the risk neutral dynamics of the spot price. Observe that according to (2.1), when S_t is increasing due to increment dW^1 , thanks to positive correlation ρ , δ_t is also likely to increase. In turn, this reduces the drift of the spot. Note that this is a second-order effect. At the risk of adding to the confusion, one should add that mean reversion of the historical dynamics of the spot is commonly accepted as an empirical fact, see for example [GS90], and values of $\rho \sim 0.3 - 0.7$ are viewed as reasonable. The theory of storage developed back in the 1950s [Bre58] shows that the endogeneous economic link is through inventory levels: when inventories are low, shortages are likely, causing high prices, as well as valuable optionality of holding the physical asset.

It is convenient to introduce a special notation for the logarithm of the spot price $X_t = \log S_t$. Indeed, the model (2.1) is linear in the state vector $Z_t = [X_t, \delta_t]$.

2.1. Model Estimation. Statistical estimation of a model can only be done from observations, consequently, and it cannot be done directly from the prescriptions of a risk neutral model as given by (2.1). We need to work backward and recover the dynamics of the state variables under the objective or historical probability structure. This requires us

to make assumptions on the market prices of risk, say λ for S_t and λ_δ for the unobserved convenience yield (each random source must have its own market price of risk). The simplest possible assumption is to assume that:

ASSUMPTION 2. λ and λ_δ are constant.

Under this assumption, the historical dynamics of the two factors under the is exactly of the same form as for the risk neutral measure, and estimation will be plain. Indeed, the model is still linear (at least for the logarithm of the spot) and hence conditionally Gaussian, and we can estimate all the parameters empirically using the standard Kalman filter method. The historical dynamics of δ_t are of the form:

$$(2.2) \quad d\delta_t = [\kappa(\theta - \delta_t) - \lambda_\delta]dt + \gamma d\widetilde{W}_t^2$$

for some bivariate Brownian motion under \mathbb{P} , say $(\widetilde{W}^1, \widetilde{W}^2)$, having the same correlation as (W^1, W^2) under \mathbb{Q} . So the dynamical equation is the same as before as long as we replace θ by $\hat{\theta} = \theta - \lambda_\delta/\kappa$. Consequently, after a straightforward discretization of the time, our measurement and transition equations become:

$$(2.3a) \quad X_n = X_{n-1} + \left(r_n - \delta_n - \frac{\sigma^2}{2} + \lambda \right) \Delta t + \xi_n, \quad \xi_n \sim \mathcal{N}(0, \sigma_S^2 \Delta t),$$

$$(2.3b) \quad \delta_n = e^{-\kappa \Delta t} \delta_{n-1} + (1 - e^{-\kappa \Delta t}) \hat{\theta} + \eta_n, \quad \eta_n \sim \mathcal{N}\left(0, \frac{\gamma^2}{2\kappa} (1 - e^{-2\kappa \Delta t})\right),$$

and they are clearly amenable to a classical Kalman filter analysis. Figure 1 shows the result of such Kalman filtering on crude oil data (see Section 5.1 for description of our data set). We estimate the model parameters $(\kappa, \theta, \gamma, \rho, \lambda)$ using the prediction error decomposition of the likelihood function coming from the filter [Har89]. As is usual in such models, the estimates of κ and λ are very imprecise. Overall, we find that $\kappa \in [0.1, 0.4]$, $\gamma \in [0.4, 0.5]$, $\theta \in [-0.3, -0.1]$, $\rho \in [0.5, 0.7]$, $\lambda \in [-0.1, 0.1]$. The spot volatility σ is estimated by the historical volatility of the time series, which is about 45% for our data set. For interest rates we use the current 3-month LIBOR rate.

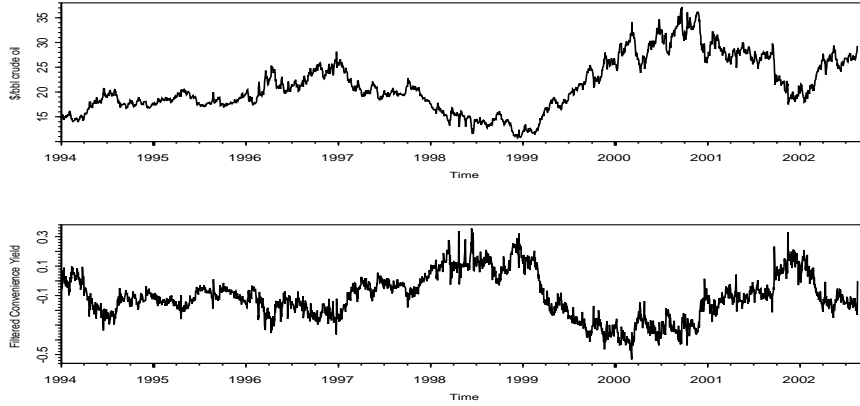


FIGURE 1. Filtered convenience yield for crude oil, 1994-2002. $\kappa = 0.2, \gamma = 0.5, \theta = -0.15, \rho = 0.7, \lambda = 0$.

2.2. Lack of Consistency. Unfortunately, the basic model (2.1) is not consistent with the *forward* curve. In commodity markets, on any given day t , the forward (sic) curve defines the term structure of futures contracts. It is defined as the graph of the function $T \mapsto F(t, T)$. In the particular case of the risk neutral model (2.1), the convenience yield is (at least conditionally) a Gaussian process, and the conditional expectation giving the value of $F(t, T)$ in (1.1) can be computed explicitly. A direct calculation gives:

$$(2.4) \quad F(t, T) = F(t, T; Z_t) = S_t e^{\int_t^T r_s ds} e^{B(t, T)\delta_t + A(t, T)} \quad \text{where}$$

$$(2.5) \quad B(t, T) = \frac{e^{-\kappa T} - 1}{\kappa},$$

$$(2.6) \quad A(t, T) = \frac{\kappa\theta + \rho\sigma_s\gamma}{\kappa^2} (1 - e^{-\kappa(T-t)} - \kappa(T-t)) + \frac{\gamma^2}{\kappa^3} (2\kappa(T-t) - 3 + 4e^{-\kappa(T-t)} - e^{-2\kappa(T-t)}).$$

Similar formulae hold in the slightly more general setting of affine processes. The fact that these exponential affine models admit analytical futures and options prices in closed form is emphasized by Bjork and Landen in [BL01]. In any case, one sees that under \mathbb{Q} , the forward contract follows the dynamics given by the stochastic differential equation:

$$(2.7) \quad dF(t, T) = F(t, T) \left[r_t dt + \sigma dW_t^1 + \gamma \frac{e^{-\kappa T} - 1}{\kappa} dW_t^2 \right].$$

Also, recall that the convenience yield follows an Ornstein-Uhlenbeck (OU) process. Conditional on \mathcal{F}_s , δ_t is Gaussian with

$$\delta_t | \mathcal{F}_s \sim \mathcal{N} \left((1 - e^{-\kappa(t-s)})\delta_s + e^{-\kappa(t-s)}\theta, \frac{\gamma^2}{2\kappa} (1 - e^{-2\kappa(t-s)}) \right).$$

Using the values of the estimated parameters in (2.2) together with the price of a traded futures contract, say $F(t, T_i)$, we can solve for the convenience yield δ_t . The values of δ_t

extracted from the observed futures prices $F(t, T_i)$ are said to be *implied* by the futures contract with maturity T_i . Figure 2 gives the plot of the value of the convenience yield implied by a three month futures oil contract. Comparing this plot with the results in Figure 1, it is clear that the predictive power of the Gibson-Schwartz model is very limited in practice. Furthermore, each futures contract seems to carry its own source of risk as evidenced by the sharp spikes in Figure 2 which result due to sudden uncorrelated movements in the futures and the spot. As Figure 3 shows, there is a further inconsistency between the implied δ_t 's from futures contracts of different maturities. In the next sections we will attempt to resolve this problem in two ways – either by considering time-dependent parameters or by taking into account the forward curve and enlarging the state space.




Figure 2: Implied convenience yield using a 3-month futures and same parameters as in Figure 1.




Figure 3: Difference in implied convenience yields between 3- and 12-month futures.

3. Miltersen Extension

Miltersen [Mil03] proposed to extend (2.1) by allowing some of the parameters to be a function of time. This is a direct analogue of the extension of the Vasicek model for interest rates done by Hull and White [HW93].

The simplest choice is to make the mean-reversion level θ time-dependent, $\theta = \vartheta(t)$. This allows to fit the initial futures prices directly, in the same way that the Hull-White model can fit the initial term structure of bond prices [BM01]. The calibration is performed by letting the spot yield be "dragged around" its changing mean $\vartheta(t)$.

Integrating the equation for δ_t in (2.1) we obtain

$$(3.1) \quad \delta_t = \delta_s e^{-\kappa(t-s)} + \int_s^t e^{-\kappa(t-u)} \vartheta(u) du + \gamma \int_s^t e^{-\kappa(t-u)} dW_u^2.$$

If we define

$$(3.2) \quad \alpha_t := \int_s^t e^{-\kappa(t-u)} \vartheta(u) du,$$

it follows that we have the deterministic shift decomposition $\delta_t = a_t + \alpha_t$, where a_t follows the mean-zero OU process

$$da_t = -\kappa a_t dt + \gamma dW_t^2.$$

To complete the calibration we must take $\vartheta(t)$ to match a chosen set of observed futures prices $F(0, T_i), i = 1, 2, \dots, n$. For this purpose define $\epsilon(0, t)$ via

$$(3.3) \quad F(0, T_i) = S_0 e^{\int_0^{T_i} (r_s - \epsilon(0, s)) ds}.$$

In HJM-type models [MS98], $\epsilon(0, t)$ is called the initial term structure of futures convenience yields. Then it can be shown that

$$(3.4) \quad \vartheta(t) = \frac{\epsilon_T(0, t)}{\kappa} + \epsilon(0, t) + \frac{\gamma^2}{2\kappa^2}(1 - e^{-2\kappa t}) - \frac{\rho\sigma\gamma}{\kappa}$$

or alternatively $\alpha_t = \epsilon_T(0, t) + \frac{\gamma^2}{\kappa^2}(1 - e^{-\kappa t})^2$.

Solving for $\epsilon(0, t)$ in (3.3) we obtain

$$(3.5) \quad \epsilon(0, t) = r_t - \frac{\partial \log(F(0, t))}{\partial t}.$$

Of course our data consists of just $\{F(0, T_i)\}$ so we must interpolate those smoothly and then take partial derivatives to infer the implied $\epsilon(0, t)$'s. We recommend fitting the forward curve with B-splines using 5-8 degrees of freedom. B-splines provide accurate fit to market data and also generate smooth first and second derivatives w.r.t. T (note that for (3.4) we also need $\frac{\partial \epsilon(t, T)}{\partial T}$). Figures 4 and 6 provide examples of such B-spline fitting with six degrees of freedom. As our input we use various 28-month crude oil forward curves from the last few years. The parameter values are the same as in Figure 1.




Figure 4: Crude oil forward curve and the interpolated term structure of convenience yields. Backwardation of 11/21/2000.




Figure 5: Term structure of mean reversion level ϑ_t , 11/21/2000.




Figure 6: Crude oil forward curve and the interpolated term structure of convenience yields. Contango observed on 1/18/2002.




Figure 7: Term structure of mean reversion level ϑ_t , 1/18/2002.

From a more general calibration point of view, we now fit exactly a set of benchmark futures prices and instead concentrate on modelling the term structure of volatility. The empirical work in this direction is complicated by the need to estimate volatility and the numerical difficulty of estimating the rest of the parameters given time-dependent models. Also, it is rather unnatural to fit the initial term-structure, while remaining in a spot paradigm since this leaves the model exposed to inconsistencies in time (i.e. the need to constant re-calibration).

4. Enlarging the Observation Equation

The inconsistency alluded to in the end of Section 2 is serious. We have an explicit formula (2.2) for the value of the futures price as a function of the spot and the convenience yield, and this formula is invertible. Hence, at least in principle, we can recover the exact value of δ_t from futures prices. Unfortunately, there are many futures traded, and if we compare the implied δ_t from a 3-month futures it does not agree with the implied δ_t implied from a 6-month contract (cf. Figure 3). The obvious solution to this dilemma is to assign to each contract, its own source of idiosyncratic risk due to temporary miss-pricing, bid-ask spread and liquidity concerns. Such an assumption is made by Schwartz [Sch97]. Considering (2.2), we see that fixing S_t , the entire forward curve is just 1-dimensional, which is certainly not what one observes in practice. A more elaborate idea is to assume a whole term structure of convenience yields, one for each maturity. This leads to HJM-type constructions [MS98].

The natural solution to model consistency is to have the futures (or any other vanilla derivatives that are liquidly traded) as part of our observation. After including the futures $F_s^i = F(s, T_i)$, $i = 1, 2, \dots, n$, the observable filtration is now $\mathcal{F}_t^O = \sigma\{S_s, F_s^i : 0 \leq s \leq t\}$.

By enlarging the set of observable instruments we are also able to address other limitations of the model. Recall that the behavior of the market price of risk λ for the spot was so far "swept under the rug". By (A2) λ is constant. Ignoring for a second the validity of such assumption, we are still faced with the problem of estimating its value. While the market price of risk can be in principle inferred from prices of traded assets, this method is cumbersome and imprecise. An alternative suggestion elaborated by Runggaldier et al. in a series of papers [BCR01, GJR02, GR01, Run03] is to

make λ_t also a stochastic process. Because a filter is robust to model specification, this approach is valid even if λ_t is deterministic or just constant.

Runggaldier [Run03] suggests using another OU process for λ_t since mean reversion and the resulting stationarity is a desirable feature:

$$(4.1) \quad d\lambda_t = \kappa_\lambda(\bar{\lambda} - \lambda_t)dt + \sigma_\lambda dW_t^3.$$

It is reasonable to assume that the market price of risk carries its own Brownian motion and that W^3 is correlated with W^1 , but not with W^2 the noise of the convenience yield. Intuitively, the market risk premia are independent from storage costs of the commodity. We also expect that there is negative correlation between the spot and the market price of risk $\rho_{S\lambda} < 0$. This is to strengthen the empirical mean-reversion in the spot [CCD03].

Working in an incomplete market we pick the minimal martingale measure \mathbb{Q} so that the Girsanov transformation corresponding to λ_t affects only W^1 (\widetilde{W}^1 under \mathbb{P}):

$$dW_t^1 = d\widetilde{W}_t^1 - \lambda_t dt.$$

Our extended state is now $Z_t = [X_t, \delta_t, \lambda_t]$ which is again Markovian, and in fact conditionally Gaussian. Thus we preserve the exponential affine model which allows for easy pricing of futures and options.

Following the standard martingale method for pricing derivatives, we assume that the price of a futures contract is a function of the state vector

$$(4.2) \quad F_t^i = F^i(t, Z_t) = \mathbb{E}_{\mathbb{Q}}[S_{T_i} | \mathcal{F}_t].$$

Then applying Itô's formula and using the fact that discounted traded asset prices are \mathbb{Q} -martingales we must have

$$(4.3) \quad dF_t^i = r_t F_t^i dt + \sigma S_t \frac{\partial F_t^i}{\partial S} dW_t^1 + \gamma \frac{\partial F_t^i}{\partial \delta} dW_t^2 + \sigma_\lambda \frac{\partial F_t^i}{\partial \lambda} dW_t^3.$$

In our case we already have the explicit expression (2.2) which held for the Gibson-Schwartz model. This was derived by a replication argument. Since λ only affects the distribution under \mathbb{P} , the replication argument still goes through in the extended model. In other words, $\frac{\partial F_t^i}{\partial \lambda} = 0$ and (4.3) simplifies to

$$(4.4) \quad \frac{dF_t}{F_t} = r_t dt + \left(\sigma + \rho\gamma \frac{e^{-\kappa T} - 1}{\kappa} \right) \lambda_t dt + \sigma d\widetilde{W}_t^1 + \gamma \frac{e^{-\kappa T} - 1}{\kappa} d\widetilde{W}_t^2 + \alpha dW_t^F.$$

The last term αdW_t^F is the idiosyncratic risk associated with F_t and used to smooth out the data. We expect α to be an order smaller than the other volatilities in (4.4). Summarizing, the complete filtering model under the real world probability \mathbb{P} is given by

(4.5a)

$$d\lambda_t = \kappa_\lambda(\bar{\lambda} - \lambda_t) dt + \sigma_\lambda dW_t^3$$

(4.5b)

$$d\delta_t = \kappa(\hat{\theta} - \delta_t) dt + \gamma d\tilde{W}_t^2$$

(4.5c)

$$dS_t = (r_t - \delta_t + \sigma\lambda_t)S_t dt + \sigma S_t d\tilde{W}_t^1$$

(4.5d)

$$dF_t^i = \left(r_t + \sigma\lambda_t + \rho\gamma \frac{e^{-\kappa T_i} - 1}{\kappa} \lambda_t \right) F_t^i dt + \sigma F_t^i d\tilde{W}_t^1 + \gamma F_t^i \frac{e^{-\kappa T_i} - 1}{\kappa} d\tilde{W}_t^2 + \alpha dW_t^{F^i}.$$

According to the setup, δ_t and λ_t are not observable, S_t is fully observed and F_t^i is imperfectly observed in the market.

Notice that after taking logarithms of observed prices, the entire system is still linear and hence amenable to Kalman filtering. We note that this linearity is more an artifact of the model rather than its goal. Our choice of OU processes for λ and δ has been motivated by heuristic arguments, not by modelling convenience.

Figure 8 shows the result of applying (4.5) to crude oil data. Once again we estimate the parameters $(\kappa, \hat{\theta}, \gamma, \rho, \kappa_\lambda, \bar{\lambda}, \sigma_\lambda, \rho_\lambda)$ using the prediction error decomposition from the filter. We find that the filter is relatively insensitive to the parameters of the λ process. The range of λ appears to be much smaller (on the order of 0.05) and that the correlation between λ and the spot is essentially zero. The presence of λ also reduces the estimated correlation between spot and δ_t to around 0.5. The empirical correlation between λ_t and δ_t is 0.258. The idiosyncratic futures noise α corresponds to a small error of about 0.05. Empirically, this is sufficient to remove the undesirable spikes of Figure 2. However, the empirical results are unsatisfactory. Even with the noise term $\alpha dW_t^{F^i}$, the filtered convenience yield exhibits severe spikes which contradict the initial assumption of δ_t following an OU process. Furthermore, in Figure 8, δ_t does not exhibit any sort of persistence, contrary to economic intuition. In short, while the theoretical model is superior, the empirical results fall short of the simpler Gibson-Schwartz case.

$\kappa = 0.2$	$\kappa_\lambda = 0.4$
$\gamma = 0.5$	$\sigma_\lambda = 0.35$
$\hat{\theta} = -0.1$	$\bar{\lambda} = 0.02$
$\rho = 0.5$	$\rho_\lambda = -0.1$

TABLE 1. Estimated parameter values for Model (4.5).

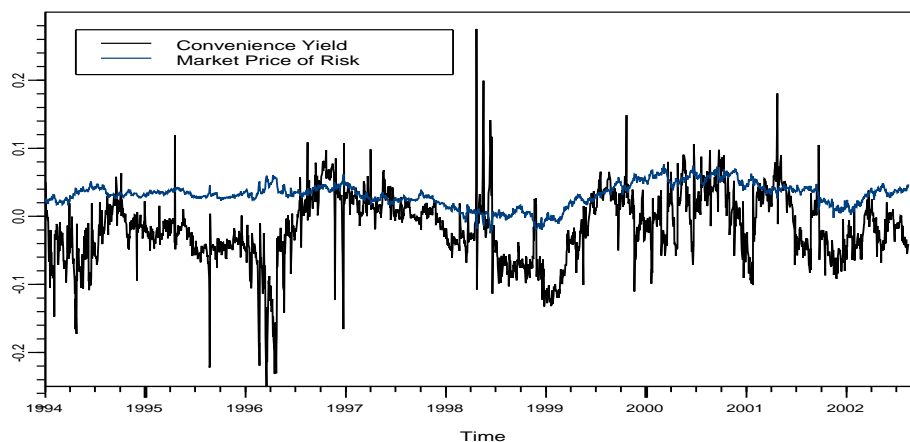


FIGURE 8. Filtered convenience yield for crude oil using the spot and 3-month futures. Parameter values are given in Table 1.

5. Empirical Implementation

5.1. Description of Oil Data. For our data set we choose the West Texas Intermediate oil contracts during the period Jan 1994 - July 2002. Crude oil is one of the most mature energy markets with a well-developed forward curve spanning several years.

The WTI futures contract trades on NYMEX since 1993. Monthly futures contracts are traded with maturities of 1, 2, 3, ..., 23, 24, 30, 36, 48, and 60 months. The longer maturities usually trade as a set of strips (for example a January-June strip of 6 months) and consequently exhibit a high correlation. A single futures contract is for 1,000bbl to be delivered anytime during the delivery month at Cushing, Oklahoma. Trading terminates on the third business day prior to the 25th calendar day of the month preceding the delivery month. The spot price is defined as the Balance-of-the-Month (BOM) contract and has varying liquidity.

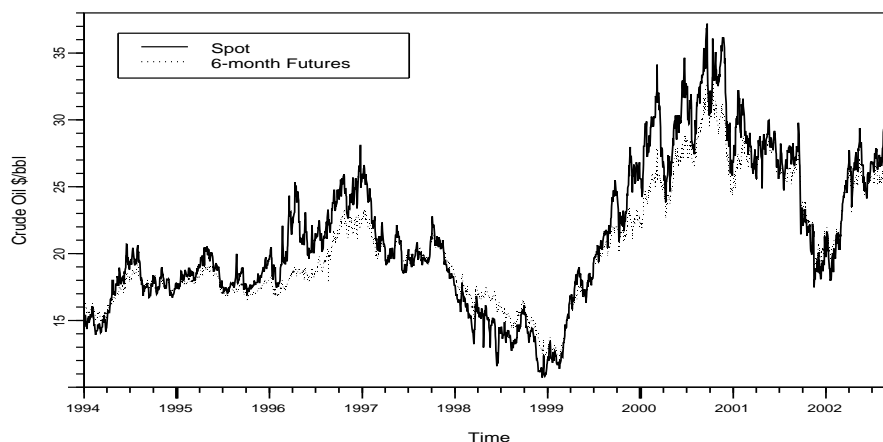


FIGURE 9. Time series for crude oil spot and 6-month futures, 1994-2002.

As opposed to many other energy assets, crude oil exhibits little seasonality and the prices tend to be somewhat stable. Consequently the value of having the physical commodity at hand is small and the net convenience yield is dominated by storage costs. As a result, the forward curve is usually in *backwardation*, i.e. $T \rightarrow F(t, T)$ is a decreasing function. Occasionally we have a flip and the forward curve appears as a *contango*. Litzenberger and Rabinowitz [LR95] report that 80%-90% of the time the oil forward curve is in backwardation. Historically, periods of contango when futures prices are above the spot price are highly correlated with periods of low prices.

Oil futures also exhibit the well-known Samuelson effect. When the contracts are first started and are far away from maturity (2+ years), they are thinly traded and exhibit low volatility. As the maturity nears both trading volume and volatility increase. As a rule of thumb, the nearest month contract is both the most liquid and the most volatile. The spot contract is rather distinct, since it is used for different purposes. Specifically, spot contracts are usually used for balancing day-to-day needs and consequently exhibit high volatility and medium volume. Thus, the term structure of oil futures volatility $T \rightarrow \sigma(t, T)$ is usually decreasing.

5.2. Results. First, we compare the consistency of the three-factor model (4.5) with respect to the forward curve. We find that our estimate of convenience yield is stable when using futures contracts with different maturities as inputs. Figure 10 shows the estimated δ_t using three different sets of two futures contracts as inputs. All three estimates are quite close to each other. Thus our model succeeds in removing the inconsistency of Figure 3.

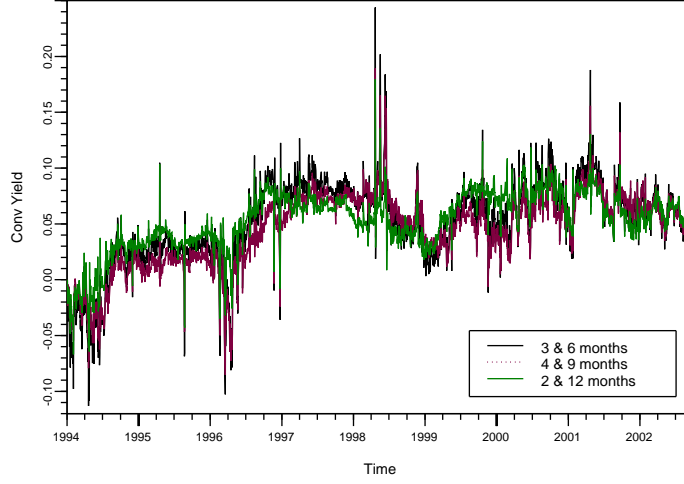


FIGURE 10. Estimates of δ_t using three different pairs of futures contracts.

Secondly, we test the notion that the convenience yield is the flow of services accruing to the holder of the physical asset but not to the holder of a futures contract. This interpretation implies that when the convenience yield is positive $\delta_t > 0$, the forward curve should be in backwardation, and conversely when the convenience yield is negative we should be in contango. Thus, if our model fits well, the estimated δ_t can be used as an indicator to predict the switches from contango and backwardation. Such market transitions are very important for actual trading. In this vein we compare the performance of our δ_t -indicator versus two other indicators proposed recently by Borovkova [Bor03b, Bor03a].

Borovkova proposes to use the moving-average inter-month spread indicator:

$$(5.1) \quad I_1(t) = \frac{1}{M} \sum_{k=0}^M \sum_{i=1}^n \gamma^i (F_{i+1}(t - k\Delta t) - F_i(t - k\Delta t))$$

and the PCA indicator:

$$(5.2) \quad I_2(t) = \frac{1}{M} \sum_{k=0}^M \sum_i w_i \bar{F}_i(t - k\Delta t).$$

Above γ is the discount factor used in weighing different inter-month spreads. The implicit assumption is that the closer contracts are more important and so the weights can be written as $w_k = \gamma^k$. Similarly, w_i are the factor loads corresponding to the first factor from PCA performed on a de-meaned forward curve. This factor generally corresponds to "slope". In [Bor03b] $I_2(t)$ is anticipative because the PCA is performed on the entire data sample.

Figure 5.2 compares the performance of these two indicators, as well as the negative filtered convenience yield $-\delta_t$, on the oil data from 1994-2002. The areas of contango/backwardation were identified a priori by looking directly at the forward curves. As can be seen, all three indicators perform well. The times when the indicators are near zero generally correspond to periods of market transition or uncertainty. Of the three the convenience yield one usually gives the clearest signal, but also gave a false backwardation signal in early 1995.

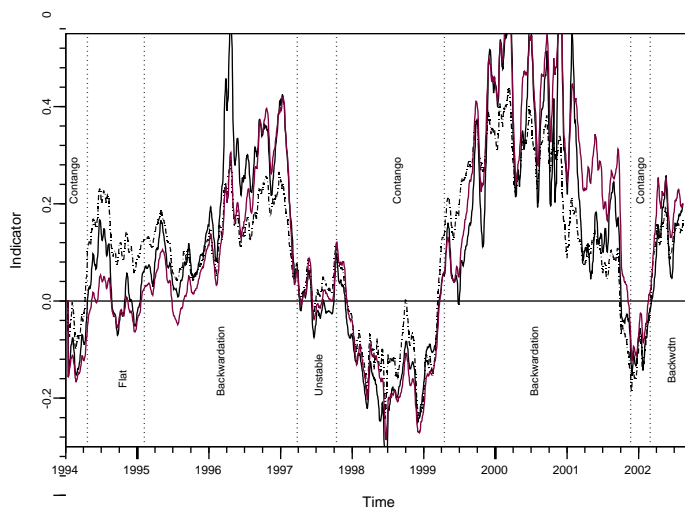


FIGURE 11. Three different indicators for forward curve state: I_1 , I_2 (gray) and $I_{GS} \delta_t$ (dotted line) using a 3-month futures from Gibson-Schwartz model. Parameter values are the same as in Figure 1.

We find that the forward curve indicator from the extended stochastic- λ model is not a good predictor for market transitions. From Figures 8 and 10 we can see the estimated convenience yield is too close to zero to give a clear signal on the state of the forward curve.

6. Conclusion

In this paper we investigated filtering of convenience yield models. We discussed the inadequacies of the classical Gibson-Schwartz model and analyzed two new extensions. The time-inhomogeneous extension first proposed by Miltersen seems to work well but requires further investigation into the forward curve volatility structure to fully judge its consistency. This is fraught with complications because of discrepancies between realized and implied volatilities. The stochastic market price of risk extension was also studied. This approach provides a "clean" solution while maintaining the exponential affine feature of the model crucial for fast implementation. Furthermore, the model is consistent with the observed forward curve. However, the filtered convenience yield displays extreme spikeness which contradicts the assumption of a driving OU process. As well, the model is not as good as the basic one when it comes to predicting forward curve state

transitions. This fact is unsettling because it goes against the economic underpinnings of the convenience yield definition.

Our work provided some clues for future research and exposed the weaknesses of various spot models. Overall, we believe that this study shows the need for more sophisticated term-structure models in order to explain both the spot and the forward curve. In turn, this requires a more careful analysis of the term structure of futures volatility. Further work also needs to be carried on applying the models to natural gas markets where seasonality is strong. Decoupling of seasonality and unobserved state variables remains an open theoretical problem.

References

- [BCR01] R. Bhar, C. Chiarella, and W.J. Runggaldier, *Filtering equity risk premia from derivative prices*, Tech. report, UTS, Sydney, 2001, QFRG Research Paper 69.
- [BL01] T. Bjork and C. Landen, *On the term structure of futures and forward prices*, Tech. report, Stockholm School of Economics, 2001.
- [BM01] D. Brigo and F. Mercurio, *Interest rate models, theory and practice*, Springer Finance, Heidelberg, 2001.
- [Bor03a] S. Borovkova, *Backwardation and contango change indicators for seasonal commodities*, Energy Power Risk Management **VIII (4)** (2003).
- [Bor03b] ———, *Detecting market transitions: from backwardation to contango and back*, Energy Power Risk Management **VIII (3)** (2003).
- [Bre58] M.J. Brennan, *The supply of storage*, American Economic Review **48** (1958), 50–72.
- [CCD03] J. Casassus and P. Collin-Dufresne, *Maximal affine model of convenience yields implied from interest rates and commodity futures*, Tech. report, Carnegie Mellon University, 2003, Preprint.
- [GJR02] A. Gombani, S.R. Jaschke, and W.J. Runggaldier, *A filtered no arbitrage model for term structures from noisy data*, Preprint, 2002.
- [GR01] A. Gombani and W.J. Runggaldier, *A filtering approach to pricing in multifactor term structure models*, International Journal of Theoretical and Applied Finance **4** (2001), 303–320.
- [GS90] R. Gibson and E. S. Schwartz, *Stochastic convenience yield and the pricing of oil contingent claims*, Journal of Finance **XLV(3)** (1990), 959–976.
- [Har89] A.C. Harvey, *Forecasting structural time series models and the kalman filter*, Cambridge University Press, Cambridge, UK, 1989.
- [HR98] J.E. Hilliard and J. Reis, *Valuation of commodity futures and options under stochastic convenience yields, interest rates, and jump diffusions in the spot*, Journal of Financial and Quantitative Analysis **33(1)** (1998), 61–86.
- [HW93] J. Hull and A. White, *One-factor interest-rate models and the valuation of interest-rate derivative securities*, Journal of Financial and Quantitative Analysis **28(2)** (1993), 235–254.
- [LR95] R. Litzenberger and N. Rabinowitz, *Backwardation in oil futures markets: Theory and empirical evidence*, Journal of Finance **L(5)** (1995), 1517–1545.
- [Mil03] K. Miltersen, *Commodity price modelling that matches current observables: A new approach*, Quantitative Finance **3(1)** (2003), 51–58.
- [MS98] K. Miltersen and E. Schwartz, *Pricing of options on commodity futures with stochastic term structures of convenience yields and interest rates*, Journal of Financial and Quantitative Analysis **33(1)** (1998), 33–59.
- [Run03] W.J. Runggaldier, *Estimation via stochastic filtering in financial market models*, Preprint, 2003.
- [Sch97] E. Schwartz, *The stochastic behavior of commodity prices: implications for valuation and hedging*, Journal of Finance **LII(3)** (1997), 922–973.

BENDHEIM CENTER FOR FINANCE, DEPARTMENT OF OPERATIONS RESEARCH AND FINANCIAL ENGINEERING, PRINCETON UNIVERSITY, PRINCETON, NJ 08544

E-mail address: rcarmona@princeton.edu

BENDHEIM CENTER FOR FINANCE, DEPARTMENT OF OPERATIONS RESEARCH AND FINANCIAL ENGINEERING, PRINCETON UNIVERSITY, PRINCETON, NJ 08544

E-mail address: mludkovski@princeton.edu