

1200, McGill College Av.
Suite 1000
Montreal, Quebec
Canada H3B 4G7

T +1 514.861.8828
F +1 514.861.4863

Prepared by Numera's
Quant Group

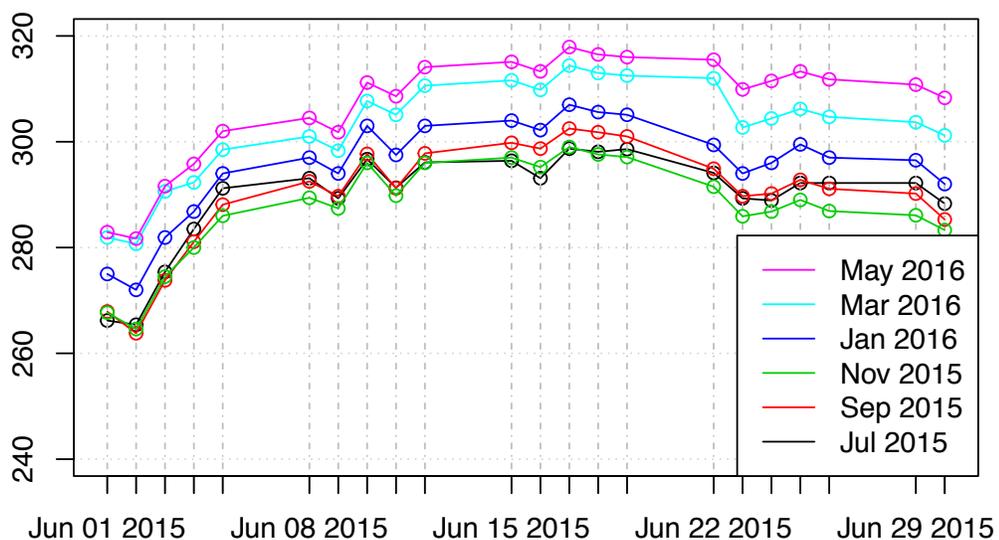
CME Lumber Futures Market: Price Discovery and Forecasting Power

In this report we explore the market for the random lengths lumber futures contract which trades on the Chicago Mercantile Exchange, as well as its role as a mechanism for price discovery and its power to forecast future prices. Using standard econometric techniques, we construct a monthly series of spot prices from the futures prices and we show that the futures price is not a more accurate predictor of future spot prices than even the current spot price.

The Market and the Futures Curve

Futures contracts for so-called *random lengths* lumber have traded on the Chicago Mercantile Exchange since 1969. They involve locking in a delivery price (in USD per thousand board feet) ahead of time for a fixed maturity and a standardized lot of a nominal 110 000 board feet (see the official CME Rulebook for specifications). They are settled at expiry through physical delivery of the lumber and can be used as a hedging tool for both producers and consumers. It is also possible to avoid physical delivery by closing a position before expiry, which allows market participants to expose themselves to lumber market risk without directly participating in the physical lumber business.

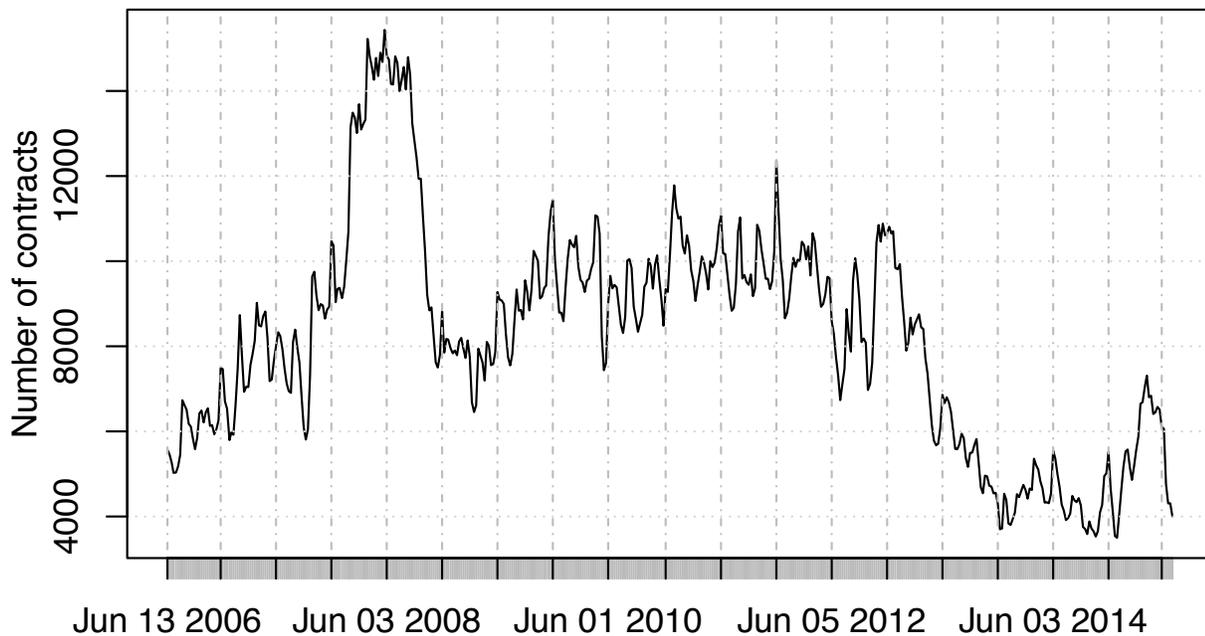
Recent Lumber Futures Prices by Contract



Source: CME Group

The lumber futures market is quite small. Maturities range up to about a year in two-month increments, as contracts expire on the last trading day before the 16th of January, March, May, July, September and November. At its peak, the current front-month contract (LBN2015, expiring in July 2015) represented an open interest of about 5000 contracts. If we compare to crude oil (one of the largest commodities futures markets), the equivalent contract (CLN2015) had interest almost 100 times larger. There are also many more crude oil contracts (every month in the short term), as well as longer-dated contracts until December 2022.

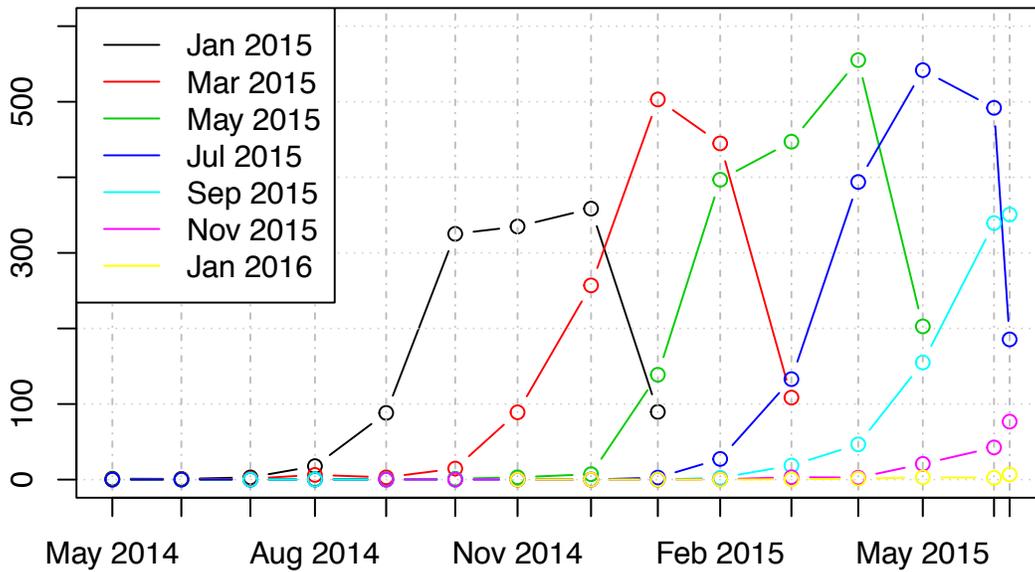
Open Interest For Lumber Futures



Source: CFTC

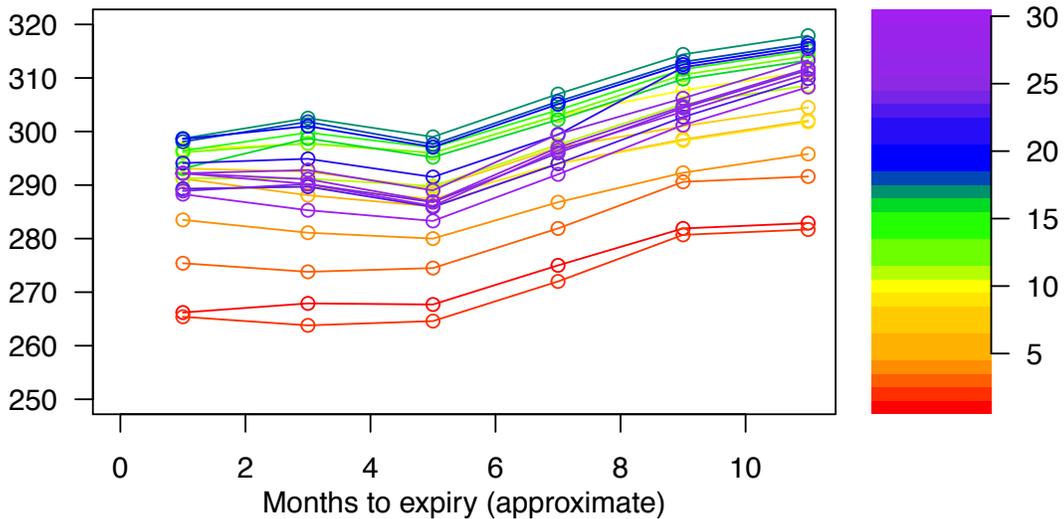
The total open interest (across all maturities) is currently sitting at roughly 4000 contracts for lumber, and 1.7 million for crude oil. The corresponding option market represents an open interest of an additional 3 million contracts for crude oil, and is thoroughly unpopular for lumber with only 700 contracts. This may be because of less interest in complex speculative positions in lumber markets.

Average Daily Volume by Contract



The majority of the interest rests on the shortest contracts. The trading volume typically peaks between one to three months before expiry and, for more than four months, there is almost no volume. This means that only the first two contracts are actually actively traded. The volume in the last month is also quite small, as most of the activity has already shifted to the next contract.

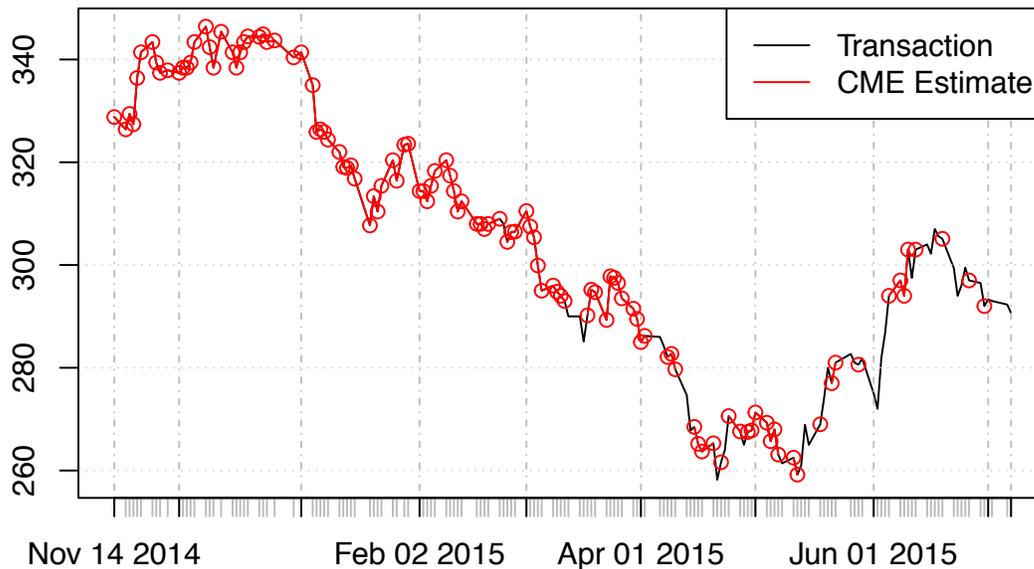
Daily Lumber Futures Curve (June 2015)



Source: CME Group

On futures markets, participants often look for information embedded in the *futures curve* (and in particular, in its slope), which is the plot of futures prices across maturities at a fixed point in time (each curve above represents prices at the end of a given day, in spectral order from red to violet). It describes the term structure of the spot price, similar to the yield curve for fixed income markets. In the case of lumber, it only has six points at any given time, and only the first two are really useful. This is because the prices for the further maturities are "stale": CME documentation describes a procedure for determining an estimated settlement price when no transaction takes place in a given trading day, and this is what is shown in the long part of the curve on *most days*. For example, there are settlement prices listed for the January 2016 contract starting on November 14th, 2014, but it did not see any volume at all until February 24th, 2015, and still does not trade regularly now.

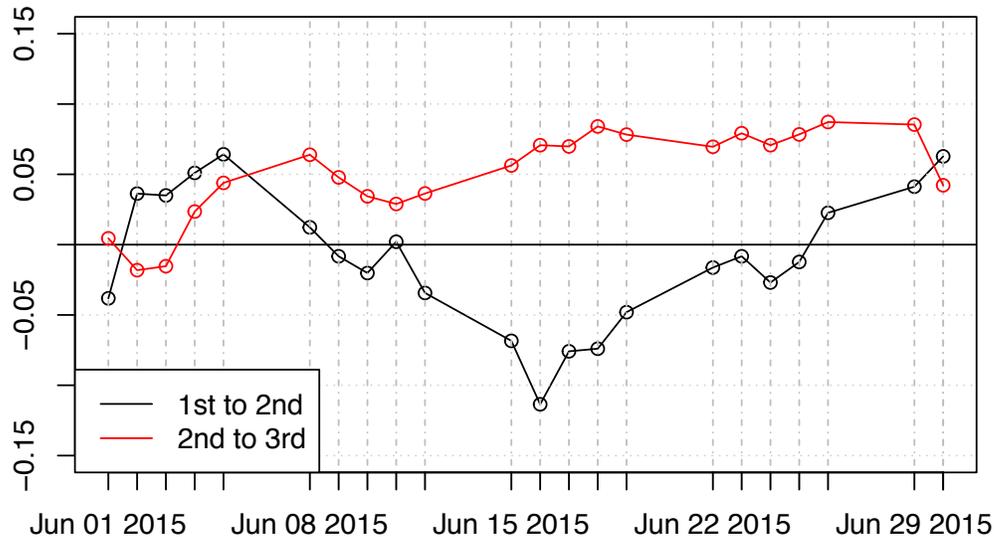
Daily Settlement Price For January 2016 Contract



Source: CME Group

The estimated settlement procedure involves shifting the price up or down by the same amount as the nearest shorter-term traded contract (*a parallel shift*). This is a simple rule but it is unlikely to give a good estimate because these curves do not tend to shift in parallel, empirically. For example, when spot prices are increasing, usually long-term futures prices tend to increase less, and so the slope becomes less positive (or even negative). You can actually see this in the futures curve plot above: on the upswing (up to around the 10th, red to yellow on the graph), the slope between the first two contracts is negative, whereas on the downswing (green to violet), it is positive. This type of behavior is seen as evidence of mean-reversion: part of any shock to the spot price is expected to be reverted over time, so that long-term contracts should always move less than short-term contracts. Parallel shifts would correspond to perfect correlation across maturities.

Futures Curve Slope Between Contracts



Source: CME Group

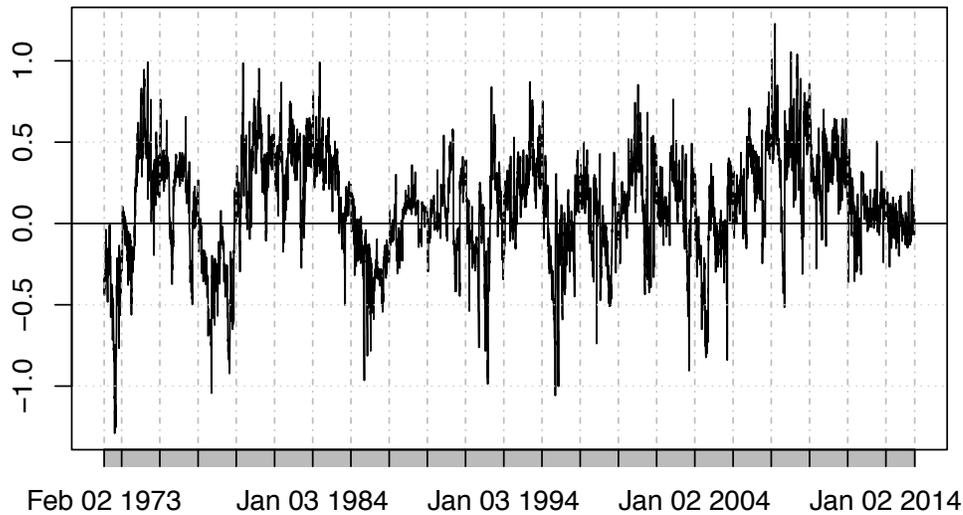
For commodities in general, the slope of the futures curve is increasing with interest rates and storage costs, and decreasing with the *convenience yield*. The convenience yield is an abstraction that is meant to convey the benefits accumulated from having a commodity on hand as opposed to having to go to the market to buy it later (in terms of avoiding supply shortages, for example). It is intuitively like a *commodity-denominated* interest rate that we would charge if we loaned out the commodity, and formally equivalent to the dividend yield for equities, or the foreign interest rate on currency markets. Its influence on the futures curve is related to the fact that the holder of the futures contract does *not* accumulate the convenience yield that an alternative buy-and-hold strategy would produce. He also does not pay storage costs and can defer capital outlays (which may earn interest in the meantime) until expiry.

Because of this relationship, the slope of the futures curve tends to be negative when interest rates and storage costs are very low and when the convenience yield is high. This is a common definition of backwardation, and the opposite situation is called *contango*. Note that there are at least two other related but subtly different definitions of these terms that are in common use: the one used in formal economic theory (involving *expectations* of future spot prices), and the one used informally in the context of the very profitable 2009 crude oil carry trade (involving risk-neutral expectations and *arbitrage*). The definition we use is purely descriptive, whereas the other two involve assumptions about future behavior and the detection of mispriced futures (which would need to be modelled).

For lumber, the slope between the first two contracts is sometimes positive and sometimes negative, but usually positive in the longer term (although, as mentioned, the long-term slope is less accurate). As seen below, the short-term slope is very persistent, with long periods of almost exclusively positive or negative slopes. However, since 2012, it is very nearly flat. Aside from the later period, the behavior for the short term slope of crude oil is similar.

This is because both commodities have similar characteristics from the point of view of storage and convenience yield. Other commodity futures may behave very differently if they are perishable (like milk or eggs) or otherwise not storable (electricity), or if they are produced on a very rigid seasonal basis (corn).

Futures curve slope (1st to 2nd contract)

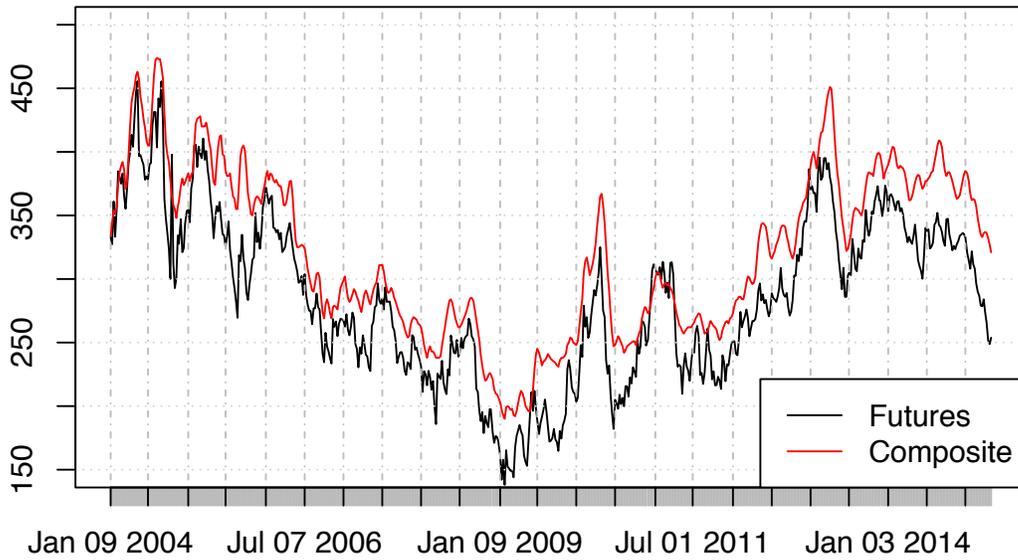


Source: CME Group

The Spot Price

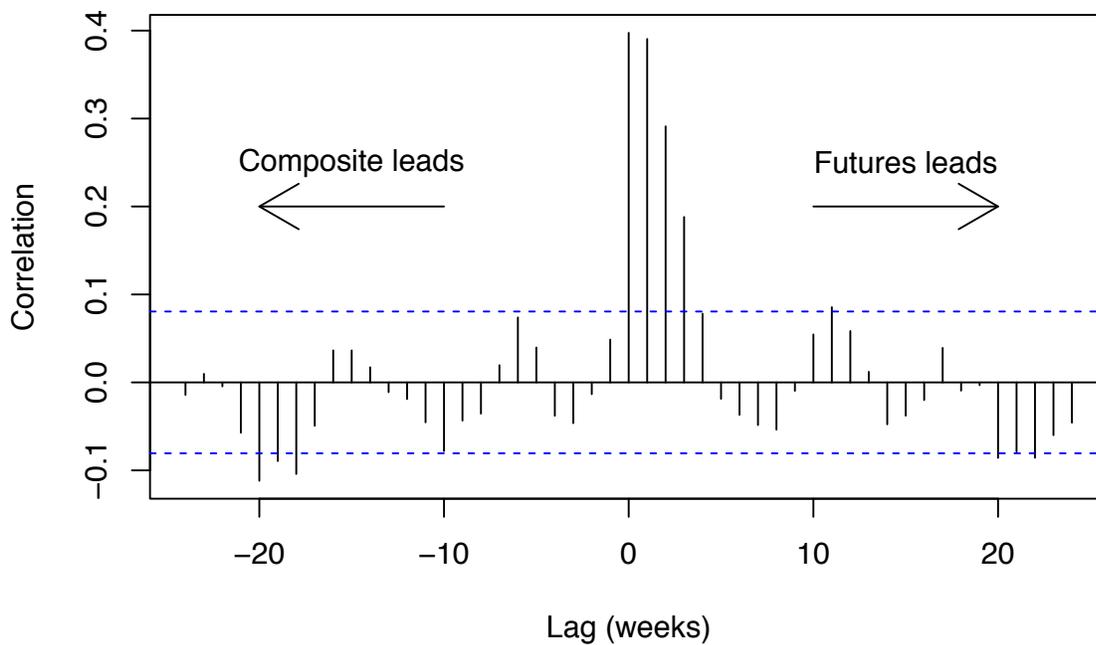
The *spot price* of lumber is unobservable in general. If we need to estimate it, one possibility is to refer to a general benchmark price such as the "Framing Lumber Composite Price" which is published by Random Lengths Publications. This series is designated as the reference price in the Softwood Lumber Products Export Charge Act which determines tariff levels for Canada-US trade. This Composite Price seems to follow the front-month futures price relatively closely, but is typically about 30-40\$ more expensive, even at expiry (aside from a short period following a price surge in 2011). This could have to do with product differences: the Composite Price covers slightly different species, grades, formats and origins of lumber, and the (undocumented) implied distribution of these probably does not match the format distribution prescribed in the futures contract or the distribution across the other attributes which is not prescribed by CME Group but is actually time-varying and market-driven. There may also be differences in the way various costs are accounted for. Also, it is worth noting that the weekly Composite Price changes appear to *lag* behind the other series by 1-3 weeks (see below). This may cast some doubt on whether we can actually use the Composite Price as a proxy for the spot price.

Composite Framing Lumber Price



Source: Random Lengths Publications & CME Group

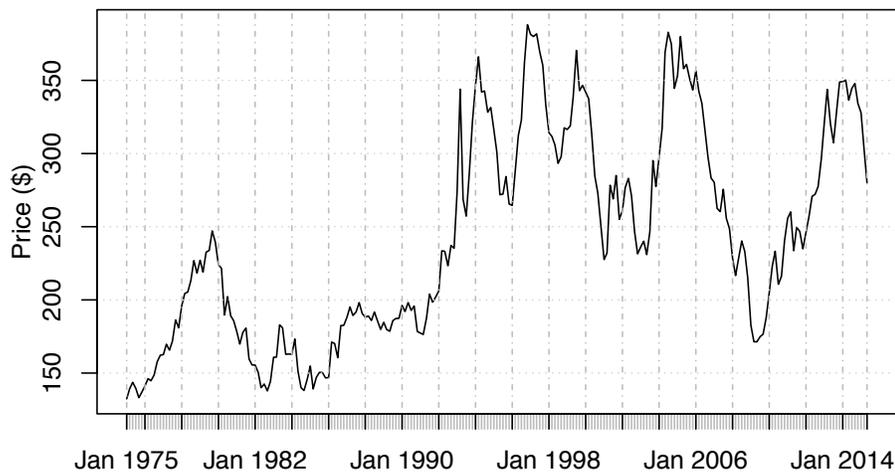
Cross-correlogram



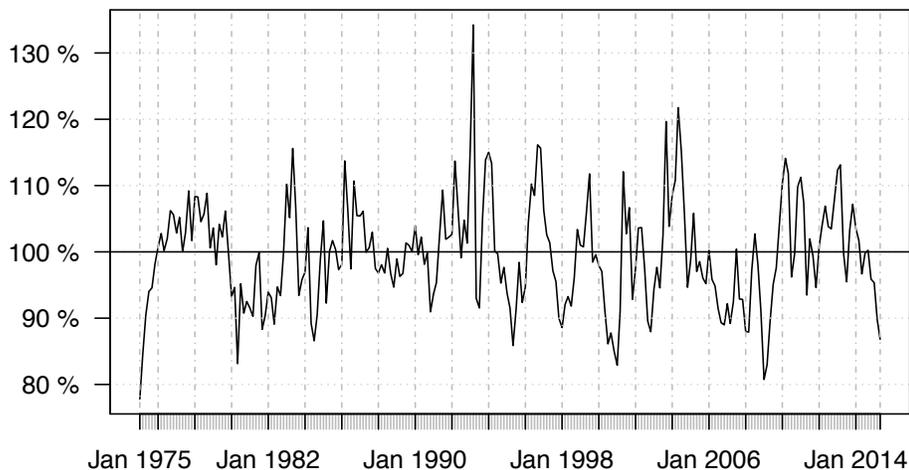
On the other hand, the futures settlement price at expiry is an actual spot price used in a physical transaction: it is the price that all the futures buyers pay for the lumber, for exactly the commodity that underlies the futures (see CME Rulebook, section 20103.E). Unfortunately, using the expiring futures only gives us the spot price once every two months (because they only expire in the middle of January, March, May, July, September and November).

We can gain some insight and try to complete this price series by using a model. In particular, the short-term/long-term model of (Schwartz and Smith 2000) is popular for commodities. It involves two factors: one which represents the long-term trend for the price (a geometric random walk), and the other that represents the short-term deviations from this trend. The short-term deviations are assumed to revert over time, and are quoted multiplicatively (i.e. 100% corresponds to a price which matches the trend).

Long-term Factor

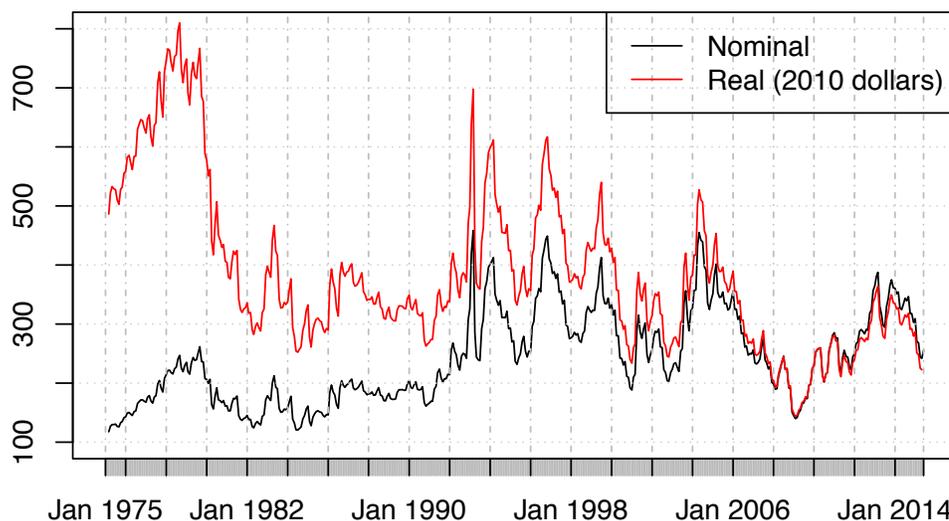


Short-term Factor



We can use the Kalman filter to estimate these two factors from expiring futures prices (January 1975 to May 2015), and the Kalman smoother to interpolate the missing data (i.e. for February, April, June, August, October and December of each year). What we find is that the spot price of lumber grows at roughly 3.5% per year on average (in *nominal* terms), although this is obviously very variable from year to year, and due in great part to inflation (especially during the 1980s). Also, we find that short-term deviations revert with a half-life of 2.5 months. It is also possible to estimate the model indirectly on multiple, longer-dated futures prices simultaneously instead, which is also described in (Schwartz and Smith 2000), but there are technical issues which could make the results less relevant to the present discussion (it requires the additional parametrization and estimation of the so-called Girsanov kernel, which is related to the aggregate risk preferences of market participants). It would be necessary if we needed to evaluate other derivative contracts, but we avoid doing this here.

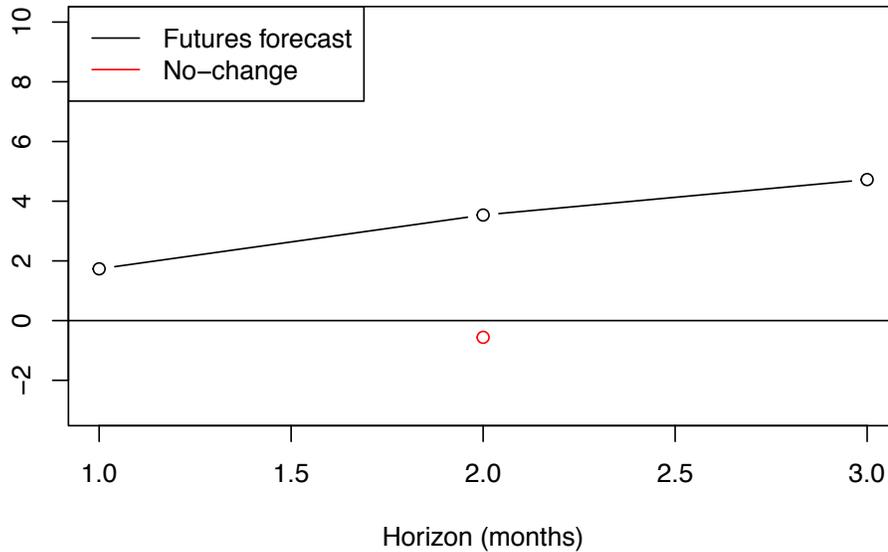
Interpolated Monthly Spot Price (via Kalman Smoother)



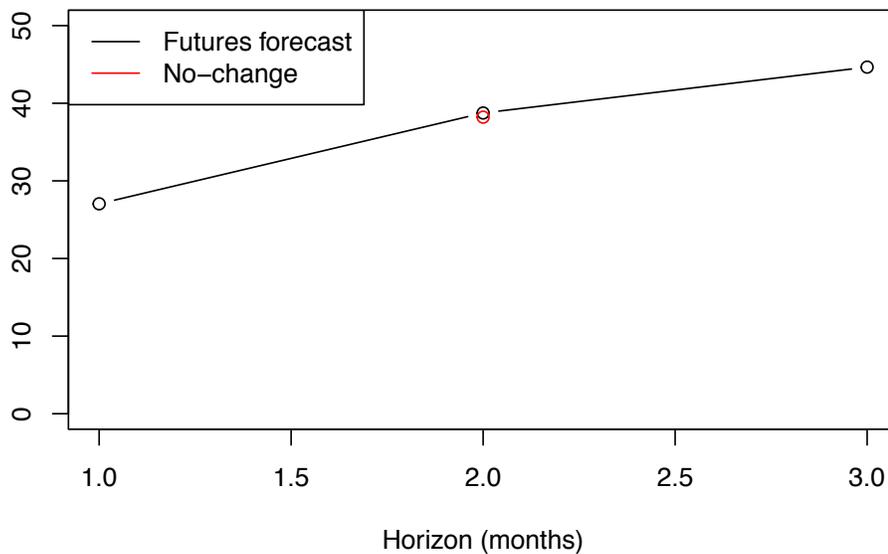
Tests of Forecasting Power

There is a commonly held belief that futures prices are somehow good predictors of the future spot price (for all asset classes, not just commodities). To test this, we take the price of the futures contract at expiry as our spot price, and then compare its value with the price of futures contracts at various horizons. In general, we find that the futures price is an increasingly biased (upwards) estimate of the future spot price over the next three months. It is also not very precise: as a rough guide, the RMSE implies that the forecast is within 50, 75 and 90\$ of the true price at one, two and three months, respectively, approximately 95% of the time. We did not consider the futures further out because they are so thinly traded.

Mean Prediction Error



Root-mean-squared Prediction Error



We would like to compare these futures-based forecasts to others which would be based on the spot price alone. We don't attempt this for an odd number of months ahead because we don't observe the spot price often enough (i.e. we did not use the "interpolated" values from the previous section, so as to minimize the assumptions involved). However, for the two-month horizon, an alternate forecast (in red above) can be obtained by simply copying the current value of the spot price. We find that this no-change forecast is much less biased, but just as precise (in terms of RMSE). We do not find evidence

that the futures-based forecast is more accurate according to the test from (Diebold and Mariano 1995), either. We applied other forecasting models (like the model from (Schwartz and Smith 2000) mentioned previously, and various ARIMA-GARCH models) which were a better fit to the spot price data, but did not ultimately provide better forecasts according to out-of-sample tests. We do not show these results here.

Part of the bias in the futures-based forecast is due to the fact that the futures price is theoretically the *risk-neutral* expectation of the future spot price, and so explicitly does not take into account the risk premium. We can't assume that it doesn't exist in general, and studies like (Fama and French 1987) demonstrated that lumber futures specifically show evidence of time-varying risk premia. This means that when that premium is included, it needs to be forecasted as well. Our other forecasts implicitly include the risk premium, which decreases the bias, but does not improve accuracy partly because the premium is also difficult to forecast.

These findings are consistent with a similar study concerning crude oil futures (Alquist and Killian 2010), where it was found that the no-change forecast was much less biased than the futures price, and that none of a dozen models and expert surveys could outperform it. Because their conclusions are similar, we cannot infer that our findings for lumber are a consequence of the much smaller market size. Also, similar findings are well-known for other asset classes, such as foreign currency (see (Meese and Rogoff 1983)).

On the other hand, contrary to the findings of (Alquist and Killian 2010) on crude oil, we find some evidence that the lumber futures price correctly predicts the direction of spot price moves more often than a coin toss for one and two months ahead (but not three). This finding is statistically significant (according to the test from (Pesaran and Timmermann 1992)), but may not be of practical significance: it is only correct 73% and 62% of the time at one and two months, respectively. Also, even if it was correct *all* of the time, so that a futures price higher than the current spot price always implied a future spot price increase, the payoff of the long futures depends on the relative size of these differences: if the futures is much more expensive than current spot prices and the spot price only increases a bit, the futures position will lose money. The same kind of relationship holds for the reverse situation and the short payoff. Empirically, if we were to enter into a long 2-month futures when in contango, and a short one when in backwardation, on average the payoff (ignoring marking-to-market) would not be positive.

None of the other econometric models we used give directional information which is correct as often as the futures-based forecast, but there may be a technical explanation for this. Typically, the forecast from these models is the *conditional expectation*, which is the best forecast with respect to mean-squared error. There is very little theory on forecasts which are optimal with respect to a directional accuracy metric instead, but it could be possible to derive such a forecast from an econometric model that fits the data well.

Also note that the no-change forecast, which performs so well in terms of bias and accuracy, almost never gives the correct direction *by construction* (the predicted change is always zero). Even though the futures price gives more accurate directions, it gets the magnitude of those changes wrong too often (increases are too big and decreases are too small).

Conclusion

We showed that the futures price at expiry is indeed a spot price (by definition), and that its behavior is similar to that of a popular benchmark price, but that the benchmark is consistently higher and seems to lag behind the price of the front-month futures. We also argued that the longer-term futures contracts may be misleading because they are estimated as parallel shifts due to lack of trading activity, and the evidence for mean-reversion is incompatible with these shifts. Finally, we demonstrated that the futures price is not a particularly good forecast of future spot prices, unless we are only concerned with the direction of changes.

References

Alquist, Ron, and Lutz Killian. 2010. "What Do We Learn from the Price of Crude Oil Futures?" *Journal of Applied Econometrics* 25: 539–73.

Chicago Mercantile Exchange. "CME Rulebook." <http://www.cmegroup.com/rulebook/CME/>.

Diebold, F.X., and R.S. Mariano. 1995. "Comparing Predictive Accuracy." *Journal of Business and Economic Statistics* 13: 253–63.

Fama, Eugene F, and Kenneth R. French. 1987. "Commodity Futures Prices: Some Evidence on Forecast Power, Premiums, and the Theory of Storage." *The Journal of Business* 60: 55–73.

Meese, R., and K. Rogoff. 1983. "Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?" *Journal of International Economics* 14: 3–24.

Pesaran, MH, and A. Timmermann. 1992. "A Simple Nonparametric Test of Predictive Performance." *Journal of Business and Economic Statistics* 10: 461–65.

Schwartz, Eduardo, and James E. Smith. 2000. "Short-Term Variations and Long-Term Dynamics in Commodity Prices." *Management Science* 46: 893–911.