Economic Modelling 64 (2017) 312–321

Contents lists available at ScienceDirect

Economic Modelling

journal homepage: www.elsevier.com/locate/econmod

Catching the curl: Wavelet thresholding improves forward curve modelling

Gabriel J. Power[a,*], James Eaves[b], Calum Turvey[c], Dmitry Vedenov[d]

[a] Department of Finance, Insurance, and Real Estate, FSA ULaval, Université Laval, 2325 rue de la Terrasse, Pavillon Palasis-Prince, Université Laval, Quebec City, QC, Canada G1H 3V8
[b] Department of Management, FSA ULaval, Université Laval, Canada
[c] Dyson School of Applied Economics, College of Business, Cornell University, USA
[d] Department of Agricultural Economics, Texas A&M University, USA

ARTICLE INFO

Keywords:
Forward curve
Futures prices
State-space
Kalman filter
Wavelets
Commodities

ABSTRACT

Modelling futures term structures (price forward curves) is essential for commodity-related investments, portfolios, risk management, and capital budgeting decisions. This paper uses a novel strategy, wavelet thresholding, to de-noise futures price data prior to estimation in a state-space framework in order to improve model fit and prediction. Rather than de-noise the raw data, this method de-noises only wavelet coefficients linked to specific timescales, minimizing the amount of information that is accidentally removed. Our findings are that, for the first five futures maturities in our sample data, in-sample (tracking) and 5-day-ahead out-of-sample (forecasting) Root Mean Squared Errors (RMSEs) are smaller both (i) when we increase the number of factors from one to four, and (ii) when we de-noise the data using wavelet thresholding. The improvement due to wavelet thresholding is often greater than the improvement from adding one more factor to the model, which is important because going beyond four factors does not improve model fit. Wavelet-based de-noising thus has the potential to improve considerably the estimation of various economic time series models, helping practitioners and policymakers with better forecasting and risk management.

1. Introduction

Exchange-traded futures contracts and over-the-counter forward contracts have long been essential instruments for price discovery and risk management (Tomek and Peterson, 2001; Williams, 2001), and their importance has only increased since 2000. Indeed, the trade volume and notional value of commodity forward and futures contracts has increased substantially over the period 2003–2008, sometimes referred to as the “financialization” of commodities, with investment inflows rising from very small amounts to about $250 billion (Irwin and Sanders, 2011; Cheng and Xiong, 2013).

Futures contracts with liquid volume are traded for a large number of maturities, in many cases every month for the first year ahead and at a lower frequency for up to five years into the future. For a given commodity, such as crude oil or corn, this constellation of futures price quotes is called forward curve or futures term structure. It represents aggregate trader information about price expectations and market participant risk aversion (e.g., Schwartz, 1997).[1]

Practitioners and policymakers have great interest in better understanding the entire forward curve or futures term structure (see e.g. Benth et al., 2007; Cortazar et al., 2016; Lautier, 2005), but most academic research focuses on studying the nearby or front-month contract price. Improving the modelling of commodity forward curves, in particular, is essential to traders and portfolio managers, and also matters for the capital budgeting and risk management decisions of corporate firms–especially those involved in oil and gas or in other commodities (e.g. Geman, 2009). Indeed, such firms are known to use strategies such as “pricing against the forward curve”, which is essentially valuation using certainty-equivalent cash flows (e.g. Titman and Martin, 2014). This is because, for purposes of valuation, futures or forward prices can be used instead of forecasted spot prices to obtain certainty-equivalent cash flows, which are then discounted at the risk-free rate instead of a risk-adjusted rate such as the cost of equity from the CAPM. Even if a firm does not hedge using futures, this approach provides the correct valuation for a commodity investment project.

Thus, the in-sample tracking and out-of-sample forecasting of commodity forward curves is an important and challenging problem in economic modelling, and it has significant practical ramifications (e.g., Cortazar et al., 2016; Cortazar et al., 2015; Schwartz and Smith,

* Corresponding author.
E-mail address: Gabriel.power@fsa.ulaval.ca (G.J. Power).

[1] Indeed it is well understood that futures prices represent expected future spot prices, but under the risk-neutral probability measure (e.g. Cox and Ross, 1976; Harrison and Kreps, 1979).

http://dx.doi.org/10.1016/j.econmod.2017.03.032
Received 14 February 2017; Received in revised form 27 March 2017; Accepted 30 March 2017

2000). This paper's contribution is to show how to improve this modelling using a new approach based on wavelet thresholding, which is then applied to data on commodity futures contract prices traded at the Chicago Mercantile Exchange (formerly Chicago Board of Trade contracts).

Indeed, the better we can separate the signal from the noise in time series data, the more useful the data becomes for making predictions. This explains why there is a large literature concerned with de-noising data. Wavelet thresholding (Donoho and Johnstone, 1994a, 1994b; Donoho,1995) is a de-noising method that has been mostly overlooked by economists. But what is innovative about it is that instead of de-noising the raw data, it de-noises the data's wavelet coefficients at the finest time scale, which reduces the amount of information that is accidentally removed. The intermediate step (see the Appendix) involves applying to the raw data a discrete wavelet transform prior to de-noising, and then applying an inverse wavelet transform to the de-noised wavelet coefficients, resulting in efficiently filtered data that can be used for estimation.

This paper demonstrates the potential of wavelet thresholding by improving the modelling of forward curves, using a multi-factor price model with several correlated sources of risk. For our sample, de-noising the data this way improves the tracking and forecasting results in most cases, suggesting that the approach should be seriously considered by commodity investment decision-makers, whether for investments, risk management, or capital budgeting.

The economic intuition behind wavelet thresholding in our setting of futures contracts is that price variations occurring below some threshold is only noise. Indeed if variation contributed to the price signal, it should be linked to longer time scales. This reasoning is similar to the argument made by Hasbrouck (2013) who shows how to use wavelet variance decomposition to identify and measure micro-structure volatility and noise. Since price changes contain both information and noise, filtering out the noisy portion prior to fitting the model will improve the efficiency of the estimation. Our empirical results confirm this.[2]

It is fair to ask whether what we call “noise” may in fact be economic news. However, there are at least two strong reasons why it is worth finding better ways to “de-noise” price data prior to fitting a forward curve model. First, there is also an entire literature based on Shiller (1981) that argues (theoretically and empirically) that traders over and under-react to information. So there is strong reason to believe that price changes contain noise and, consequently, increasing the signal/noise ratio is helpful. Thus, if it improves forecasting it must be increasing the signal/noise ratio. Second, the wavelet thresholding approach we propose removes only the component that has a one-day horizon (i.e., the one-day horizon wavelet function). The previous literature (e.g. Hasbrouck, 1991) shows that price changes are least partially noise, and that price changes caused by the arrival of new trades are partly noise and partly information. Our hypothesis is that the portion of price changes that is noise can be identified using the wavelet one-day horizon (timescale) and it is then is removed. Meanwhile the part that is information is identified by the wavelet two-day horizon or longer timescales, and these components of the data are not removed. Thus, we plausibly filter out noise but not useful information.

2. Literature review

2.1. The term structure of commodity futures prices

The futures price F_t for a given date t and maturity T equals the time t expectation of the spot price S_T at maturity T under the risk-neutral probability measure Q (Black, 1976; Cox et al., 1981, 1979; Harrison and Kreps, 1979). It is well understood, therefore, that the futures price is a risk-adjusted forecast of spot price and thus it reflects both the market's expectations as well as a risk adjustment.

$$F(x_t, t, T) = \mathrm{E}_t^Q(S_T) \tag{1}$$

In a simple model of the forward price curve for storable commodities, the following relationship holds at all times:

$$F(t, T) = F(t, t)\exp^{(r+c-\delta)(T-t)} \tag{2}$$

where *F(t,t)* is the futures price for a contract expiring today (i.e., equal to the spot price notwithstanding basis risk), r is the risk-free rate of interest (e.g. 3-month U.S. Treasury bill), c is the cost of carry and δ is the convenience yield. The shape of the forward curve depends only on the net convenience yield: $r+c-\delta$. If $r+c > \delta$, contango results, and if $r+c < \delta$, backwardation results. The convenience yield represents the economic value of holding physical stocks of a commodity, e.g. the benefits of holding inventories to maintain a smooth running commercial operation and avoiding the risk of stock-outs. This definition has, however, been debated in the literature (e.g. Brennan et al., 1997; Williams, 2001). This concept provides a useful way to link commodity inventory levels with the shape of the forward curve. An example of a commodity futures price term structure is presented in Fig. 1 for Chicago Board of Trade corn futures on 6/17/2004, a period when the market was in contango. For some energy and agricultural commodities, the forward curve is also affected by seasonal cycles (Tomek, 1997, 2000; Fackler and Roberts, 1999).

Under the risk-neutral measure, asset price dynamics imply the following relationship (e.g. Fackler and Roberts, 1999):

$$\mu + \delta = r + \sigma\lambda \tag{3}$$

where μ is the actual drift term, δ is the convenience yield, r is the risk-free rate of interest, σ is the diffusion term, and λ is the market price of risk for the state variable in question. The equation may be rearranged to give:

$$\mu - \sigma\lambda = r - \delta \tag{4}$$

which implies that the risk-adjusted drift μ - σλ equals the risk-free rate minus the convenience yield, r - δ. The convenience yield can be estimated, because the left-hand side parameters are estimated from the data using the above model (2), while the 3-month US Treasury bill provides a good proxy for the risk-free interest rate. For multi-factor models as the one described below, additional parameters must be incorporated in the equation, but the approach is similar. .

The most popular approaches used to model the term structure of futures prices are called “reduced-form” and describe the stochastic process generating the futures price term structure using a small number of “factors” each of which is described by a stochastic process.[3] The first approach aims to estimate the unobservable convenience yield of a real or financial asset, which helps explain the shape of the forward curve (e.g., Brennan and Schwartz, 1985; Gibson and Schwartz, 1990). The second approach, which is more general, models the asset price as an affine function of state variables, which are usually unobservable (e.g., Schwartz and Smith, 2000). We therefore use this more general approach to model the forward curve.

2.2. Wavelet-based methods for economic time series

Our main contribution to the literature is to show how wavelet thresholding, applied to the futures price data prior to estimation of the forward curve model, improves the tracking and forecasting perfor-

[2] Although this study is the first to our knowledge to use this empirical strategy, it is worth noting that Haven et al. (2012) use other wavelet methods to de-noise option prices.

[3] An entirely different approach, called “structural”, specifies stochastic processes separately for supply and demand side shocks (see e.g. Pirrong, 2011). This approach is computationally more challenging.