

Scarcity risk premium*

Thibault Lair[†]

This version: December, 2019

Abstract

This paper revisits the cost-of-carry model and proposes a decomposition of the futures basis that disentangles the seasonality risk premium from the scarcity risk premium. The contribution of this paper to the asset pricing literature is threefold. First, it brings novel insights on the fundamental relationship between the futures basis and inventory dynamics. Empirical evidence shows that the seasonality risk premium captures expectations about the inventory seasonalities while the scarcity risk premium reflects the excess supply and demand imbalances over the expected seasonal fluctuations. Second, this paper investigates the pricing of expectations within the futures basis. Results suggest that the seasonality risk premium is priced-in and that the scarcity risk premium carries all the predictive power embedded in the futures basis. The third contribution of this paper is to provide evidence that the scarcity risk premium, and ultimately the futures basis, is mostly a compensation for the unexpected increase in the risk of stock-out and that the associated return predictability finds its origin in the slow diffusion of information and underreaction to abnormal changes in inventories, above and beyond seasonal dynamics.

Keywords: Commodities, theory of storage, seasonality, scarcity, risk premium, underreaction, market anomalies, factors.

*I would like to thank Frank Fabozzi, Abraham Lioui, Guido Baltussen, Alexandre de Roode, and seminar participants at the 2019 EFMA Conference and 2019 EDHEC PhD in Finance Forum for valuable discussions and useful feedback. Comments are welcomed, including references to related papers but inadvertently overlooked. The views expressed here are those of the author and not necessarily those of any affiliated institution.

[†]Lair is a PhD in Finance candidate at the EDHEC Business School and from Robeco Institutional Asset Management. Email: lair.thibault@gmail.com

Introduction

The cost-of-carry relationship is the corner stone of the theory of storage. It assumes the difference between the futures and spot prices, i.e. the futures basis, can be explained by the interest foregone when buying commodities in the physical market, the associated storage costs and a convenience yield. The latter is defined as a benefit that accrues to the commodity holder resulting from the potential productive value of his inventory. This value of physical ownership reflects market expectations about the future availability of a commodity.

A vast literature provides evidence that the futures basis is informative about futures risk premia. This paper sheds light on the origin of this return predictability by disentangling various expectations embedded in the basis. As production and demand seasonality can have a large influence on the future state of inventories, a convenience yield decomposition is proposed to dissociate known seasonal supply and demand imbalances from abnormal shocks above and beyond those transitory dynamics. Agents form expectations about the impact of inventories' seasonality on the basis (henceforth the seasonality risk premium) and adjust their conditional expectations as new information arrives. Changes in expectations resulting from unanticipated supply and demand shocks (henceforth the scarcity risk premium) are associated with the incremental risk of scarcity and capture the non-linearities in the futures basis.

The contribution of this paper to the asset pricing literature is three-fold. First, it brings novel insights on the fundamental relationship between the futures basis and inventory dynamics. While the literature has documented the relationship between the basis and the level of inventories, as well as the presence of seasonality in both the basis and the inventories, the question of how various inventory dynamics are priced in the cross-section of commodity markets remains an open research topic. This paper investigates whether the predictable seasonality in the convenience yield is related to known seasonal fluctuations in inventories, and whether the residual, or seasonally-adjusted, convenience yield is associated with unexpected supply and demand shocks. It provides strong empirical evidence that the convenience yield embeds expectations about those two different inventory dynamics. On one hand, the seasonal risk premium captures expectations about the inventory seasonalities. On the other hand, the scarcity risk premium reflects the excess supply and demand imbalances over the expected seasonal fluctuations.

Second, this paper contributes to the understanding of the risk premia in the cross-section of commodity markets by investigating the pricing of expectations within the futures basis. It provides strong empirical evidence that the seasonality risk premium is priced-in as unconditional expectations about the seasonal inventory dynamics carry no predictive power. On the contrary, the conditional expectations about the impact of abnormal supply and demand shocks earn a premium and carry predictive power for future returns. In fine, the scarcity risk premium carries all the predictive power embedded in the futures basis. Other basis components, i.e. the foregone interest and the net cost of storage, have no information content. Those findings are robust to the introduction of commodity sectors and seasons as control variables.

The third contribution of this paper is to show that the scarcity risk premium, and ultimately the futures basis, is a compensation for the un-

expected increase in the risk of stock-out and that the associated return predictability finds its origin in the slow diffusion of information and underreaction to abnormal changes in inventories, above and beyond seasonal dynamics. It presents strong evidence of underreaction along the futures curve following a marginal change in the risk of stock-out. The underreaction is broad based along the curve but more pronounced in the far contracts, leading to futures curve twists. These results suggest that market participants also underreact to the risk that the inventory depletion happens at faster rate than anticipated, as well as the risk that inventory imbalances might resorb at a lower speed than expected due to the slow adjustment of supply. Despite the slow diffusion of information and the unexpected nature of the supply and demand shocks, information decays fast and the return predictability is transient. The diminishing statistical significance of excess returns with the length of the holding horizon indicates that the information is progressively diffused throughout the market up until it is fully incorporated in futures prices. The scarcity risk premium has significant predictive power up to a 3 months horizon, which is too short to be associated with a fundamental resolution of abnormal supply and demand imbalances.

This paper is organized as follows. Section 1 reviews the literature on the theory of storage and on the seasonality in commodity markets. Section 2 proposes a decomposition of the futures basis that disentangles expected seasonal shocks from unexpected shocks to inventories. Section 3 presents the data and methodologies used in this paper. Section 4 investigates the fundamental relationship between the futures basis and the inventory dynamics. Section 5 provides empirical evidence on the pricing of expectations in the cross-section of commodity markets. The last section analyses the origins of the return predictability associated with the scarcity risk premium. To conclude, future avenues of research are discussed.

1 Literature review

The cost-of-carry relationship describes the no-arbitrage condition between the spot and the futures price and defines the futures basis. In commodity markets the traditional definition is augmented by a term very specific to perishable goods, the convenience yield. This value of physical ownership is the corner stone of the theory of storage and reflects market expectations about the future availability of a commodity. More importantly, the future basis has been found to be informative about futures risk premia. As production and demand seasonality can have a large influence on the anticipations about the future state of inventories, this paper aims at shedding light on the origin of this predictability by disentangling various expectations embedded in the basis.

1.1 Theory of storage

The theory of storage which finds its foundation in the work of Kaldor (1939) and Working (1948) has been largely documented in the literature on commodity markets. Its corner stone, the cost-of-carry relationship,

assumes the difference between futures prices and spot prices can be explained by the interest foregone when buying commodities in the physical market, the associated storage costs and a convenience yield. The latter is defined as a benefit that accrues to the commodity holder resulting from the potential productive value of this inventory. Various model specification have been proposed in the literature to define this cost-of-carry relationship. Fama and French (1987) consider a model with fixed marginal storage cost and convenience yield. Szymanowska et al. (2014) propose an alternative definition with proportional storage costs that accrues per period.

$$F_t^{(n)} = S_t \left((1 + RF_t^{(n)})^n (1 + U_t^{(n)})^n - C_{t+n} \right)$$

$F_t^{(n)}$ is the futures price at time t expiring in n -periods and S_t is the spot price at time t . $RF_t^{(n)}$ is the per-period risk-free rate at time t with maturity $t + n$ and known at t . $U_t^{(n)}$ is the per-period physical storage costs, for storing commodities over n periods, expressed as a percentage of the spot price and known at t . Finally, the convenience yield C_{t+n} is defined as a cash payment occurring at time $t + n$ and is also known in t .

The convenience yield can be interpreted as the net income, valued at time t , that the physical holder requires to sell his inventory at time $t + n$ at the price $F_t^{(n)}$, once he has been compensated for the interest foregone and for the storage costs he is facing. Likewise, the convenience yield represents the net amount the futures investor is willing pay, beyond the current spot price S_t , i.e. the interest and storage costs compensation required by the commodity seller, in exchange for settling the purchase of the physical commodity at a price $F_t^{(n)}$. As such the convenience yield is foremost an implied quantity that rules out any arbitrage opportunity in the cost-of-carry model and describes the current equilibrium in futures markets.

A related concept is the futures basis which ties the current futures price $F_t^{(n)}$ maturing in n periods to the current spot price S_t and summarizes the cost-of-carry relationship. The futures basis has also been coined the futures carry or roll-yield in the literature on commodities and is similar to other asset classes carry definition.¹ The below equation defines the n -period log or percentage basis $y_t^{(n)}$, i.e. the per-period futures carry for maturity n .

$$F_t^{(n)} = S_t \exp \left\{ y_t^{(n)} \times n \right\}$$

Accordingly, the basis is a reflection of the foregone interest, the storage costs and the convenience yield and can be directly estimated from observed futures and spot prices.

$$y_t^{(n)} = \frac{1}{n} \ln \left\{ \left((1 + RF_t^{(n)})^n (1 + U_t^{(n)})^n - \frac{C_{t+n}}{S_t} \right) \right\} = \frac{1}{n} \ln \left\{ \frac{F_t^{(n)}}{S_t} \right\}$$

A large part of the literature on the theory of storage has focused on trying to model the dynamics of supply and demand that could explain the empirical distribution of commodity prices. Most commodities display non-normal distributions with high positive skewness and kurtosis, price

¹Koijen et al. (2018) provide evidence of the presence of a carry factor across multiple asset classes.

jumps, non-linearities in prices and in conditional variance, as well as autocorrelation. The main hypothesis put forward to explain the behavior of prices are the inelasticity of demand in the presence of supply shocks with a non-linear functional form, the slow adjustment of supply and the lower boundaries on the level of inventories.

The impossibility to carry negative inventories plays a crucial role in the model of storage dynamics by Deaton and Laroque (1992) as it introduces the non-linearity able to characterize both the observed price jumps and conditional variance. Routledge et al. (2000) propose an equilibrium model for the commodity forward curve in which the non-negativity constraint on inventory levels introduces an immediate consumption timing option in spot prices. The value of this option fluctuates with inventories and transitory shocks to supply and demand.

Gorton et al. (2013) provide the first comprehensive study of the relationship between physical inventories and commodity futures. They show that the shape of the futures curves is associated with the level of inventories and that this relationship becomes highly non-linear as the risk of a stock-out increases. This confirms the model predictions of Routledge et al. (2000) in which the elasticity of the convenience yield to changes in inventory levels is a decreasing function of inventory levels. Hevia et al. (2018) find that the non-linearity between inventories and commodity prices is negatively related to the maturity of the futures contract.

1.2 Seasonality

Seasonality plays an important role in the theory of storage as both production and demand can exhibit seasonal patterns that impact inventories. The resulting inventory seasonality should influence directly the convenience yield embedded in the basis and thus both the futures price and its carry. Figure 1 illustrates the presence of seasonality in both commodity futures prices as well as in the futures basis. Panel a of Figure 1 plots the futures curve of the Gasoline RBOB contract on 2016-01-05. Panel b of Figure 1 shows the autocorrelation coefficients of the futures basis in excess of the foregone interest for Lean Hogs futures contracts over 252 daily lags.

Finding support in the theory of storage and in the work of Brennan (1958) and Telser (1958) on the relationship between the convenience yield and the level of physical inventories, Fama and French (1987) test for the presence of seasonal variation in the basis using seasonal dummies. They find reliable evidence of seasonality in some agricultural and animal commodities but none in metals. Moreover, they provide evidence that the basis has forecasting power for future spot changes of seasonal commodities. More recently Brooks et al. (2013) have confirmed those findings using a larger sample size and a broader universe of commodities. They provide additional robustness checks on the presence of seasonality in the basis and show that the forecasting power of the basis cannot be related to the magnitude of the seasonal patterns a commodity exhibits.

Fama and French (1987) argue that for commodities that exhibit seasonality in supply or demand, the predictability of future spot variation should be increasing in storage costs. Indeed, high costs should deter the build-up of inventories whose main function is to smooth demand and supply imbalances. The presence of forecasting power for some of the

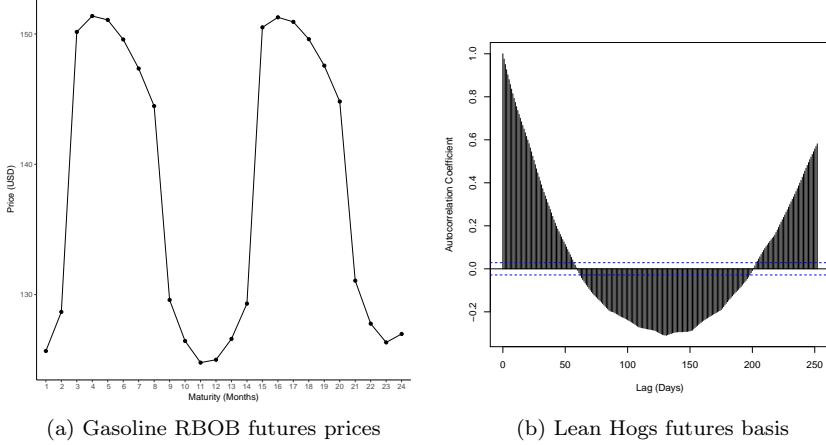


Figure 1: Illustration of seasonality in commodity contracts prices and futures basis

perishable commodities (i.e. with high storage costs) and its absence for metals should support this hypothesis.

French (1986) proposes a two-period model of production, consumption, and storage to describe the relation between the expected seasonal variation in spot prices, the convenience yield as well as the underlying seasonality of demand and supply. Intuitively, in this model, the sensitivity of the convenience yield to the level of inventories will influence how pronounced the seasonal patterns in spot prices are. Moreover, the elasticity of the convenience yield to changes in inventory levels is a decreasing function of inventory levels.

Following the Fama and French studies, a large strand of the literature has focused on incorporating seasonality in the modeling of the futures curve dynamics. A number of publications have proposed multi-factor state-space model representations under no-arbitrage conditions and constant risk premium, which include seasonal terms to model seasonal variations both in spot prices and volatilities. Among those, Sørensen (2002) introduces a deterministic seasonal component in the pricing of futures. He prefers as an alternative to the standard dummy variable estimation of the deterministic component an evaluation through a linear combination of trigonometric functions, following the approach of Hannan et al. (1970).

Next to the deterministic seasonal component present in spot prices, Geman and Nguyen (2005) incorporates in the dynamics of the deseasonalized stochastic component of the spot price a deterministic term to model the seasonal variation observed in the volatility of commodity futures. This follows from Deaton and Laroque (1992) who document that the conditional variance of commodity prices is increasing in the level of prices. If spot price increases are driven by seasonal variations in the level of inventories, then it is likely that the coincident volatility spikes originate from the seasonality observed in physical inventories.

Both papers share a limitation when it comes to the estimation of the seasonal component. More specifically, the later is function of the

estimation date rather than of the maturity of the contract. Borovkova and Geman (2006) tackle this issue and consider a cost-of-carry model with a deterministic seasonal premium within a stochastic convenience yield. By relying on an average futures price along the curve, this approach does not disentangle the interest foregone from the seasonal component in the estimation.

More recently Hevia et al. (2018) develop a multi-factor affine model of commodity futures with stochastic seasonal fluctuations. In their nine factors approach, the seasonal shocks are driven by two unobserved factors which only explain marginally the observed risk premia but still are non-negligible. Most of the risk premia originate from the spot factor and the three factors (i.e. level, slope, curvature) describing the cost-of-carry term structure. Interest rate factors' influence increases with the level of interest rate. Finally, the authors reveal the importance of allowing for time variation in the estimation of the seasonal component to match the time-variation of seasonal patterns observed in inventories. More importantly this helps to avoid attributing erroneously those dynamics to other factors. The presence of non-linearities between inventories and the net convenience yield supports the theory of storage.

This paper differs from this strand of the literature in that it focuses solely on the identification of a seasonal risk premium within the futures basis rather than the modeling the futures curve dynamics. In that sense, it does not impose any structure neither on the forward curve nor on the spot or futures price dynamics, beyond the cost-of-carry relationship. Its contribution to the literature on seasonality within commodity markets is twofold. First, this paper addresses some of the concerns noted above on the estimation of the seasonal premium. More specifically, we consider the delivery month of futures contracts to define seasons and exclude the foregone interest from the estimation. This allows one to clearly exclude any term premium that would originate from the interest rate compensation. Moreover, using all contracts along the futures curve provides robust estimates of the seasonal premium per commodity.²

Another strand of literature addresses the presence of seasonality within general equilibrium models. Most notably, Hirshleifer (1991) investigates the hedging decision with sequential arrival of information and agents maximizing utility derived from a multi-good consumption. Contrary to most partial equilibrium models, the futures price is here a martingale.³ The author shows that the optimal hedging policy is affected by amongst others, the correlation of the producer's output with the aggregate, the relative sensitivity of his production to the environment versus other producers and the demand price elasticity. In this setup, the sequential arrival of information resolves uncertainty and defines the optimal hedging policy of producers. This in turn drives the seasonal patterns observed in commodity futures markets. Interestingly, Hirshleifer refutes the general pre-conception of a short hedging pressure resulting from producers willing to hedge their bottom line volatility. He argues for example that when demand shocks are arising from changes in aggregate wealth, the optimal hedging policy of producers with non-stochastic output might be to remain unhedged or even to take a long hedge position.

²Fama and French (1987)

³This is based on the assumption of additively separable preferences, complete markets and non-stochastic endowment of the numeraire.

Finally, a more recent strand of literature has focused on providing evidence of seasonality in returns across asset classes and of the profitability of the strategies aiming at exploiting these return dynamics.

Keloharju et al. (2016) document the presence of seasonality in returns across various stock portfolios and factors as well as in commodities using same month past returns as predictor. The authors investigate the source of the seasonality observed at the stock level. They suggest it could either emanate from the underlying risk factors' seasonality, thus having a systematic origin, or resides within each individual security, being therefore essentially of idiosyncratic nature. An alternative explanation to the recurrent variation in securities returns is that seasonality is merely the result of serial-correlation in innovations. Within stock markets, the authors find supportive evidence for the hypothesis that seasonality at the stock level is induced by their exposure to the underlying risk factors own seasonality. Finally, they show that return-based seasonal strategies are economically significant, cannot be explained by exposure to macroeconomic risks and are resilient to investors sentiment.

This study adds to this literature by investigating the systematic nature of returns seasonality in commodity markets. Moreover, its main contribution is to characterize and quantify the the risk premium attached to seasonality while dissociating it from the documented return-based seasonal strategies.

2 Supply and demand expectations

Seasonality plays an important role in the theory of storage as both production and demand can exhibit seasonal patterns that impact inventories. The resulting inventory seasonality influences directly the convenience yield embedded in the basis and thus both the futures price and its carry. It is only fair to assume that agents and market participants have expectations about the seasonal impact on the basis resulting from seasonality in inventories. Without any information on the future state of supply and demand, they can form priors based on the information available to them at any point in time, i.e the current filtration. As new information arrives agents adjust their conditional expectations away from their original prior to account for the anticipated marginal change in supply and demand imbalance. These adjustments in expectations for "unexpected" shocks to inventories characterize the scarcity risk. In order to differentiate the seasonality premium from the scarcity premium, I propose in this section a decomposition of the futures basis in excess of the long-run futures cost-of-carry that allows to disentangle the effect on the convenience yield of expected seasonal shocks versus unexpected shocks to inventories.

2.1 Definitions

Let us define the convenience yield C_{t+n} with Z_{t+n} the seasonal impact on the basis resulting from seasonality in inventories and X_{t+n} the residual convenience yield in excess of the seasonal basis contribution. This decomposition allows to disentangle the effect on the convenience yield of expected seasonal shocks versus unexpected shocks to inventories.

$$C_{t+n} = Z_{t+n} + X_{t+n}$$

Here Z_{t+n} can be interpreted as the premium futures investors are willing to pay at time $t + n$ respectively to buy commodities at a forward date in the presence of known transitory supply and demand imbalances. When inventories are expected to be low, commodity holders value the future productive capacity of their current inventory and are requiring a premium to sell their inventory forward (i.e. a negative Z_{t+n}). When inventories are expected to be high, the future productive value of inventories diminishes and commodity holders are willing to offer their current stock at a forward discount (i.e. a positive Z_{t+n}).

The residual convenience yield X_{t+n} arises from unexpected supply and demand imbalances. Unexpected imbalances are here understood as the expectations given the current set of information on future risks beyond "normal" seasonal expectations Z_{t+n} . X_{t+n} defines the scarcity risk premium and captures the non-linearities observed in the basis⁴.

I assume that agents learn about seasonal patterns over time. They have the ability to forecast temporary supply and demand imbalances based on their knowledge about the production and demand cycles or simply by evaluating their impact on the futures basis. Assuming a m periods seasonality cycle, the expectation conditional on information up to time t of the seasonal Z_{t+n} is thus equal to last year same period seasonal convenience yield. It is a martingale.

$$E_t[Z_{t+n}] = Z_{t+n - m_n^+}$$

$$m_n^+ = m \times \lceil \frac{n}{m} \rceil$$

Here m_n^+ defines the number of complete seasonal cycles we need to look backward in order to obtain the last observable seasonal convenience yield with the same seasonal period as time $t + n$. The notation $\lceil i \rceil$ describes the largest integer not less than i . It is assumed that the observations have the same frequency as the seasonal periodicity.

2.2 Basis decomposition

The basis can be further decomposed to incorporate this distinction between the expected seasonal convenience yield and the unexpected residual convenience yield. In the rest of the paper log prices will be denoted using lower cases.

⁴Routledge et al. (2000) propose an equilibrium model of commodity forward curve in which the non-negativity constraint on inventory levels introduces an immediate consumption timing option in spot prices. The value of this option fluctuates with inventories and transitory shocks to supply and demand. Gorton et al. (2013) provide the first comprehensive study of the relationship between physical inventories and commodity futures. They show that the shape of the futures curves is associated with the level of inventories and that this relationship becomes highly non-linear as the risk of a stock-out increases. This confirms the model predictions of Routledge et al. (2000) in which the elasticity of the convenience yield to changes in inventory levels is a decreasing function of inventory levels. Also, the authors document that both the basis and the momentum factors are loading on commodities with current low physical inventories. Both risk premia are thus a compensation for bearing the risk of further deterioration in inventory levels.

$$\begin{aligned}
y_t^{(n)} &= \frac{1}{n} \ln \left\{ \left(1 + RF_t^{(n)} \right)^n \left(1 + U_t^{(n)} \right)^n - \frac{C_{t+n}}{S_t} \right\} \left(\right. \\
&= \frac{1}{n} \ln \left\{ \left(1 + RF_t^{(n)} \right)^n \left(1 + U_t^{(n)} \right)^n \left(1 - \frac{Z_{t+n} + X_{t+n}}{S_t \left(1 + RF_t^{(n)} \right)^n \left(1 + U_t^{(n)} \right)^n} \right) \right\} \left(\right. \\
&= r f_t^{(n)} + u_t^{(n)} + \frac{1}{n} \ln \left\{ \left(1 - \frac{Z_{t+n} + X_{t+n}}{S_t \left(1 + RF_t^{(n)} \right)^n \left(1 + U_t^{(n)} \right)^n} \right) \right\} \left(\right)
\end{aligned}$$

The last term of the equation can be log-linearized following the approach used by Campbell and Shiller (1988). I thus proceed to a first-order Taylor series expansion around the long-run mean of both components of the convenience yield, respectively \bar{Z} and \bar{X} .

$$\begin{aligned}
\ln \left\{ \left(1 - \frac{Z_{t+n} + X_{t+n}}{S_t \left(1 + RF_t^{(n)} \right)^n \left(1 + U_t^{(n)} \right)^n} \right) \right\} &\simeq \ln \left\{ \left(1 - \frac{\bar{Z} + \bar{X}}{S_t \left(1 + RF_t^{(n)} \right)^n \left(1 + U_t^{(n)} \right)^n} \right) \right\} \left(\right. \\
&\quad \left. + \frac{(Z_{t+n} - \bar{Z}) + (X_{t+n} - \bar{X})}{S_t \left(1 + RF_t^{(n)} \right)^n \left(1 + U_t^{(n)} \right)^n \left(- (Z_{t+n} + X_{t+n}) \right)} \right)
\end{aligned}$$

Furthermore, I impose a mild distributional restriction on the convenience yield by assuming its unconditional mean is zero. This assumption is easily met by detrending the convenience yield and requiring that any structural benefit the commodity owner might earn from physically holding inventories is reflected in a reduced net storage cost $U_t'^{(n)}$.

$$E[C] = 0$$

Given the mean across seasons of an additive seasonal component is zero for a detrended series, the unconditional mean of Z is null.

$$E_m[Z] = \sum_{i=1}^m Z_i = 0$$

$$E[Z] = 0$$

As a result, the unconditional mean of the residual convenience yield is also zero. Thus X_{t+n} captures the convenience yield resulting from transitory inventory imbalances arising from unexpected shocks in supply and demand.

$$E[X] = 0$$

Assuming the long-run means \bar{Z} and \bar{X} equal their respective unconditional mean, the log-linearization now simplifies to a sum of two terms, which are simply the proportions of the futures price that can be explained by respectively the seasonal component and the residual convenience yield.

$$\ln \left\{ \left(1 - \frac{Z_{t+n} + X_{t+n}}{S_t \left(1 + RF_t^{(n)} \right)^n \left(1 + U_t'^{(n)} \right)^n} \right) \right\} \left(- \frac{Z_{t+n}}{F_t^{(n)}} - \frac{X_{t+n}}{F_t^{(n)}} \right)$$

The per-period log-basis now simplifies to a linear equation with four terms capturing the foregone interests, the net storage costs, the seasonal premium and the residual convenience yield in excess of this seasonal component, i.e the scarcity premium.

$$\begin{aligned} y_t^{(n)} &\simeq r f_t^{(n)} + u_t^{(n)} - \frac{1}{n} \left(\frac{Z_{t+n-m}^+}{F_t^{(n)}} + \frac{X_{t+n}}{F_t^{(n)}} \right) \\ y_t^{(n)} &= r f_t^{(n)} + u_t^{(n)} - \zeta_t^{(n)} - \chi_t^{(n)} \end{aligned}$$

Here $\zeta_t^{(n)}$ defines the per-period proportion of the futures prices explained by the seasonal premium. Similarly $\chi_t^{(n)}$ represents the per-period proportion of the futures prices explained by the scarcity premium.

$$\begin{aligned} \zeta_t^{(n)} &= \frac{1}{n} \left(\frac{Z_{t+n-m}^+}{F_t^{(n)}} \right) \\ \chi_t^{(n)} &= \frac{1}{n} \left(\frac{X_{t+n}}{F_t^{(n)}} \right) \end{aligned}$$

The foregone interests and the net storage costs can be grouped together to define the long-run cost-of-carry $\Upsilon_t^{(n)}$. The time and futures curve variation observed in $\Upsilon_t^{(n)}$ result solely from the dynamics of interest rates and the shape of the yield curve as the net storage costs are assumed constant across the curve and equal to their full sample conditional mean⁵.

$$\Upsilon_t^{(n)} = r f_t^{(n)} + u_t^{(n)}$$

The basis can thus be expressed as a function long-run cost-of-carry, a seasonal premium and a scarcity premium.

$$y_t^{(n)} = \Upsilon_t^{(n)} - \zeta_t^{(n)} - \chi_t^{(n)}$$

Ultimately, we can incorporate this basis decomposition in the cost-of-carry relationship with the futures price being defined as follows.

$$F_t^{(n)} = S_t \exp \left\{ y_t^{(n)} \times n \right\}$$

$$F_t^{(n)} = S_t \exp \left\{ \left(f_t^{(n)} - \zeta_t^{(n)} - \chi_t^{(n)} \right) \times n \right\}$$

3 Data and methodology

3.1 Data

3.1.1 Futures contracts

The data set covers 29 commodity futures contracts eligible for inclusion in the Bloomberg Commodity Index over a period going from 1983-05-31 to 2018-08-21. It contains prices, open interest and volumes as well as various static information about the contract (e.g. the maturity date, contract size or minimum tick size). The sector classification follows the Bloomberg sector indices. The data is obtained directly from the CME or ICE database or via Datastream for LME contracts. Data on short interest rates are sourced from the FRED database. I use USD Libor rates for

⁵This assumption can be relaxed to incorporate non-seasonal dynamics in the net costs of storage and any potential structural change. Any seasonal variation in the storage costs is in this model captured by the seasonal premium.

maturities up to 1 year and interpolate with swap rates for longer maturities. Data for Libor rates and swap rates are available starting respectively on 1986-01-09 and 2000-07-03. Table A1 in Appendix A provides an overview of the different commodity markets, their sector classification, as well as the year of the first futures contract.

Based on the set of available futures contracts, I create generic futures curves which allows to roll the futures contracts according to the desired rolling scheme. This allows to align the measurement of the signal based on the futures curve (e.g. the basis) with the desired implementation. For the purpose of this paper, futures contracts are rolled one day before the last trade date, defined as the minimum of the first notice date, the last tradeable date and the last delivery date in order to accommodate varying contract specifications across commodity futures.

This implementation differs from the traditional approach followed in the literature which rolls futures contracts on the last day of the previous month. This allows one to capture the dynamics of the scarcity risk up until it materializes. Indeed for contracts expiring close to month-end rolling at the beginning of the month would mean forgoing potentially valuable information about season specific expectations of demand and supply imbalances.

Some commodities exhibit a non-regular contract cycle such that there might not be an outstanding contract for each season (e.g. there is only five monthly contracts for wheat futures with delivery in March, May, July, September and December). To complement the set of available contracts, I create synthetic contracts to increase the breadth of the strategy. Upon data availability, these contracts are constructed by simple linear interpolation between a near and a far contract to obtain the desired maturity.

The futures exchange can issue two types of contract. On the one hand, most listed futures are issued for a subset of regular delivery months and constitute the contract cycle. On the other hand, the exchange can also issue serial contracts, which are "off" cycle. While the liquidity might be limited on serial contracts, it is difficult to identify those contracts through time as the choice by the exchanges of maturities up for issuance changes through time. Instead, to control for illiquidity, we handle the problem at the core and clean the data for stale pricing and impose a minimum number of pricing observations (set arbitrarily to 10).

3.1.2 Inventories

Physical inventories data for 27 commodities are collected from multiple exchanges, including the London Metal Exchange (LME), the Commodity Exchange (COMEX), the Intercontinental Exchange (ICE) and the New-York Mercantile Exchange (NYMEX), and governmental agencies, including the National Agricultural Statistics Services from the United States Department of Agriculture (NASS-USDA) and the United States Department of Energy (DOE). No inventory information is available for Brent Crude Oil. Table A2 in Appendix A provides an overview of the inventory data for each commodity market in our universe. It details the source, the starting year of the data, the reporting frequency and the reporting lag applied to the series.

3.2 Basis decomposition

3.2.1 Estimation methodology

In the absence of reliable spot data and following the literature, the basis is measured between every consecutive contract along the futures curve starting from the first active contract and is expressed in percent of the front contract price adjusted for the number of days until delivery in order to allow for a fair comparison across contracts with different maturities. Note that the term premium is defined from the second nearest contract and beyond this maturity.

The seasonal convenience yield is estimated by regression of a detrended log-basis in excess of the foregone interest on seasonal dummies. Note that given the non-linearities observed in the basis, I use a robust regression methodology and the median is preferred to the mean for detrending.

$$y_t^{(n)} - rf_t^{(n)} - u_t'^{(n)} = \sum_{i=1}^m \beta_{t,i} d_i + \epsilon_t$$

The restriction imposed on the unconditional mean of the convenience yield is enforced by setting the net storage cost $u_t'^{(n)}$ equal to the median of the log-basis in excess of the foregone interests $\tilde{v}^{(n)}$.

$$\begin{aligned} u_t'^{(n)} &= \tilde{v}^{(n)} \\ v_t^{(n)} &= y_t^{(n)} - rf_t^{(n)} \end{aligned}$$

The long-run cost-of-carry $\Upsilon_t^{(n)}$ is thus simplified to incorporate this restriction while the two components of the convenience yield, i.e. the seasonal premium and the residual convenience yield in excess of the seasonal component, are defined as follows.

$$\begin{aligned} \Upsilon_t^{(n)} &= rf_t^{(n)} + \tilde{v}^{(n)} \\ \zeta_t^{(n)} &= -\sum_{i=1}^m \beta_{t,i} d_i \\ \chi_t^{(n)} &= -\epsilon_t \end{aligned}$$

I use a panel approach along the futures curve, i.e. considering all available contracts at any point in time, and carry out the estimation on an expanding window to allow for the slow adjustment of expectations about the seasonal premium. The long-run net storage costs are estimated on the same measurement window as the seasonal convenience yield.

3.2.2 Full-sample estimates

Tables B1, B2, B3 and B4 in the Appendix show the full-sample estimates of the seasonal premium over the long-run cost-of-carry per commodity. The estimation has been carried out by means of robust regression using a panel approach along the futures curve, i.e. considering all available contracts at any point in time, using seasonal dummies. The results are reported using Newey-West standard errors, i.e. heteroskedasticity and

autocorrelation consistent estimate of the covariance matrix of the coefficient estimates.

Those results provide interesting insights on the heterogeneity of commodities with regards to the seasonal risk premium and the influence of time-variation in supply and demand. First, looking at the adjusted R^2 of the regression allows to characterize the influence of the seasonal factor as a driver of the futures basis. We see that highly seasonal commodities like natural gas (NG) have a R^2 close to 0.6 while for aluminum this number is barely different from zero. The more pronounced the seasonal variations are, i.e. the amplitude of the premium across seasons, the higher the adjusted R^2 of the regression as they are large contributors to the variance of the basis. This is thus a useful indicator to classify commodities as seasonal or non-seasonal commodities.

Second, the estimated premia per season are usually highly statistically significant, i.e. at the 1% confidence level, for seasonal commodities. Non-seasonal commodities like platinum (PL) can also have specific seasons with a statistically significant premium. While the alternative and equivalent trigonometric approach⁶ to estimating seasonal components is convenient to understand the components of the seasonal cycle, the chosen simple dummy regression approach suffices in characterizing the premium attached to futures contract maturing in specific seasons.

3.2.3 Futures curve decomposition

The basis measured over various investment horizons n fully describes the shape of the futures curve. Alternatively measuring the basis between every subsequent futures along the curve delivers an equivalent representation.⁷ Given the above basis decomposition, we can assess how the time to maturity influences each component independently along the curve. For illustration purpose, Figure 2 shows the log-basis decomposition of the Natural Gas futures curve on January 5, 2016 when measured from consecutive contracts.

⁶See, for example, Hannan et al. (1970), Sørensen (2002), Borovkova and Geman (2006) or Hevia et al. (2018) for more details on the trigonometric specification of seasonality. The results indicate the need to tailor the inclusion of the harmonics beyond the fundamental frequency to differentiate between commodities.

⁷Such a representation is conceptually similar to the notion of forward rates along fixed income yield curves.

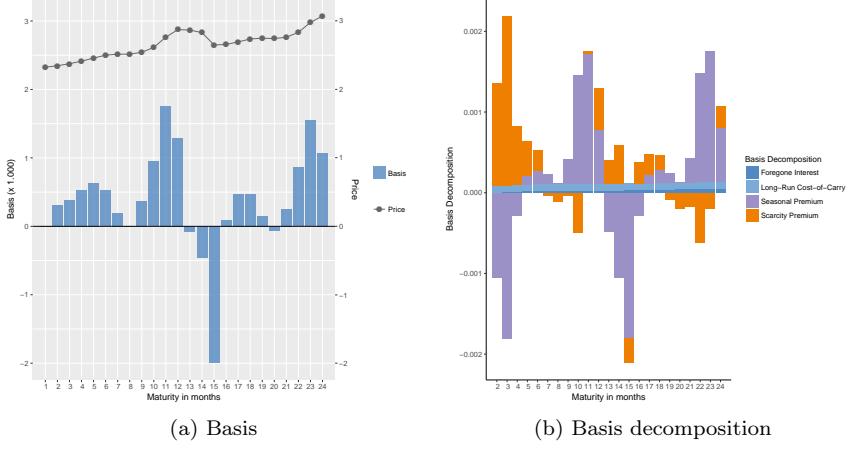


Figure 2: Natural Gas futures basis decomposition along the curve

The curve associated with the long-run cost-of-carry $\Upsilon_t^{(n)}$ captures the convexity or concavity of respectively backwardated and contangoed commodities resulting from the variation in the long-run net storage costs and in interest rates across maturities. Thus the contribution to the log-basis of the foregone interests reflects the shape of the interest rate curve. The contribution of the long-run net storage costs is assumed constant along the futures curve and seasonal variation in storage costs resulting transitory supply and demand imbalances are captured by the seasonal component. While the long-run storage costs curve is flat on any given day its level can vary over time. Figure 3 shows the contribution to the log-basis decomposition of the Natural Gas futures curve on January 5, 2016 of both the foregone interest and the long-run net storage costs.

The seasonal premium curve describes how the cyclicalities in demand and supply command a premium $\zeta_t^{(n)}$ that varies according to the specific season the maturity n is associated with. This generates the observed oscillation along the futures curve for seasonal commodities.

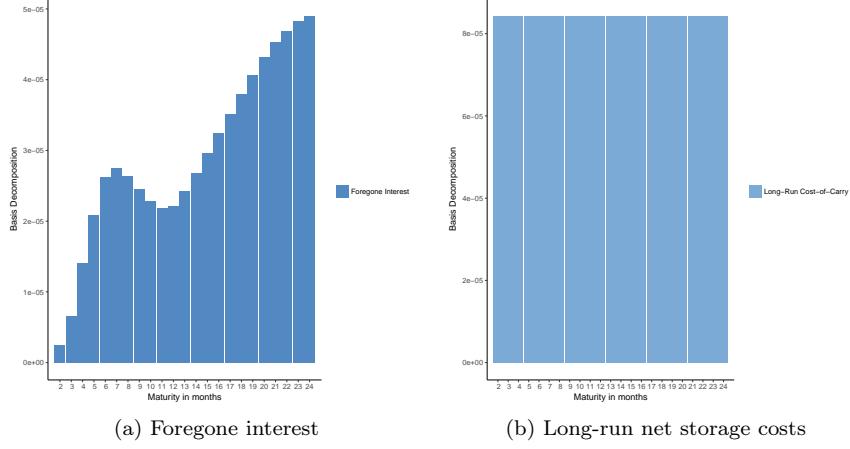


Figure 3: Long-run cost-of-carry along the curve

Next we can define the curve originating from the scarcity premium $\chi_t^{(n)}$. It carries information about how market participants are pricing recent shocks to supply and demand. Indeed the curve describes the priced persistence of those shocks by market participants, i.e. the extent to which they are expecting those shocks to be transitory in nature. Then, the scarcity premium curve also describes the pace at which investors expect those shocks to correct, i.e. the curve captures the expected resolution speed of supply and demand imbalances through time. Figure 4 shows the contribution to the log-basis decomposition of the Natural Gas futures curve on January 5, 2016 of both the seasonal and scarcity premia.⁸

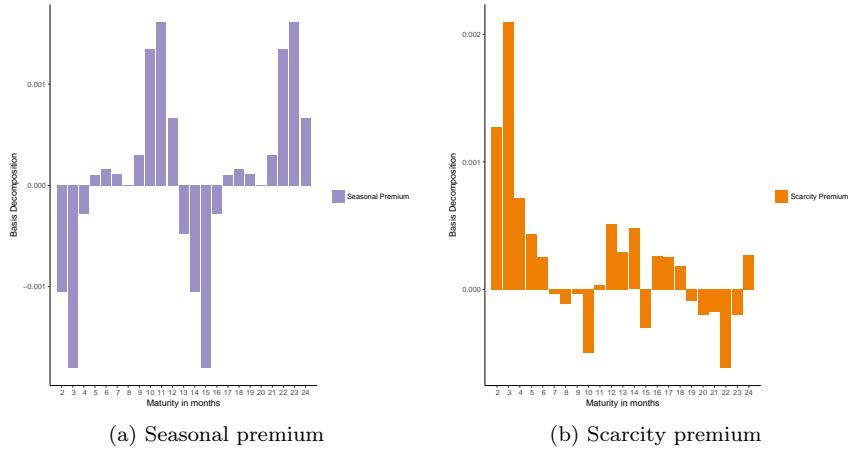


Figure 4: Seasonal and scarcity premia along the curve

⁸ As the additive decomposition of the basis is visually more appealing, the contributions to the basis in Figure 4 actually represent the negative of the seasonality and scarcity premia. See Section 2.2 for the basis decomposition and a clarification of the premia's signs.

3.3 Inventory dynamics

To analyze the relationship between the expectations embedded in the basis and the inventories dynamics, I follow Gorton et al. (2013) and define the "normal" inventory levels as the Hodrick-Prescott (HP) filtered trend of log inventories.⁹ The deviation from this long-run inventory level corresponds to the cyclical component of the inventory dynamics. Gorton et al. (2013) find strong empirical evidence of the presence of seasonality in the cyclical inventory levels. Following the same residualization approach performed on the basis in excess of the foregone interest and the long-run costs of storage, described in Section 2.2, I proceed to a seasonal adjustment of the cyclical inventory levels, above its normalized long-run level, to disentangle seasonal inventory fluctuations from abnormal variations.

The measurement of those excess variations is sensitive to the noise associated with the estimation of long-run normalized inventory level. In a recent paper, Hamilton (2017) discusses issues associated with the use of the HP filter and proposes an alternative filtering approach. As the debate on the most suitable methodology to define the cyclical component of economic timeseries is beyond the scope of this paper, I consider two model-free variables, the 1 and 3 months changes in inventory levels, as proxy for the short-term inventory dynamics. I also perform a similar seasonal adjustment on those additional measures which also display large seasonal variations and can equivalently capture abnormal supply and demand shocks.

3.4 Portfolio formation

In order to analyze the pricing of expectations along the futures curve and the impact of inventory dynamics, I create quintile sorting portfolios for each basis factor. Quintiles are created for each generic futures series, i.e. one for the cross-section of front contracts, one for all the nearest contracts maturing after the active contract and so forth. While the literature usually focuses on the front contract, the use of all available contracts along the futures curve increases substantially the sample size. This leads to an increase in power for testing the statistical significance of factor exposures. I follow the traditional portfolio formation methodology with portfolios being rebalanced at month's end. Positions are equally-weighted and portfolios do not contain synthetic contracts. I require a minimum breadth of 5 futures contract in order to create quintile portfolios for a generic futures series.

3.5 Multiple testing bias

The risk of p-hacking has received lately increasing attention in the academic literature given the bias to publish positive results and the risk of multiple hypothesis testing.¹⁰ To address data mining considerations I

⁹As in Gorton et al. (2013) the smoothness parameter is set at 160,000, retaining peak-to-peak cycles of over 30 years, to reflect the slow adjustments of the inventories' determinants, i.e. production and storage capabilities.

¹⁰See amongst others Harvey and Liu (2015), Harvey et al. (2016), Harvey (2017), Hou et al. (2015), Hou et al. (2017), Chordia et al. (2017), and Linnainmaa and Roberts (2018).

follow the recommendation of Harvey et al. (2016) to raise the statistical significance threshold and reject the null hypothesis for t-statistics above 3.0.

4 Fundamental expectations

So far I have suggested that the basis decomposition proposed in Section 2.2 allows to disentangle unconditional expectations about the seasonal impact on the basis, resulting from known seasonality in inventories, from conditional expectations originating from abnormal shocks to inventories that characterize the scarcity risk. Without any information on the future state of supply and demand, market participants form priors based on the information available to them at any point in time, i.e the current filtration, and as new information arrives agents adjust their conditional expectations away from their original priors to account for the anticipated marginal change in supply and demand imbalance.

In this section, I investigate the fundamental nature of the expectations built-in the convenience yield. More specifically, I test the associated hypotheses that the seasonal premium is driven by the seasonality of inventories while the scarcity premium originates from the excess inventory dynamics. For a robust evaluation of those relationships, I conduct panel regression and a sorting portfolio exposure analysis using three proxies of the inventory dynamics.

4.1 Panel regressions

To analyze the non-linear relationship between the basis and inventories Gorton et al. (2013) estimate a cubic spline regression of the basis on normalized inventory levels, i.e. detrended using a Hodrick-Prescott (HP) filter to capture the lon-run inventory level, and account for seasonal variations in the basis using dummy variables. While the authors provide strong evidence of this non-linear relationship, their model specification does not allow to dissociate the effect of prior expectations, originating from known seasonal variations in supply and demand, from marginal adjustment resulting from fundamental abnormal shocks. Indeed, the normalized inventory level is not seasonally-adjusted in their approach although they find strong empirical evidence of the presence of seasonality in inventories.

To specifically analyze the influence of inventory seasonality and abnormal supply and demand shocks, I conduct panel regressions of the basis, as well as for each of its components, on those inventory dynamics. As in the vast majority of cases, the Hausman (1978) test supports the use of random effects, they are systematically controlled for in the estimation. The analysis is conducted on month-end data for the front contract and the results are robust to the use of daily data. The t-statistics rely on the heteroskedasticity and autocorrelation consistent Newey and West (1987) standard errors estimates, using the Newey and West (1994) automatic lag selection procedure. The results for the univariate regressions on the level of the three inventory dynamics proxies are reported in Appendix C.

Table 1 presents the results when the estimation is conducted on the cyclical inventory level components. We see that the basis is significantly

exposed to the abnormal cyclical inventory level in excess of its seasonal component. This exposure originates from the scarcity premium which exhibits a similar exposure and thus confirms the hypothesis.

Table 1: Regression basis factors on inventory cycle components.

	Basis	ForegoneInt	Median	Season	Scarcity
Intercept	0.0001 <i>t</i> = 3.130***	0.00004 <i>t</i> = 20.108***	-0.00004 <i>t</i> = -2.132**	0.00002 <i>t</i> = 1.165	0.0001 <i>t</i> = 3.143***
Inv Cycle Season	-0.0003 <i>t</i> = -1.466	-0.00001 <i>t</i> = -3.225***	0.00001 <i>t</i> = 2.549**	-0.0004 <i>t</i> = -2.339**	0.0001 <i>t</i> = 0.421
Inv Cycle Excess	0.0003 <i>t</i> = 4.235***	-0.00001 <i>t</i> = -1.065	0.00001 <i>t</i> = 2.075**	0.00004 <i>t</i> = 2.843***	0.0003 <i>t</i> = 4.298***
Adjusted R ²	0.026	0.002	0.005	0.026	0.027

Note: The t-statistics rely on the heteroskedasticity and autocorrelation consistent Newey-West (1987) standard errors estimates, using the Newey-West (1994) automatic lag selection procedure. *p<0.1; **p<0.05; ***p<0.01

Interestingly, the seasonal premium has a negative exposure to the seasonal component of the cyclical inventory level. The difficulty associated with the interpretation of those results is that the cyclical inventory seasonal component is a level metric while the basis is a change measure, here defined as the difference between the futures contract and the spot prices. As such while a positive seasonal inventory component is expected to lead to a relatively lower spot price and higher basis all else equal, implying a positive relationship, this measure abstracts from the expected seasonal component associated with the futures contract and thus affects the basis. If the futures curve captures the seasonality of inventories, prices along the curve should be a reflection of the seasonal inventory level, implying in a negative relationship. The positive relationship of the seasonal premium with the excess inventory level is more difficult to interpret. Indeed, it could suggest that the excess and seasonals component of the cyclical inventory level are correlated although they should be orthogonal by design. I confirmed ex-post that two components are independant with a close to null correlation. In light of the ten times lower coefficient for the excess inventory level, I conclude this is a second order effect.

Table 2: Regression basis factors on 3 months inventory changes.

	Basis	ForegoneInt	Median	Season	Scarcity
Intercept	0.0001 <i>t</i> = 3.158***	0.00004 <i>t</i> = 19.971***	-0.00004 <i>t</i> = -2.128**	0.00002 <i>t</i> = 1.048	0.0001 <i>t</i> = 3.133***
Inv Chg 3M Season	0.004 <i>t</i> = 1.533	-0.00001 <i>t</i> = -0.414	-0.00001 <i>t</i> = -0.317	0.002 <i>t</i> = 1.420	0.002 <i>t</i> = 1.545
Inv Chg 3M Excess	0.001 <i>t</i> = 1.762*	0.00001 <i>t</i> = 0.763	-0.0001 <i>t</i> = -2.072**	0.0001 <i>t</i> = 0.337	0.001 <i>t</i> = 2.351**
Adjusted R ²	0.013	-0.0002	0.002	0.011	0.007

Note: The t-statistics rely on the heteroskedasticity and autocorrelation consistent Newey-West (1987) standard errors estimates, using the Newey-West (1994) automatic lag selection procedure. *p<0.1; **p<0.05; ***p<0.01

Table 2 and Table 3 present the results when the estimation is conducted on the components of respectively the 3 months and 1 month change in inventories. In both tables we can see a similar pattern emerging with the basis and the scarcity premium having a statistically significant exposure to the excess inventory changes, while the seasonal premium is loading, although not significantly, on the seasonal component of those changes.

Table 3: Regression basis factors on 1 month inventory changes.

	Basis	ForegoneInt	Median	Season	Scarcity
Intercept	0.0001 <i>t</i> = 3.143***	0.00004 <i>t</i> = 19.969***	-0.00004 <i>t</i> = -2.128**	0.00002 <i>t</i> = 1.109	0.0001 <i>t</i> = 3.123***
Inv Chg 1M Season	0.009 <i>t</i> = 1.596	0.00000 <i>t</i> = 0.017	-0.00004 <i>t</i> = -0.600	0.007 <i>t</i> = 1.740*	0.002 <i>t</i> = 0.890
Inv Chg 1M Excess	0.001 <i>t</i> = 1.816*	0.00000 <i>t</i> = 0.260	-0.0001 <i>t</i> = -1.822*	0.0002 <i>t</i> = 0.630	0.001 <i>t</i> = 2.645***
Adjusted R ²	0.012	-0.0001	0.0002	0.027	0.001

Note: The t-statistics rely on the heteroskedasticity and autocorrelation consistent Newey-West (1987) standard errors estimates, using the Newey-West (1994) automatic lag selection procedure. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

4.2 Sorting portfolio exposures

To provide further insights on the influence of inventory dynamics on the basis and its components, this section investigates, for each basis component, the sorting portfolios exposure to inventory characteristics. Table 4 provides information on the results of this non-parametric approach. The average exposure to all three inventory dynamics proxies, in terms of level and its seasonal and abnormal components, is reported for each quintile portfolio. For each basis factor, the last row contains the t-statistic for the difference in mean exposure between the top and bottom portfolios. The t-statistics above 3.0 are shown in bold and the asterisk identifies the characteristics for which the quintiles also display a monotonic behavior.

We observe for the basis portfolios significant differences in exposures to the excess cyclical inventory level and the 1 month inventory change between the top and bottom portfolios. The seasonal premium portfolios have highly significant exposure to the level and seasonal component of both the 1 month and 3 months inventory changes. The scarcity premium has significant exposure to all level and excess components of all three proxies for inventory dynamics.

Those results add to the literature. First, they bring additional empirical evidence on the relationship between the futures basis and the inventory dynamics in the cross-section of commodity markets. The second and more relevant contribution is to provide robust support to the claim that the basis embeds expectations about two different inventory dynamics. On one hand, the seasonal premium captures expectations about the cross-sectional dispersion in inventory seasonalities. On the other hand, the scarcity premium reflects the excess supply and demand imbalances over the expected seasonal fluctuations.

Table 4: Basis Factors Quintile Portfolios - Inventory dynamics exposures

Quantile	Cyclical inventory level			Inventory changes 3 months			Inventory changes 1 month		
	Level	Season	Excess	Level	Season	Excess	Level	Season	Excess
Basis									
1	-3.83	1.02	-4.33	-0.46	0.27	-1.08	-0.18	0.38	-0.97
2	-2.55	1.03	-3.11	-0.32	0.18	-0.55	-0.43	0.06	-0.62
3	2.91	1.30	1.84	0.60	0.40	0.10	0.39	0.20	0.13
4	6.48	1.07	5.24	0.48	0.27	0.07	0.31	0.19	0.07
5	6.54	1.69	4.91	0.85	0.18	0.03	0.67	0.28	0.35
t-statistic (5-1)	2.88	0.50	3.28	1.21	1.44	2.64	0.44	2.97	3.61
Foregone Interest									
1	4.25	2.12	2.23	0.07	0.49	-0.43	0.05	0.29	-0.20
2	5.55	1.62	3.82	0.47	0.54	-0.25	0.36	0.33	-0.16
3	1.74	1.03	0.81	0.48	0.43	-0.37	0.55	0.55	-0.20
4	-1.09	0.57	-1.32	0.02	0.03	-0.08	-0.05	-0.06	-0.01
5	-0.52	0.74	-0.49	0.17	0.00	0.17	-0.17	-0.07	-0.08
t-statistic (5-1)	0.38	1.34	0.50	0.34	1.38	0.50	0.44	3.62	0.29
Long Run Cost Of Carry									
1	1.73	0.17	1.30	0.62	0.29	0.23	0.69	0.31	0.36
2	1.69	1.35	0.72	0.25	0.24	-0.02	0.15	0.00	0.01
3	0.06	1.41	-1.04	0.25	0.48	-0.29	0.11	0.27	-0.19
4	2.38	1.28	1.65	0.17	0.27	-0.49	0.11	0.26	-0.55
5	4.81	1.65	3.55	0.01	0.00	-0.56	0.05	-0.01	-0.27
t-statistic (5-1)	2.55	2.09	1.83	2.21	1.30	2.67	2.38	1.57	2.06
Season Premium									
1	2.19	1.30	1.13	-0.15	0.08	-0.70	-0.20	0.11	-0.49
2	4.75	1.33	3.64	0.15	0.17	-0.17	-0.09	0.07	-0.20
3	2.04	1.26	0.97	0.04	0.34	-0.31	-0.03	0.15	-0.21
4	0.11	0.79	-0.35	0.77	0.34	0.34	0.61	0.18	0.30
5	0.62	1.18	-0.13	0.90	0.46	-0.09	0.92	0.72	-0.06
t-statistic (5-1)	0.44	0.59	1.02	3.96	4.28	0.30	3.13*	5.19	0.32
Scarcity Premium									
1	5.09	1.05	3.96	1.39	0.38	0.40	1.34	0.55	0.67
2	7.89	1.09	6.69	0.89	0.38	0.36	0.61	0.19	0.35
3	4.24	1.27	3.10	0.57	0.32	0.13	0.28	0.15	0.11
4	-3.47	1.24	-4.24	-0.57	0.20	-0.90	-0.53	0.13	-0.90
5	-5.34	1.34	-5.99	-1.36	0.13	-1.71	-0.94	0.17	-1.61
t-statistic (5-1)	6.16	1.53	6.32	4.72*	1.23	5.97*	3.76*	0.78	5.26*

Note: The cyclical inventory level corresponds to the difference in percentage between the inventory level and the HP-filtered inventory trend. For each inventory dynamic, its Level is decomposed in a Season component, which captures the seasonal average, and an Excess component, capturing the abnormal inventory dynamics. The t-Stat (5-1) tests for the difference in mean exposure between the top and bottom quintile for each characteristic. The t-statistics above 3.0 are shown in bold and the asterisk identifies the characteristics for which the quintiles also display a monotonic behavior. All figures are expressed in percentages. Inventory changes are reported as annualized figures.

5 Priced expectations

In this section, I attempt to address the question of whether all expectations carry predictive power and, if not, to identify which expectations are priced-in the commodity futures curves. First, the influence of the various factors driving the basis on the futures spot and term risk premia is analyzed. Then the traditional sorting portfolio approach is used to document the risk and return profile of the basis factors' portfolios and investigate the statistical significance of any basis factor premia.

5.1 Spot risk premium

The spot risk premium is defined in Appendix D as the expected return of holding the nearest one-period futures contract until maturity. Table 5 sheds light on the spot premium. Following Bakshi et al. (2017), I define the market as an equally-weighted portfolio of all commodities, using

the front contracts. Acknowledging the relevance of seasons, I investigate the premium conditional on seasons. As the influence of seasonality varies across commodities and sectors, the results are also presented for all sectors. While the spot premium is significantly different from zero for the overall market, it is only significant for the industrial metals sector over the full sample. We can observe large seasonal variations in the spot premium both at the market and at the sector levels.

Table 5: Spot Risk Premium

Sector	All	Season											
		1	2	3	4	5	6	7	8	9	10	11	12
Annualized Mean Returns (%)													
Market	4.16	7.74	21.32	7.86	14.17	5.86	3.01	10.34	5.98	4.30	-7.82	-0.74	5.74
Energy	0.62	-19.33	11.19	32.06	42.55	35.11	17.75	-5.57	-25.88	-7.52	-3.75	-1.86	-42.60
Grains	1.32	11.17		8.15		-1.63		5.99	-23.27	1.55	-46.85	-4.23	-2.51
Industrial metals	9.18	21.02	35.85	23.43	10.03	11.67	-16.95	-3.45	11.97	-5.13	9.80	-1.42	-7.17
Livestock	1.96	-0.90	5.77	5.05	-8.49	13.47	-1.63	12.49	12.00	6.28	0.71	-17.78	-2.67
Precious metals	8.01	0.68	44.24	42.80	-0.75	4.49	-1.42	-32.41	16.03	9.17	6.36	1.62	6.63
Softs	1.43	-0.89		9.67		-5.65		3.33		7.56	-4.41	26.92	3.10
Annualized Standard Deviations (%)													
Market	16.33	20.81	19.09	20.04	17.47	17.22	19.48	17.55	14.95	18.35	18.80	20.88	20.07
Energy	31.31	36.29	35.48	38.07	34.27	30.63	28.23	29.36	28.59	28.92	32.84	34.83	32.72
Grains	20.61	20.69		18.93		20.86		20.73	26.87	26.75	20.11	25.58	22.40
Industrial metals	23.36	25.58	25.76	24.24	23.02	23.44	24.45	25.05	23.37	22.84	25.00	29.59	26.37
Livestock	13.87	14.91	17.47	13.12	15.58	16.66	16.08	19.22	13.66	11.25	17.92	12.74	20.55
Precious metals	23.50	24.35	21.11	20.43	23.05	25.11	23.92	23.04	21.20	21.43	28.70	24.52	24.50
Softs	19.46	30.95		22.83		19.70		19.20		23.27	23.83	34.23	24.66
t-Statistics													
Market	1.26	0.80	2.03	1.33	1.58	0.76	0.28	1.26	0.88	0.53	-1.08	-0.07	0.77
Energy	0.08	-0.79	0.47	1.24	1.89	1.74	0.96	-0.29	-1.39	-0.40	-0.17	-0.08	-1.90
Grains	0.29	0.90		1.01		-0.13		0.50	-1.05	0.10	-2.69	-0.28	-0.26
Industrial metals	1.65	1.20	2.07	1.39	0.64	0.73	-1.03	-0.20	0.76	-0.33	0.56	-0.07	-0.39
Livestock	0.60	-0.11	0.59	0.65	-1.05	1.11	-0.17	0.79	1.94	0.65	0.08	-1.34	-0.23
Precious metals	1.40	0.03	2.53	2.41	-0.04	0.22	-0.07	-1.68	0.91	0.51	0.25	0.08	0.32
Softs	0.36	-0.05		1.35		-0.56		0.34		0.64	-0.43	1.32	0.31

Note: Following Harvey et al. (2016), the statistical significance threshold for t-statistics is raised to 3.0 in order to control for the risk of multiple hypothesis testing. The tests for which the null hypothesis, of returns not being significantly different from zero, is rejected are identified in bold.

Next, I investigate the forecasting ability of the futures basis with respect to the spot risk premium. While the literature mainly focuses on the evaluation of the unbiased forward hypothesis, the objective is here to gain understanding in which factors within the basis carry predictive power. In order to better reflect the nature of the different premia driving the basis, especially the non-linearity of the scarcity premium, I focus here on daily holding returns. This should allow to capture more effectively the influence of the different factors driving the futures spot and term premia.

Table 6 reports the results of forecasting regressions of the forward excess returns relative to the market on the basis, its constituents and other control variables (e.g. sectors and seasons). Consistent with Bakshi et al. (2017) I find that the market is a major risk factor and can explain a large portion (about 25%) of the timeseries variation in the future spot risk premium. I thus control for the market as a risk factor by focusing on excess returns and keep the market as an independent variable to evaluate the efficacy of the adjustment. Results shown below are robust to this choice of specification. The t-statistics rely on the heteroskedasticity and autocorrelation consistent Newey and West (1987) standard errors estimates, using the Newey and West (1994) automatic lag selection procedure.

The first regression corresponds to the usual projection of forward returns onto the futures basis. Consistent with previous literature results, a high basis (contangoed curve) leads to negative forward returns. This re-

Table 6: Regression excess spot premium above market on basis factors

	Basis	Basis Factors	Control Sectors	Control Seasons
Constant	0.0001 <i>t</i> = 1.088	0.00003 <i>t</i> = 0.548		
MarketEW	0.00003 <i>t</i> = 0.001	0.0001 <i>t</i> = 0.003	0.0001 <i>t</i> = 0.010	-0.0001 <i>t</i> = -0.009
Basis	-0.353 <i>t</i> = -2.720***			
ForegoneInterest		-0.109 <i>t</i> = -0.124	0.035 <i>t</i> = 0.039	-0.149 <i>t</i> = -0.173
LongRunCostOfCarry		-0.912 <i>t</i> = -2.058**	-0.961 <i>t</i> = -1.649*	-0.898 <i>t</i> = -1.849*
SeasonalPremium		-0.225 <i>t</i> = -1.311	-0.280 <i>t</i> = -1.574	-0.223 <i>t</i> = -1.269
ScarcityPremium		0.572 <i>t</i> = 3.324***	0.563 <i>t</i> = 3.658***	0.592 <i>t</i> = 4.136***
Sector_Energy			-0.00001 <i>t</i> = -0.045	
Sector_Grains			0.0001 <i>t</i> = 0.545	
Sector_IndustrialMetals			0.0002 <i>t</i> = 1.808*	
Sector_Livestock			-0.0002 <i>t</i> = -1.187	
Sector_PreciousMetals			0.0002 <i>t</i> = 1.491	
Sector_Softs			-0.0001 <i>t</i> = -1.143	
Season_1				-0.00004 <i>t</i> = -0.323
Season_2				0.00000 <i>t</i> = 0.016
Season_3				0.0002 <i>t</i> = 0.923
Season_4				0.0001 <i>t</i> = 1.001
Season_5				0.0004 <i>t</i> = 2.017**
Season_6				-0.0001 <i>t</i> = -0.770
Season_7				-0.0001 <i>t</i> = -0.343
Season_8				0.00002 <i>t</i> = 0.160
Season_9				-0.0001 <i>t</i> = -0.665
Season_10				0.0002 <i>t</i> = 1.180
Season_11				-0.0001 <i>t</i> = -0.355
Season_12				0.00000 <i>t</i> = 0.007
Adjusted R ²	0.0002	0.0004	0.0004	0.0003
Residual Std. Error	0.015 (df = 119193)	0.015 (df = 119190)	0.015 (df = 119185)	0.015 (df = 119179)

Note: The *t*-statistics rely on the heteroskedasticity and autocorrelation consistent Newey-West (1987) standard errors estimates, using the Newey-West (1994) automatic lag selection procedure. **p*<0.1; ***p*<0.05; ****p*<0.01

lation is highly significant but does not pass the requirements put forward in Harvey et al. (2016). The second regression focuses on the different basis factors as explanatory variables for forwards excess returns. We see that the first factor, the foregone interest, is not statistically significant and has little influence but the sign of the coefficient is as per expectation, i.e. futures contracts trade at a premium relative to spot prices to compensate for the interests foregone by physical commodity holders. As time passes the futures contract is expected to converge to the spot price if nothing else changes.

The second factor, the long-run cost-of-carry which captures the storage costs net of any structural convenience yield, shows a large negative beta that is statistically significant at the traditional 5% level. The sign is consistent with the traditional basis sign and the expectation that commodity futures with high storage costs trade at a premium and converge to spot overtime.

The third basis factor, the seasonality premium, although not significant deserves a short comment. Indeed an interesting observation is that the seasonal premium coefficient has an opposite sign to the scarcity premium although both are components of the convenience yield. This suggests both insurance premia have different roles and relates to separate hedging demands. In order to interpret the negative coefficient of the seasonal premium let's recall from Section 2.2 that here a high value actually indicates that the futures trades at a discount. A high seasonal premium thus relates to a period of expected oversupply and the negative sign of the coefficient suggests that during those periods the futures price actually move beyond what is priced in. Hedging seasonal risk is a fruitful strategy for risk averse investors. This also suggests that this factor does not simply capture the seasonal variation in the costs of storage in which case the coefficient would be of opposite sign to the interest rate and long-run cost-of-carry factors.

Finally, the last factor, the scarcity premium, is highly significant with a t-statistic well above 3.0. One interpretation for the positive coefficient is that when market participants are pricing a high risk of scarcity, the scarcity premium pushes down futures prices which leads to positive forward returns. A probably more accurate explanation is that a high scarcity premium is actually driven by the front end of the curve and that the slow diffusion of information combined to the slow adjustment of supply and demand leads to a curve shift in the direction of the scarcity risk and thus to futures returns predictability. Further insights into the dynamics of the futures curve in presence of scarcity risk are provided in Section 6.2.

The next two regressions introduce commodity sectors and seasons as control variables but overall those dummy variables are not significant. More importantly their introduction weakens the significance of the long-run cost-of-carry factor suggesting our initial finding might not be robust and that there is a large cross-sectional dispersion across sectors (the factor is constant through seasons). The result for the scarcity premium on the other hand becomes more significant, suggesting those findings are robust to the model specification.

5.2 Term risk premium

The term risk premium is defined in Appendix D as the one-period expected holding return of a n -period futures contract in excess of the spot risk premium. Here all contracts with a maturity up to one year out from the front month contract are considered. This allows to capture all seasons at any point in time as well as to diversify the exposure to the forward rate along the curve. Table 7 sheds light on the term premia across sectors and through seasons. The results are similar to the one obtained for the spot premium. The term premium is only significant for the industrial metals and energy sectors over the full sample. We can observe large seasonal variations in the term premium both at the market and at the sector levels.

Table 7: Term Risk Premium

Sector	All	Season											
		1	2	3	4	5	6	7	8	9	10	11	12
Annualized Mean Returns (%)													
Market	0.93	0.34	2.28	0.55	2.94	2.23	3.43	2.10	2.93	0.23	2.48	1.18	2.23
Energy	5.23	3.23	4.06	2.99	5.47	6.22	6.19	6.52	6.73	5.68	5.88	4.96	5.58
Grains	1.38	-0.62		1.96		1.57		2.36	0.52	1.17	-0.86	-0.72	3.38
Industrial metals	1.39	1.15	1.34	1.72	1.74	1.25	1.68	1.49	1.49	1.29	0.90	1.48	1.16
Livestock	2.95	-3.24	4.92	-2.82	3.48	2.68	6.59	7.92	4.46	-0.14	0.78	-1.83	4.85
Precious metals	0.07	0.04	-0.01	0.32	0.11	0.43	-0.03	0.32	-0.12	0.11	-0.05	-0.05	0.18
Softs	0.35	-4.53		0.99		2.71		0.90		-1.59	3.32	-5.42	1.18
Annualized Standard Deviations (%)													
Market	4.29	3.65	4.14	4.63	3.90	4.05	4.21	4.57	3.52	4.40	5.43	4.11	4.63
Energy	9.95	11.37	10.98	10.25	10.58	10.54	10.77	10.81	11.01	10.79	11.11	10.89	11.09
Grains	5.00	7.66		5.42		5.34		5.71	6.44	5.80	7.84	9.84	6.09
Industrial metals	2.33	2.58	2.58	2.61	2.62	2.64	2.64	2.59	2.63	2.66	2.74	2.64	2.60
Livestock	8.82	6.93	13.60	7.04	10.22	13.88	13.66	18.13	12.57	6.22	11.41	6.73	14.73
Precious metals	0.54	0.98	0.69	1.54	0.69	1.45	0.66	0.98	0.73	1.08	0.67	0.71	0.74
Softs	5.68	11.63		6.21		5.57		6.08		6.00	10.57	12.52	6.40
t-Statistics													
Market	1.30	0.53	3.10	0.71	4.24	3.23	4.62	2.69	4.67	0.30	2.71	1.62	2.85
Energy	3.04	1.52	1.98	1.56	2.76	3.16	3.08	3.23	3.27	2.82	2.83	2.44	2.70
Grains	1.64	-0.44		2.06		1.64		2.31	0.43	1.16	-0.58	-0.40	3.21
Industrial metals	3.14	2.20	2.51	3.22	3.26	2.32	3.10	2.83	2.77	2.38	1.61	2.75	2.19
Livestock	1.87	-1.64	1.78	-1.62	1.65	0.86	2.38	2.10	1.75	-0.09	0.34	-1.11	1.62
Precious metals	0.80	0.22	-0.11	1.18	0.89	1.71	-0.27	1.84	-0.93	0.58	-0.37	-0.08	1.40
Softs	0.36	-2.08		0.95		2.83		0.86		-1.52	1.85	-2.27	1.09

Note: Following Harvey et al. (2016), the statistical significance threshold for t-statistics is raised to 3.0 in order to control for the risk of multiple hypothesis testing. The tests for which the null hypothesis, of returns not being significantly different from zero, is rejected are identified in bold.

Following the approach put forward when investigating the unbiased forward hypothesis, I investigate the forecasting ability of the relative futures basis with respect to the term risk premium. Table 8 reports the results of forecasting regressions of the forward relative returns on the basis differential, the difference in the basis factors and other control variables (e.g. sectors and seasons). It is worth noting that the long-run cost-of-carry factor drops off the set of independent variables across the various regressions. This follows from the basis decomposition proposed in Section 2.2 and the restriction imposed on this factor to be constant across maturities at any point in time.

Table 8: Regression term premium on basis factors differential

	Basis	Basis Factors	Control Sectors	Control Seasons
Constant	0.0001 <i>t</i> = 5.256***	0.0001 <i>t</i> = 5.378***		
MarketEW	-0.003 <i>t</i> = -1.872*	-0.004 <i>t</i> = -2.140**	-0.004 <i>t</i> = -1.993**	-0.004 <i>t</i> = -2.025**
BasisDiff	-0.184 <i>t</i> = -3.659***			
ForegoneInterestDiff		-4.504 <i>t</i> = -2.716***	-4.400 <i>t</i> = -2.204**	-4.571 <i>t</i> = -2.128**
SeasonPremiumDiff		-0.067 <i>t</i> = -1.214	-0.067 <i>t</i> = -1.138	-0.072 <i>t</i> = -1.169
ScarcityPremiumDiff		0.377 <i>t</i> = 6.462***	0.362 <i>t</i> = 4.862***	0.379 <i>t</i> = 4.783***
Sector_Energy			0.0002 <i>t</i> = 2.956***	
Sector_Grains			0.0001 <i>t</i> = 2.492**	
Sector_IndustrialMetals			0.0001 <i>t</i> = 3.682***	
Sector_Livestock			0.0001 <i>t</i> = 1.672*	
Sector_PreciousMetals			0.00003 <i>t</i> = 2.416**	
Sector_Softs			0.00004 <i>t</i> = 1.133	
Season_1				0.0001 <i>t</i> = 3.393***
Season_2				0.0001 <i>t</i> = 4.072***
Season_3				0.0001 <i>t</i> = 4.499***
Season_4				0.0001 <i>t</i> = 4.231***
Season_5				0.0001 <i>t</i> = 6.287***
Season_6				0.0001 <i>t</i> = 3.825***
Season_7				0.0001 <i>t</i> = 5.050***
Season_8				0.0001 <i>t</i> = 4.872***
Season_9				0.0001 <i>t</i> = 3.130***
Season_10				0.00005 <i>t</i> = 1.820*
Season_11				0.00004 <i>t</i> = 1.649*
Season_12				0.0001 <i>t</i> = 3.235***
Adjusted R ²	0.001	0.001	0.002	0.002
Residual Std. Error	0.005 (df = 942219)	0.005 (df = 942217)	0.005 (df = 942212)	0.005 (df = 942206)

Note: The t-statistics rely on the heteroskedasticity and autocorrelation consistent Newey-West (1987) standard errors estimates, using the Newey-West (1994) automatic lag selection procedure. *p<0.1; **p<0.05; ***p<0.01

The first specification regresses the forward futures return relative to the front contract return against the market and the basis differential. The basis differential is highly significant and in line with expectations a high spread predicts negative relative returns. Although only significant at the 10% threshold it is worth highlighting the influence of the market factor on the term premium whereby a rise in commodity markets leads to a negative term premium. To some extent this result is a reflection of the correlation structure and relative risk along the futures curve, i.e. contracts further along the curve exhibit lower beta to spot changes. This effect is consistent across the various regressions performed. Supportive evidence is provided by Table 9 which shows the panel full-sample correlation estimates of the first twelve futures contract along the curve with the active contract as well as the volatilities of the contracts as we move along the curve (e.g. column 1 is the active front contract while column 3 refers to the third contract along the curve).

Table 9: Risk and correlation along the futures curve

	1	2	3	4	5	6	7	8	9	10	11	12
Correlation	1.00	0.99	0.98	0.97	0.96	0.95	0.94	0.94	0.93	0.93	0.92	0.91
Volatility (%)	29.69	28.69	27.46	25.77	25.00	24.49	24.09	23.81	23.56	23.64	23.72	23.91

The second regression focuses on the various basis factors' differentials as explanatory variables of the forward relative returns. The first factor differential, the foregone interest one, is found to be significant at the 1% level. The sign of the coefficient is in line with expectations and the previous findings on the spot risk premium. The second factor differential, the seasonal premium one, is here as well not significant but the result is consistent with the findings on the spot premium. This provides comfort in the interpretation put forward above. The last factor differential, the scarcity premium one, is highly significant with a t-statistic well above 3. Here as well the results are consistent with earlier findings.

The next two regressions introduce commodity sectors and seasons as control variables. As opposed to the results on the spot risk premium, here the control variables are often significant at the 5% level or lower. This hints at a large cross-sectional dispersion across sectors. With regards to seasons while these results suggest the term premium remains exposed to seasonal fluctuations it is worth mentioning this control variable only captures the season exposure of the far contract. The result for the scarcity premium are left unchanged suggesting those findings are robust to the model specification.

5.3 Factor portfolios

For all basis factors, Table 10 sheds light on the risk, return and t-statistic of quintiles portfolios implemented on the front contracts, together with those of the top-minus-bottom portfolio.

This table carries a few insightful results. First and foremost, the results confirm earlier literature findings that the futures basis carries information about future returns. The t-statistics for the bottom quintile portfolio, which equally-weights the commodity markets with the most negative basis (backwardated curves) is highly significant with a t-statistic of 3.26. It outperforms the top-minus-bottom portfolio which exhibits a

slightly lower level of significance at -2.57 as the top portfolio with the most contangoed curves also earns a barely positive return.

The second key finding, which comes as a confirmation of the results presented just above, is that the scarcity risk premium distilled from cost-of-carry relationship is the sole risk premium embedded in the basis. None of the other basis components are informative suggesting the related expectations are priced-in the futures curve and do not carry predictive power. Note that both the seasonal and scarcity premium are of opposite sign relative to the basis as a high scarcity premium leads to a highly negative basis. As a result, the top quintiles for those two premiums relate to the bottom quintiles for the basis and its other components. The top quintile portfolio has a t-statistic of 3.10 and drives the top-minus-bottom portfolio which is more informative than the basis one with a t-statistic of 2.83 as the bottom scarcity portfolio earns a negative premium.

Table 10: Basis Factors Quintile Portfolios

Statistics	Quintile Portfolios					
	1	2	3	4	5	5 - 1
Basis						
Return	20.68	12.34	7.47	5.12	0.89	-19.64
Volatility	26.53	23.89	20.96	25.76	29.73	32.04
t-statistic	3.26	2.16	1.49	0.83	0.13	-2.57
Foregone Interest						
Return	7.74	10.28	1.81	7.30	12.53	4.66
Volatility	22.41	26.70	23.46	25.25	21.59	23.29
t-statistic	1.44	1.59	0.32	1.19	2.41	0.84
Long Run Cost Of Carry						
Return	3.19	14.98	10.37	9.19	2.08	-1.12
Volatility	30.97	26.58	19.64	23.15	24.40	31.72
t-statistic	0.43	2.36	2.21	1.66	0.36	-0.15
Season Premium						
Return	7.13	9.13	12.30	6.87	3.05	-4.09
Volatility	29.15	24.34	18.65	24.58	29.60	30.45
t-statistic	1.03	1.57	2.77	1.17	0.43	-0.56
Scarcity Premium						
Return	-3.93	-1.86	8.83	16.07	18.65	22.39
Volatility	30.21	27.12	20.86	23.09	25.13	33.24
t-statistic	-0.55	-0.29	1.78	2.91	3.10	2.83

Note: Following Harvey et al. (2016), the statistical significance threshold for t-statistics is raised to 3.0 in order to control for the risk of multiple hypothesis testing. The tests for which the null hypothesis, of returns not being significantly different from zero, is rejected are identified in bold.

All in all, the results presented in this section add to the asset pricing literature and contribute to the understanding of the risk premia and their drivers in the cross-section of commodity markets. This paper brings novel insights on the pricing of expectations within the futures basis and provides strong empirical evidence that unconditional expectations about the seasonal inventory dynamics are priced-in. On the contrary, the conditional expectations about the impact of abnormal supply and demand shocks earn a premium and carry predictive power for future returns.

6 Return predictability

The scarcity risk premium has been identified above as the key driver beyond the futures spot and term premia and as the sole risk premium

embedded in the futures basis but little is known on the origin of this return predictability. The efficient market hypothesis¹¹ suggests that in the presence of shocks to supply and demand market participants adjust their expectations and incorporate instantaneously this new information in prices such that return predictability is precluded. Under such hypothesis the scarcity risk premium should reflect the incorporation of new information and the resulting expectations into the futures curve but should not be able to forecast forward returns.

At the same time commodity markets are characterized by the slow adjustment of demand and supply which leads to shock persistence. On the one hand, price inelastic demand originates from the absence of substitutes in the short run. On the other hand, the supply of commodities is subject to the inelasticity of production in the short run driven by seasonal production cycles for perishable commodities and the long run effects of investments and productivity increase on the total supply. Those two effects combined result in a lengthy resolution of supply and demand imbalances.

Two potential competing but not mutually exclusive hypotheses could explain the observed return predictability. The first one would be the presence of autocorrelation in unexpected net supply shocks such that selecting commodities on the basis of previous shocks would provide valuable information about future expected returns. The second hypothesis would be that the slow diffusion of information and/or market participants underreaction to new information lead to return predictability.¹²

6.1 Abnormal shocks serial correlation

In order to asses the first hypothesis of recurring shocks, I consider in Figure 5 the average autocorrelation coefficients up to one year across all combinations of market and season for the level and the change in the scarcity risk premium as it captures unexpected shocks to the net supply. While the dynamics of the level coefficients corroborates the persistence of shocks and the slow adjustment of fundamentals, the autocorrelation coefficients for the changes in the scarcity risk premium cannot attest of the presence of a systematic recurrence of shocks through time. Indeed, the absence of large and significant autocorrelation coefficient for the change of the scarcity risk premium suggests unexpected shocks to supply and demand are not autocorrelated and I thus dismiss this first hypothesis.

¹¹See amongst others Malkiel and Fama (1970).

¹²Both behavioral phenomena have been widely documented in equity markets. See amongst others the seminal paper of Jegadeesh and Titman (1993) for the underreaction of market participants. See amongst others Hong and Stein (1999) and Hong et al. (2000) for the slow diffusion of information. Dissociating those two potential sources of return predictability is beyond the scope of this paper.

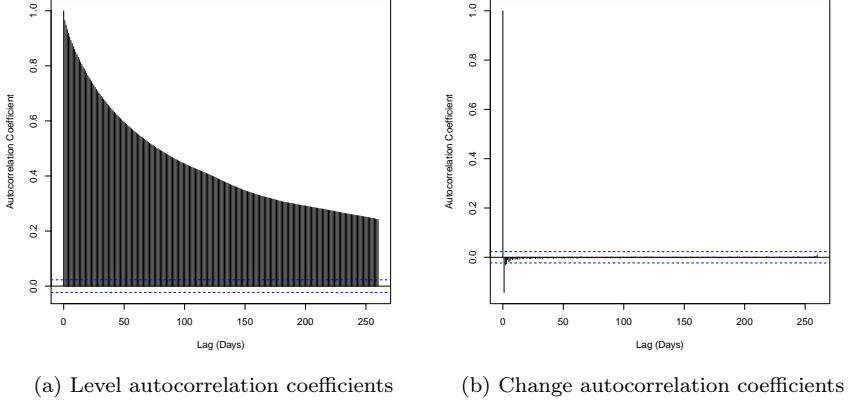


Figure 5: Scarcity risk premium dynamics

6.2 Underreaction and slow diffusion of information

The second hypothesis of slow diffusion of information and market participants underreaction leverages on a growing body of literature whereby behavioral and cognitive biases¹³ as well as limits to arbitrage¹⁴ impair market efficiency. To corroborate this alternative hypothesis, I investigate the futures curve dynamics affecting the basis around the emergence of a scarcity premium.

Table 11 sheds light on those dynamics by conducting an event analysis. It considers the risk and return profile of the near and far contracts, i.e. both legs involved in the measurement of the scarcity risk premium, coincident on the identification of this abnormal convenience yield premium in excess of the seasonal premium, and, on the following day, for all contracts along the curve. The first two columns display the contemporaneous average performance of the near and far contracts, conditional on the scarcity premium being positive or negative. We see that a positive (negative) scarcity premium is accompanied with a more positive (negative) return of the front contract relative to the far contract. The next columns report forward information (return, volatility and t-statistic) for the near and far contracts. Following the emergence of a positive scarcity premium, i.e. a risk of stock-out, we see that the futures curve continues to shift upward and flattens. Likewise the formation of a negative scarcity premium, i.e. a risk of oversupply, leads to a continuation of the downward shift and a flattening of the futures curve. This continuation in the parallel shift of the futures curve, in the direction of the scarcity premium, thus confirms the underreaction and slow diffusion of information hypothesis.

It worth noticing that the forward returns of the near and far contracts are statistically significant when the scarcity premium is positive and that those results are consistent across seasons and sectors. When the

¹³See amongst others Hirshleifer (2001).

¹⁴See amongst others Shleifer and Vishny (1997).

scarcity premium is negative the forward returns are not statistically significant and results are not consistent through sectors or seasons. These results provides evidence that market participants underreact to the risk of scarcity and suggest the scarcity premium is truly a compensation for the risk of stock-out.

Table 11: Curve Dynamics

<i>Information</i>		<i>Contemporaneous</i>			<i>Forward</i>			
Contract	Near	Far	Near			Far		
Sector/Season	Return	Return	Return	Volatility	t-statistic	Return	Volatility	t-statistic
Positive scarcity premium								
Average	17.50	14.29	5.92	12.22	2.89	7.98	12.05	3.95
Energy	17.96	14.71	7.33	22.24	1.85	9.45	21.47	2.47
Grains	12.02	8.20	4.05	18.59	1.29	7.13	18.29	2.30
Industrial metals	23.25	19.60	10.74	22.44	2.22	12.62	22.08	2.65
Livestock	10.99	4.95	3.10	12.47	1.37	6.49	12.05	2.96
Precious metals	3.20	1.68	0.23	23.73	0.05	1.71	23.69	0.39
Softs	15.34	10.76	-0.05	18.06	-0.01	2.58	17.76	0.86
1	18.05	15.14	5.88	14.41	2.38	8.02	14.24	3.28
2	19.42	16.65	7.08	14.27	2.88	8.72	14.05	3.60
3	18.43	15.67	8.61	13.93	3.63	10.42	13.79	4.44
4	21.88	18.55	6.16	13.82	2.58	8.72	13.52	3.73
5	19.47	16.22	7.19	13.91	3.01	9.68	13.70	4.12
6	18.19	14.07	4.86	13.74	2.05	7.65	13.58	3.27
7	16.82	13.96	7.15	14.31	2.92	9.11	14.16	3.76
8	20.29	16.54	6.17	13.76	2.60	8.69	13.46	3.74
9	16.62	12.95	6.22	13.39	2.71	8.41	13.20	3.71
10	16.28	13.19	3.77	13.35	1.65	5.21	13.13	2.32
11	14.67	12.44	3.53	13.13	1.57	5.27	12.96	2.37
12	12.59	10.60	3.60	13.10	1.61	5.22	13.01	2.36
Negative scarcity premium								
Average	-10.18	-7.54	-0.72	12.53	-0.34	-2.13	12.35	-1.03
Energy	-1.53	1.73	8.07	25.19	1.79	6.18	24.31	1.42
Grains	-6.36	-3.03	-0.93	18.25	-0.30	-2.88	17.94	-0.95
Industrial metals	-3.34	-1.62	1.81	21.82	0.42	0.92	21.64	0.21
Livestock	-2.73	3.54	4.28	13.06	1.79	0.67	12.33	0.30
Precious metals	-4.61	-2.94	-0.09	21.28	-0.02	-1.78	21.22	-0.50
Softs	-13.88	-11.15	-1.31	15.58	-0.50	-2.78	15.23	-1.09
1	-13.35	-10.75	-2.63	13.09	-1.16	-4.10	12.88	-1.84
2	-11.78	-9.54	-2.64	13.41	-1.15	-3.57	13.23	-1.58
3	-8.98	-6.74	-2.29	13.46	-1.00	-3.29	13.30	-1.46
4	-11.24	-8.27	0.74	13.14	0.33	-0.83	12.93	-0.37
5	-9.25	-6.24	0.81	13.55	0.35	-0.73	13.39	-0.32
6	-9.91	-6.68	-0.07	14.23	-0.03	-2.07	14.13	-0.86
7	-10.36	-7.85	-2.15	14.10	-0.88	-3.52	13.97	-1.46
8	-11.81	-8.89	0.49	13.95	0.21	-1.48	13.67	-0.63
9	-9.69	-6.78	-0.64	13.89	-0.27	-2.48	13.68	-1.06
10	-10.82	-8.60	1.18	14.52	0.48	-0.25	14.37	-0.10
11	-9.01	-6.83	0.96	14.10	0.40	-0.27	13.88	-0.12
12	-10.01	-7.89	-1.72	13.98	-0.72	-2.80	13.84	-1.19

Note: Following Harvey et al. (2016), the statistical significance threshold for t-statistics is raised to 3.0 in order to control for the risk of multiple hypothesis testing. The tests for which the null hypothesis, of returns not being significantly different from zero, is rejected are identified in bold.

Another interesting result relates to the curve twisting, i.e. flattening or steepening, originating from a more pronounced underreaction in the far contract. This implies that the underestimation of the probability of

stock-out is larger in the seasons neighboring the potential materialization of the event. In equity markets the slow diffusion of information has been put forward to explain underreaction and the momentum factor.¹⁵ Assuming market participants can trade the whole futures curve, this theory could explain a parallel shift in the futures curve but not the dynamics of local underreaction around the event risk. An alternative interpretation is that market participants are subject to a framing bias whereby they are not able to extrapolate the implications of the risk of stock-out. This leads hedging and liquidity demand to be primarily concentrated in the seasonal contract where the event is located and to the mispricing of both the risk that the inventory depletion happens at faster rate than anticipated as well as the risk that inventory imbalances might resorb at a lower speed than expected due to the slow adjustment of supply. As the risk of scarcity rises further, investors would recalibrate their probabilities and reassess the risks surrounding the event.

6.3 Information decay

To further investigate the second hypothesis of slow diffusion of information and market participants underreaction, I consider the persistence of excess returns relative to the market over multiple holding horizons and along the curve for various distance to maturity for both the top and bottom quintile scarcity risk portfolios. Figure 6 shows a heatmap of the t-statistics of the excess returns over multiple horizons (from 1 day to 12 months) for the different contracts along the curve, where the contract number indicates its location on the curve and increases with the distance to maturity. The colors allows to classify the t-statistics in various categories depending on its level (e.g. following Harvey et al. (2016) blue and red correspond to t-statistics respectively above 3 and below -3) while the colour shading allows to distinguish for various t-statistic levels within a category.

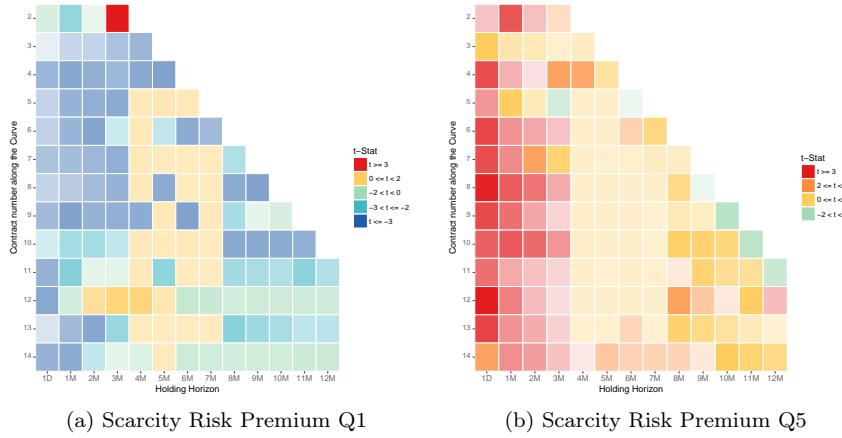


Figure 6: Excess performance persistence

¹⁵See amongst others Hong and Stein (1999) and Hong et al. (2000)

Beyond allowing to detect the presence of excess returns' persistence, associated with forecasting power, this representation carries valuable information. First, by conditioning on the quintile number it allows to distinguish the direction of the excess return predictability between high and low scarcity risk. Second, by evaluating multiple holding horizons, this analysis allows to assess the decay in the excess return over time and thus to define the return predictability horizon. Third, by considering the whole futures curve, we can identify whether the excess returns are uniformly distributed along the curve or whether they are concentrated in a specific location (e.g. the front end) of the curve.

We see from Figure 6a that an unexpected risk of over-supply, i.e. a negative scarcity risk premium, leads to negative highly significant excess returns over the market for a holding horizon up to 3 months. This pattern is mostly observable in the front end of the curve, for contract with a time to maturity up to 9 months. Figure 6b shows that unexpected risks of scarcity, i.e. a positive premium is associated with highly significant positive excess returns over the market across the whole curve for a holding horizon up to 3 months.

In both figures we observe that the excess returns are transient and most significant right after the unexpected shocks, suggesting the information is largely incorporated by market participants without delay. Their diminishing significance with the length of the holding horizon indicates that the information is progressively diffused throughout the market up until it is fully incorporated in futures prices. More importantly, their short persistence through time, i.e. a high alpha decay, seems inconsistent with the expected timelines of a fundamentally driven resolution of supply and demand imbalances.

The results presented in this section bring convincing evidence that the scarcity risk premium is a compensation for the risk of stock-out. It finds its origin in the slow diffusion of information and underreaction to abnormal changes in inventories, above and beyond seasonal dynamics. The underreaction is broad based along the curve and is more pronounced in the far contracts, leading to curve twist and suggesting market participants also underreact to the risk that the inventory depletion happens at faster rate than anticipated as well as the risk that inventory imbalances might resorb at a lower speed than expected due to the slow adjustment of supply. While most of the information is incorporated shortly after the materialization of the risk of stock-out, the return predictability horizon is about 3 months which is long enough to support the underreaction and slow diffusion of information hypothesis and too short to be associated with a fundamental resolution of abnormal supply and demand imbalances.

Conclusion

The cost-of-carry relationship is the corner stone of the theory of storage. It assumes the difference between the futures and spot prices, i.e. the futures basis, can be explained by the interest foregone when buying commodities in the physical market, the associated storage costs and a convenience yield. The latter is defined as a benefit that accrues to the

commodity holder resulting from the potential productive value of his inventory. This value of physical ownership reflects market expectations about the future availability of a commodity.

A vast literature provides evidence that the futures basis is informative about futures risk premia. This paper sheds light on the origin of this return predictability by disentangling various expectations embedded in the basis. As production and demand seasonality can have a large influence on the future state of inventories, a convenience yield decomposition is proposed to dissociate known seasonal supply and demand imbalances from abnormal shocks above and beyond those transitory dynamics. Agents form expectations about the impact of inventories' seasonality on the basis (henceforth the seasonality risk premium) and adjust their conditional expectations as new information arrives. Changes in expectations resulting from unanticipated supply and demand shocks (henceforth the scarcity risk premium) are associated with the incremental risk of scarcity and capture the non-linearities in the futures basis.

The contribution of this paper to the asset pricing literature is three-fold. First, it brings novel insights on the fundamental relationship between the futures basis and inventory dynamics. While the literature has documented the relationship between the basis and the level of inventories, as well as the presence of seasonality in both the basis and the inventories, the question of how various inventory dynamics are priced in the cross-section of commodity markets remains an open research topic. This paper investigates whether the predictable seasonality in the convenience yield is related to known seasonal fluctuations in inventories, and whether the residual, or seasonally-adjusted, convenience yield is associated with unexpected supply and demand shocks. It provides strong empirical evidence that the convenience yield embeds expectations about those two different inventory dynamics. On one hand, the seasonal risk premium captures expectations about the inventory seasonalities. On the other hand, the scarcity risk premium reflects the excess supply and demand imbalances over the expected seasonal fluctuations.

Second, this paper contributes to the understanding of the risk premia in the cross-section of commodity markets by investigating the pricing of expectations within the futures basis. It provides strong empirical evidence that the seasonality risk premium is priced-in as unconditional expectations about the seasonal inventory dynamics carry no predictive power. On the contrary, the conditional expectations about the impact of abnormal supply and demand shocks earn a premium and carry predictive power for future returns. In fine, the scarcity risk premium carries all the predictive power embedded in the futures basis. Other basis components, i.e. the foregone interest and the net cost of storage, have no information content. Those findings are robust to the introduction of commodity sectors and seasons as control variables.

The third contribution of this paper is to show that the scarcity risk premium is a compensation for the unexpected increase in the risk of stock-out and that the associated return predictability finds its origin in the slow diffusion of information and underreaction to abnormal changes in inventories, above and beyond seasonal dynamics. It presents strong evidence of underreaction along the futures curve following a marginal change in the risk of stock-out. The underreaction is broad based along the curve but more pronounced in the far contracts, leading to futures curve twists. These results suggest that market participants also under-

react to the risk that the inventory depletion happens at faster rate than anticipated, as well as the risk that inventory imbalances might resorb at a lower speed than expected due to the slow adjustment of supply. Despite the slow diffusion of information and the unexpected nature of the supply and demand shocks, information decays fast and the return predictability is transient. The diminishing statistical significance of excess returns with the length of the holding horizon indicates that the information is progressively diffused throughout the market up until it is fully incorporated in futures prices. The scarcity risk premium has significant predictive power up to a 3 months horizon, which is too short to be associated with a fundamental resolution of abnormal supply and demand imbalances.

As the probability of stock-out increases, the net hedging demand rises. On one hand, commodity consumers need to hedge out the scarcity risk by buying spot and storing the commodity or by taking long futures positions with delivery at the time of likely scarcity. On the other hand, producers would benefit from unwinding their hedges or taking outright long commodity exposures when inventory levels are low and eroding.¹⁶ At the same time, other informed market participants, e.g. speculators, know about the predictive power embedded in the basis. The question of who might be a willing counterparty to the hedging and speculative demand in the presence of scarcity risk thus deserves further research.

Another avenue of research relates to the pricing of expectations along the futures curve. On one hand, this paper provides empirical evidence that the seasonality risk premium is priced-in and carries no predictive power. On the other hand, Keloharju et al. (2016) find strong evidence of the presence of returns seasonality in U.S. and international equities as well as in commodities. Although the authors find that for the later the difference between the high and low seasonal portfolios is only marginally statistically significant but economically important, the questions of what the drivers of the return seasonality factor are and how it relates to the seasonality risk premium defined in this paper warrant additional research.

¹⁶Hirshleifer (1991) refutes the general pre-conception of a short hedging pressure resulting from producers simply willing to hedge their bottom line volatility. He argues that when demand shocks are arising from changes in aggregate wealth, the optimal hedging policy of producers with non-stochastic output might be to remain unhedged or even to take a long hedge position.

Appendices

Appendix A: Data description

Table A1 presents the universe of commodity futures contracts considered in this paper and provides information concerning the starting year of the sample data, the source, the contract code letter for the available contract delivery months and the classification in commodity sectors.

Table A1: Futures contracts

Contract Code	Description	Sector	Exchange	Start Year	Months	Source
BO	Soybean Oil	Grains	CBOT	1961	FGHJKMNQUVZ	CME
C	Corn	Grains	CBOT	1960	HKNUZ	CME
CC	Cocoa	Softs	ICE	1970	HKNUZ	ICE
CL	WTI Crude Oil	Energy	NYMEX	1983	FGHJKMNQUVXZ	CME
CO	Brent Crude Oil	Energy	ICE	1993	FGHJKMNQUVXZ	ICE
CT	Cotton	Softs	ICE	1972	HKNVZ	ICE
FC	Feeder Cattle	Livestock	CME	1974	FGHJKMNQUVX	CME
GC	Gold	Precious metals	COMEX	1975	GJMQVZ	CME
HG	Copper	Industrial metals	COMEX	1989	FGHJKMNQUVXZ	CME
HO	Heating Oil	Energy	NYMEX	1986	FGHJKMNQUVXZ	CME
HU	Unleaded Gazoline	Energy	NYMEX	1987	FGHJKMNQUVXZ	CME
JO	Orange Juice	Softs	ICE	1967	FHKNUX	ICE
KC	Coffee C	Softs	ICE	1973	HKNUZ	ICE
KW	KC HRW Wheat	Grains	CBOT	1976	HKNUZ	CME
LA	Aluminium	Industrial metals	LME	1997	FGHJKMNQUVXZ	Datostream
LC	Live Cattle	Livestock	CME	1965	GJMQVZ	CME
LH	Lean Hogs	Livestock	CME	1987	GJMNQVZ	CME
LL	Lead	Industrial metals	LME	1998	FGHJKMNQUVXZ	Datostream
LN	Nickel	Industrial metals	LME	1997	FGHJKMNQUVXZ	Datostream
LT	Tin	Industrial metals	LME	1998	FGHJKMNQUVXZ	Datostream
LX	Zinc	Industrial metals	LME	1997	FGHJKMNQUVXZ	Datostream
NG	Natural Gaz	Energy	NYMEX	1990	FGHJKMNQUVXZ	CME
PL	Platinum	Industrial metals	NYMEX	1987	FJNV	CME
S	Soybean	Grains	CBOT	1970	FHKNUX	CME
SB	Sugar No. 11	Softs	ICE	1964	HKNV	ICE
SI	Silver	Precious metals	COMEX	1975	FHKNUZ	CME
SM	Soybean Meal	Grains	CBOT	1964	FHKNUVZ	CME
W	Wheat	Grains	CBOT	1959	HKNUZ	CME
XB	Gasoline RBOB	Energy	NYMEX	2006	FGHJKMNQUVXZ	CME

Table A2 provides information on the inventory dataset. More specifically, it specifies the starting year of the sample data, the source, the corresponding futures contract code, the reporting frequency and the reporting lag applied to the timeseries. For a more detailed description of the inventory data, I refer to the Appendix B of Gorton et al. (2013).

Table A2: Commodity inventories

Contract Code	Description	Source	Frequency	Start Year	Lag
BO	Soybean Oil	USDA	Monthly	1964	1 Month
C	Corn	USDA	Monthly	1960	1 Month
CC	Cocoa	ICE	Daily	2012	1 Day
CL	WTI Crude Oil	DOE	Weekly	1982	1 Week
CO	Brent Crude Oil				
CT	Cotton	ICE	Daily	2002	1 Day
FC	Feeder Cattle	USDA	Monthly	1994	1 Month
GC	Gold	COMEX	Daily	1992	1 Day
HG	Copper	COMEX	Daily	1992	1 Day
HO	Heating Oil	DOE	Weekly	1993	1 Week
HU	Unleaded Gasoline	DOE	Weekly	2004	1 Week
JO	Orange Juice	USDA	Monthly	1964	1 Month
KC	Coffee C	ICE	Daily	2010	1 Day
KW	KC HRW Wheat	CBOT	Weekly	2008	1 Week
LA	Aluminum	LME	Daily	1978	1 Day
LC	Live Cattle	USDA	Monthly	1917	1 Month
LH	Lean Hogs	USDA	Monthly	1917	1 Month
LL	Lead	LME	Daily	1970	1 Day
LN	Nickel	LME	Daily	1979	1 Day
LT	Tin	LME	Daily	1970	1 Day
LX	Zinc	LME	Daily	1970	1 Day
NG	Natural Gaz	DOE	Monthly	1973	1 Month
PL	Platinum	NYMEX	Daily	1995	1 Day
S	Soybean	USDA	Monthly	1964	1 Month
SB	Sugar No. 11	USDA	Monthly	1996	1 Month
SI	Silver	COMEX	Daily	1992	1 Day
SM	Soybean Meal	USDA	Monthly	1964	1 Month
W	Wheat	USDA	Monthly	1960	1 Month
XB	Gasoline RBOB	DOE	Weekly	2004	1 Week

Appendix B: Estimation of the seasonal premium

Tables B1, B2, B3 and B4 show the full-sample estimates of the seasonal premium over the long-run cost-of-carry per commodity. The estimation has been carried out by means of robust regression using a panel approach along the futures curve, i.e. considering all available contracts at any point in time, using seasonal dummies. The results are reported using Newey-West standard errors, i.e. heteroskedasticity and autocorrelation consistent estimate of the covariance matrix of the coefficient estimates.

Table B1: Regression of the convenience yield on seasonal dummies

	BO	C	CC	CL	CT	CO	FC
Season 1	0.021 <i>t</i> = 3.579***	0.107 <i>t</i> = 7.964***	-0.001 <i>t</i> = -0.327	0.025 <i>t</i> = 6.306***	0.047 <i>t</i> = 7.183***	0.051 <i>t</i> = 6.659***	-0.037 <i>t</i> = -3.002***
Season 2				0.001 <i>t</i> = 0.186			0.004 <i>t</i> = 0.368
Season 3	-0.010 <i>t</i> = -1.423	-0.055 <i>t</i> = -3.842***		-0.001 <i>t</i> = -0.197		-0.0001 <i>t</i> = -0.005	-0.005 <i>t</i> = -0.279
Season 4				-0.001 <i>t</i> = -0.117			0.074*** <i>t</i> = 4.770***
Season 5	-0.014 <i>t</i> = -1.900*	-0.058 <i>t</i> = -4.124***	0.007 <i>t</i> = 1.260	0.001 <i>t</i> = 0.058	-0.006 <i>t</i> = -0.662	0.003 <i>t</i> = 0.247	0.087*** <i>t</i> = 4.953***
Season 6				-0.002 <i>t</i> = -0.179		0.003 <i>t</i> = 0.232	
Season 7	-0.011 <i>t</i> = -1.541	-0.084 <i>t</i> = -5.869***	0.005 <i>t</i> = 0.842	-0.001 <i>t</i> = -0.091	-0.022 <i>t</i> = -2.562***	0.002 <i>t</i> = 0.203	
Season 8	-0.048 <i>t</i> = -4.863***			0.007 <i>t</i> = 0.876		0.003 <i>t</i> = 0.239	0.129*** <i>t</i> = 0.029***
Season 9	-0.051 <i>t</i> = -5.054***	-0.267 <i>t</i> = -7.202***	0.005 <i>t</i> = 0.967	-0.0005 <i>t</i> = -0.069		0.002 <i>t</i> = 0.144	0.045 <i>t</i> = 3.039***
Season 10	-0.079 <i>t</i> = -5.051***			0.005 <i>t</i> = 0.692	-0.135 <i>t</i> = -4.338***	0.0002 <i>t</i> = 0.020	0.046 <i>t</i> = 3.119***
Season 11	-0.008 <i>t</i> = -0.379			0.003 <i>t</i> = 0.446		0.004 <i>t</i> = 0.350	0.064*** <i>t</i> = 2.769***
Season 12	-0.024 <i>t</i> = -2.567***	-0.153 <i>t</i> = -7.295***	-0.003 <i>t</i> = -0.428	0.003 <i>t</i> = 0.453	-0.079 <i>t</i> = -4.913***	-0.0001 <i>t</i> = -0.006	
Adjusted R ²	0.075	0.270	0.004				
Residual Std. Error	0.089 (df = 81587)	0.133 (df = 57322)	0.055 (df = 51736)	0.096 (df = 161832)	0.155 (df = 63712)	0.077 (df = 99610)	0.120 (df = 32764)
Note	The t-statistics rely on the heteroskedasticity and autocorrelation consistent Newey-West (1987) standard errors estimates, using the Newey-West (1994) automatic lag selection procedure. *p<0.1; **p<0.05; ***p<0.01						

Table B2: Regression of the convenience yield on seasonal dummies (cont.)

	GC	HG	HO	JO	KC	KW	LA	LC
Season 1	$t = 5.004^{***}$	$t = 2.732^{***}$	$t = 0.014$	$t = 1.781^*$	$t = 0.048$	$t = 0.014$	$t = 0.025$	$t = 0.011$
Season 2	$t = -0.007$	$t = -0.002$	$t = -0.254$	$t = -6.397^{***}$	$t = -0.080$	$t = 3.074^{***}$	$t = 4.887^{***}$	$t = 2.164^{**}$
Season 3	$t = -3.271^{***}$	$t = -0.002$	$t = -0.008$	$t = -10.319$	$t = -0.226$	$t = 0.020$	$t = 0.0002$	$t = -3.429^{***}$
Season 4	$t = -0.008$	$t = 0.001$	$t = -0.008$	$t = -11.934^{***}$	$t = -0.254$	$t = 2.040^{**}$	$t = 0.002$	$t = 0.138$
Season 5	$t = -0.001$	$t = 0.001$	$t = 0.002$	$t = -7.325^{***}$	$t = 0.023$	$t = 2.205^{**}$	$t = 0.006$	$t = 5.358^{***}$
Season 6	$t = -0.008$	$t = 0.005$	$t = -3.674^{***}$	$t = 0.562$	$t = -0.138$	$t = 6.773^{***}$	$t = 0.485$	$t = -0.084$
Season 7	$t = -0.001$	$t = 0.001$	$t = -0.229$	$t = 0.001$	$t = -0.051$	$t = 0.015$	$t = 0.002$	$t = -5.342^{***}$
Season 8	$t = -4.043^{***}$	$t = 0.005$	$t = -4.043^{***}$	$t = 0.577$	$t = 0.003$	$t = 3.419^{***}$	$t = 1.434$	$t = -0.214$
Season 9	$t = 0.001$	$t = -0.001$	$t = 0.390$	$t = -0.069$	$t = 0.039$	$t = 0.003$	$t = 0.015$	$t = -7.325^{***}$
Season 10	$t = -0.008$	$t = 0.002$	$t = -3.610^{***}$	$t = 0.249$	$t = 0.042$	$t = 3.383^{***}$	$t = 0.271$	$t = 2.205^{**}$
Season 11	$t = -0.001$	$t = 0.003$	$t = -0.353$	$t = 0.031$	$t = 0.037$	$t = -2.578^{***}$	$t = 0.001$	$t = -0.031$
Season 12	$t = -0.008$	$t = -0.001$	$t = -3.389^{***}$	$t = -0.064$	$t = 0.035$	$t = 2.876^{***}$	$t = -0.153$	$t = 6.486^{***}$
Adjusted R ²	0.027	0.001	0.062 (df = 75912)	0.132 (df = 133629)	0.404	0.035 (df = 59076)	0.095 (df = 54091)	0.131 (df = 39788)
Residual Std. Error	0.009 (df = 75912)	0.062 (df = 133629)	0.132 (df = 59076)	0.098 (df = 54091)	0.131 (df = 39788)	0.052 (df = 118991)	0.052 (df = 118991)	0.124 (df = 40122)

Note The t-statistics rely on the heteroskedasticity and autocorrelation consistent Newey-West (1987) standard errors estimates, using the Newey-West (1994) automatic lag selection procedure. *p<0.1; **p<0.05; ***p<0.01

Table B3: Regression of the convenience yield on seasonal dummies (cont.)

	LH	LL	LN	LT	LX	NG	PL	S
Season 1	0.318 <i>t</i> = 6.019***	-0.008 <i>t</i> = -1.491	-0.005 <i>t</i> = -0.646	0.001 <i>t</i> = 0.452	-0.009 <i>t</i> = -1.228	0.260 <i>t</i> = 15.613***	0.002 <i>t</i> = 0.571	0.024 <i>t</i> = 12.156***
Season 2	-0.003 <i>t</i> = -0.337	0.001 <i>t</i> = 0.102	0.003 <i>t</i> = 0.660	0.001 <i>t</i> = 0.073	-0.526 <i>t</i> = -6.548***	0.011 <i>t</i> = 2.382***		
Season 3	-0.004 <i>t</i> = -0.051	-0.001 <i>t</i> = -0.058	0.005 <i>t</i> = 1.033	0.001 <i>t</i> = 0.119	-0.728 <i>t</i> = -7.855***	0.010 <i>t</i> = 2.378***	-0.012 <i>t</i> = -2.791***	
Season 4	-0.135 <i>t</i> = -2.070**	0.001 <i>t</i> = 0.130	-0.001 <i>t</i> = -0.134	0.005 <i>t</i> = 1.159	-1.105 <i>t</i> = -13.161***	0.001 <i>t</i> = 0.107	0.001 <i>t</i> = 0.229	
Season 5	0.598 <i>t</i> = 7.607***	0.002 <i>t</i> = 0.301	-0.0001 <i>t</i> = -0.014	0.004 <i>t</i> = 0.912	-0.375 <i>t</i> = -16.475***	0.008 <i>t</i> = 1.626	-0.026 <i>t</i> = -5.343***	
Season 6	0.222 <i>t</i> = 3.818***	0.004 <i>t</i> = 0.482	0.002 <i>t</i> = 0.242	0.004 <i>t</i> = 0.667	0.006 <i>t</i> = 0.577	-0.218 <i>t</i> = -12.081***	0.008 <i>t</i> = 1.984***	
Season 7	-0.353 <i>t</i> = -6.170***	0.006 <i>t</i> = 0.831	0.0004 <i>t</i> = 0.042	0.005 <i>t</i> = 1.065	-0.192 <i>t</i> = -10.802***	0.003 <i>t</i> = 0.250	-0.016 <i>t</i> = 0.137	-0.492***
Season 8	-0.497 <i>t</i> = -8.918***	0.003 <i>t</i> = 0.396	-0.001 <i>t</i> = -0.127	0.004 <i>t</i> = 0.980	-0.220 <i>t</i> = 0.339	0.003 <i>t</i> = -12.654***	-0.112 <i>t</i> = 1.352	-0.798***
Season 9	0.006 <i>t</i> = 0.698	0.003 <i>t</i> = 0.315	0.006 <i>t</i> = 1.391	0.006 <i>t</i> = 0.644	-0.264 <i>t</i> = -15.064***	0.007 <i>t</i> = 0.044	-0.011 <i>t</i> = 2.310*	-0.250 <i>t</i> = -11.052***
Season 10	-0.936 <i>t</i> = -15.317***	0.003 <i>t</i> = 0.369	0.001 <i>t</i> = 0.131	0.003 <i>t</i> = 0.688	0.001 <i>t</i> = 0.074	-0.130 <i>t</i> = -5.571***	-0.001 <i>t</i> = -0.103	
Season 11	-0.001 <i>t</i> = -0.104	-0.001 <i>t</i> = -0.106	0.002 <i>t</i> = 0.485	0.004 <i>t</i> = 0.395	0.271 <i>t</i> = 5.062***	0.008 <i>t</i> = 2.064***	-0.106 <i>t</i> = -8.403***	
Season 12	-0.413 <i>t</i> = -6.712***	-0.001 <i>t</i> = -0.114	0.001 <i>t</i> = 0.083	0.003 <i>t</i> = 0.703	0.004 <i>t</i> = 0.336	0.389 <i>t</i> = 8.914***	0.011 <i>t</i> = 2.151*	
Adjusted R ²	0.676	0.059 (df = 53927)	0.059 (df = 95539)	0.081 (df = 118636)	0.035 (df = 72037)	0.065 (df = 118850)	0.321 (df = 149053)	0.594 (df = 12506)
Residual Std. Error	0.293 (df = 53927)	0.002 (df = 118636)	0.002 (df = 72037)	0.001 (df = 118850)	0.002 (df = 149053)	0.033 (df = 12506)	0.130 (df = 71299)	0.310 (df = 71299)

Note The *t*-statistics rely on the heteroskedasticity and autocorrelation consistent Newey-West (1987) standard errors estimates, using the Newey-West (1994) automatic lag selection procedure. **p*<0.1; ***p*<0.05; ****p*<0.01

Table B4: Regression of the convenience yield on seasonal dummies (cont.)

	SB	SI	SM	W	XBW
Season 1	0.032 <i>t</i> = 0.816	-0.007 <i>t</i> = -4.276***	0.010 <i>t</i> = 2.346**	0.028 <i>t</i> = 4.380***	0.072 <i>t</i> = 8.323***
Season 2		0.016 <i>t</i> = 7.644***			0.062 <i>t</i> = 5.733***
Season 3	0.038 <i>t</i> = 0.929	0.005 <i>t</i> = 1.982*	-0.002 <i>t</i> = -0.276		0.112 <i>t</i> = 9.740***
Season 4		0.014 <i>t</i> = 4.982***			0.854 <i>t</i> = 10.284***
Season 5	-0.075 <i>t</i> = -1.799*	0.003 <i>t</i> = 1.171	-0.021 <i>t</i> = -2.777***	-0.086 <i>t</i> = -3.889***	0.019 <i>t</i> = 1.762*
Season 6		0.060 <i>t</i> = 5.789***			-0.054 <i>t</i> = -4.548***
Season 7	-0.100 <i>t</i> = -2.341**	-0.003 <i>t</i> = -0.956	0.006 <i>t</i> = 0.898	-0.161 <i>t</i> = -4.683***	-0.105 <i>t</i> = -8.125***
Season 8		0.015 <i>t</i> = 5.951***	-0.083 <i>t</i> = -6.851***		-0.131 <i>t</i> = -10.167***
Season 9	-0.048 <i>t</i> = -1.118	0.003 <i>t</i> = 1.385	-0.147 <i>t</i> = -6.334***	0.011 <i>t</i> = 1.391	-0.180 <i>t</i> = -13.003***
Season 10	-0.013 <i>t</i> = -0.324	0.014 <i>t</i> = 6.529***	-0.216 <i>t</i> = -7.522***		-0.702 <i>t</i> = -22.468***
Season 11		0.015 <i>t</i> = 6.547***			-0.191 <i>t</i> = -14.866***
Season 12		0.002 <i>t</i> = 0.982	-0.005 <i>t</i> = -0.811	0.031 <i>t</i> = 4.357***	-0.111 <i>t</i> = -8.429***
Adjusted R ²	0.148 (df = 41501)	0.107 (df = 91352)	0.113 (df = 78655)	0.166 (df = 49092)	0.795 (df = 89852)
Residual Std. Error					0.169 (df = 49092)
Note	The t-statistics rely on the heteroskedasticity and autocorrelation consistent Newey-West (1987) standard errors estimates, using the Newey-West (1994) automatic lag selection procedure. *p<0.1; **p<0.05; ***p<0.01				

Appendix C: Univariate regression of basis factors on inventory dynamics

Table C1: Regression basis factors on inventory cycle.

	Basis	ForegoneInt	Median	Season	Scarcity
Intercept	0.0001 <i>t</i> = 3.132***	0.00004 <i>t</i> = 20.072***	-0.00004 <i>t</i> = -2.132**	0.00002 <i>t</i> = 1.104	0.0001 <i>t</i> = 3.110***
Inv Cycle Level	0.0002 <i>t</i> = 3.677***	-0.00001 <i>t</i> = -1.266	0.00001 <i>t</i> = 2.577***	-0.00002 <i>t</i> = -0.940	0.0002 <i>t</i> = 3.617***
Adjusted R ²	0.013	0.002	0.005	0.0004	0.025

Note: The t-statistics rely on the heteroskedasticity and autocorrelation consistent Newey-West (1987) standard errors estimates, using the Newey-West (1994) automatic lag selection procedure. *p<0.1; **p<0.05; ***p<0.01

Table C2: Regression basis factors on 3 months inventory change.

	Basis	ForegoneInt	Median	Season	Scarcity
Intercept	0.0001 <i>t</i> = 3.144***	0.00004 <i>t</i> = 19.952***	-0.00004 <i>t</i> = -2.127**	0.00002 <i>t</i> = 1.069	0.0001 <i>t</i> = 3.117***
Inv Chg 3M	0.002 <i>t</i> = 1.894*	0.00001 <i>t</i> = 0.523	-0.0001 <i>t</i> = -1.900*	0.001 <i>t</i> = 1.199	0.001 <i>t</i> = 2.429**
Adjusted R ²	0.009	0.0001	0.002	0.003	0.007

Note: The t-statistics rely on the heteroskedasticity and autocorrelation consistent Newey-West (1987) standard errors estimates, using the Newey-West (1994) automatic lag selection procedure. *p<0.1; **p<0.05; ***p<0.01

Table C3: Regression basis factors on 1 month inventory change.

	Basis	ForegoneInt	Median	Season	Scarcity
Intercept	0.0001 <i>t</i> = 3.153***	0.00004 <i>t</i> = 20.003***	-0.00004 <i>t</i> = -2.128**	0.00002 <i>t</i> = 1.078	0.0001 <i>t</i> = 3.128***
Inv Chg 1M	0.002 <i>t</i> = 2.173**	0.00000 <i>t</i> = 0.265	-0.0001 <i>t</i> = -1.804*	0.001 <i>t</i> = 1.534	0.001 <i>t</i> = 2.899***
Adjusted R ²	0.005	0.0001	0.0004	0.005	0.001

Note: The t-statistics rely on the heteroskedasticity and autocorrelation consistent Newey-West (1987) standard errors estimates, using the Newey-West (1994) automatic lag selection procedure. *p<0.1; **p<0.05; ***p<0.01

Appendix D: Futures returns and risk premia

Following Szymańska et al. (2014) and Hevia et al. (2018), some key concepts about futures returns and risk premia are reviewed below in the context of the basis decomposition proposed in Section 2.2.

Spot premium

The expected hold-until-maturity return of a n -period futures contract is defined as a series of one-period expected holding returns of the futures plus a settlement return at maturity.

$$E_t \left[r_{f_t^{(n)}} \right] = E_t \left[\ln(S_{t+n}) - \ln(F_t^{(n)}) \right] = E_t \left[s_{t+n} - f_t^{(n)} \right]$$

$$= E_t \left[(s_{t+n} - f_{t+n-1}^{(1)}) + (f_{t+n-1}^{(1)} - f_{t+n-2}^{(2)}) + \dots + (f_{t+1}^{(n-1)} - f_t^{(n)}) \right]$$

The special case of the nearest one-period futures contract holding return is then defined as follows.

$$E_t \left[r_{f_t^{(1)}} \right] = E_t \left[\ln(S_{t+1}) - \ln(F_t^{(1)}) \right] = E_t \left[s_{t+1} - f_t^{(1)} \right]$$

The spot risk premium $\pi_{s,t}^{(1)}$ is defined as the expected return from holding the nearest one-period futures contract until maturity. An alternative equivalent representation is the expected spot return in excess of the futures basis. Essentially it captures the expected return above and beyond what is priced in the futures curve.

$$\pi_{s,t}^{(1)} = E_t \left[r_{f_t^{(1)}} \right] = E_t \left[r_{s_{t+1}} \right] - y_t^{(1)}$$

The spot risk premium corresponds thus to the expected return of an outright long position in the front contract. In the absence of expected spot change or parallel shift in the futures curve, the spot risk premium would correspond to the negative of the futures basis, that is the convergence of the futures price to the spot price.

Incorporating the proposed basis decomposition, the spot premium correspond to the expected spot return in excess of the long-run cost-of-carry, the next period seasonal and scarcity premia.

$$\pi_{s,t}^{(1)} = E_t \left[r_{s_{t+1}} \right] - \Upsilon_t^{(1)} - \zeta_t^{(1)} - \chi_t^{(1)}$$

Term premium

As in Hevia et al. (2018), the term risk premium is defined as the one-period expected holding return of a n -period futures contract in excess of the spot risk premium.

$$E_t \left[(\ln(f_{t+1}^{(n-1)}) - \ln(F_t^{(n)})) - (\ln(S_{t+1}) - \ln(F_t^{(1)})) \right] = E_t \left[(f_{t+1}^{(n-1)} - f_t^{(n)}) - (f_{t+1}^{(1)} - f_t^{(1)}) \right]$$

$$= E_t \left[(f_{t+1}^{(n-1)} - f_t^{(n)}) \right] \left(\pi_{s,t}^{(1)} \right)$$

$$= \pi_{y,t}^{(n)}$$

Expanding the definition provides interesting insights into the factors driving the term premium.

$$\begin{aligned}
\pi_{y,t}^{(n)} &= E_t \left[\left(f_{t+1}^{(n-1)} - f_t^{(n-1)} + f_t^{(n-1)} - f_t^{(n)} \right) \right] \pi_{s,t}^{(1)} \\
&= E_t \left[\left(f_{t+1}^{(n-1)} - f_t^{(n-1)} \right) \right] \left(y_t^{(n)} - \pi_{s,t}^{(1)} \right) \\
&= E_t \left[\left(f_{t+1}^{(n-1)} - f_t^{(n-1)} \right) \right] - E_t \left[r_{s,t+1} \right] - \left(y_t^{(n)} - y_t^{(1)} \right) \left(\dots \right)
\end{aligned}$$

The first term on the right-hand side of the last equation corresponds to the expected one-period change for a futures contract of constant $n-1$ maturity while the second term is the expected spot change. Those two terms describe both the parallel shift and twist in the futures curve. Expectations about a parallel shift in the curve simplifies the term premium to the last term while a steepening of the futures curve would result from the changes in the back-end of the curve dominating the front-end and lead to a rise in the term premium. Those two terms thus correspond to the expected return of a curve trade long a far contract and short the spot, also coined a calendar spread, betting on the steepening of the futures curve.

The last term of the equation describes the basis differential between the n -period futures contract and the front contract maturing one period ahead. This corresponds to the one-period expected return of the calendar spread *ceteris paribus*. When the basis term structure is flat, this term drops out. For a curve in contango the term is positive and impacts negatively the term premium while for a backwardated curve the term is negative and would contribute positively to the term premium.

Akin to the spot premia, the term premia is defined by the one-period expected return of a curve trade above and beyond what is priced in the futures curve. This interpretation deviates from the one proposed in Szymanowska et al. (2014) where the term premium is defined as the expected deviation from the expectation hypothesis of the basis term structure.

$$\begin{aligned}
n \left(y_t^{(n)} \right) \left(\dots \right) &= \left(y_t^{(1)} \right) \left(\dots \right) - \left(n-1 \right) E_t \left[\left(y_{t+1}^{(n-1)} \right) \right] \pi_{y,t}^{(n)} \\
&= \left(n-1 \right) E_t \left[\left(y_{t+1}^{(n-1)} \right) - y_t^{(n)} \right] - \left(y_t^{(1)} - y_t^{(n)} \right) \left(\dots \right)
\end{aligned}$$

The above definition of the term premium also captures in the first term in the right-hand side of the equation the expected slope change of the futures curve as measured by the change in the basis as time passes for the long maturity contract. This definition imposes a slightly more restrictive assumption on the dynamics of the futures curve as it assumes the basis term structure is flat and would thus fail to capture the impact of the current shape of the curve on the estimation of the term premium.

To get a further understanding of what drives the term premium let's now focus on the last term of the equation and incorporate the basis decomposition. We see below that it captures the differential between the front and far contract in terms of the long-run cost-of-carry, which captures the influence of the interest rate curve. For constant net storage cost and a flat forward rate curve, this term drops out. It also reflects the relative influence of season between contracts. When the term premium is measured from contracts with the same season the difference is null. The last term captures the scarcity risk differential and thus expectations about the relative influence of supply and demand shocks beyond usual seasonal effects. For persistent supply and demand shocks this term disappears.

$$\left(y_t^{(n)} - y_t^{(1)} \right) = \left(\gamma_t^{(n)} - \gamma_t^{(1)} \right) - \left(\zeta_t^{(n)} - \zeta_t^{(1)} \right) - \left(\chi_t^{(n)} - \chi_t^{(1)} \right) \left(\dots \right)$$

In the absence of expected slope change of the futures curve, the term risk premium would correspond to the cumulative differential between the

different components of the basis. We can thus rewrite the term premium as follows.

$$\pi_{y,t}^{(n)} = E_t \left[\left(f_{t+1}^{(n-1)} - f_t^{(n-1)} \right) \right] \left(E_t [r_{s,t+1}] - \left(\zeta_t^{(n)} - \gamma_t^{(1)} \right) + \left(\zeta_t^{(n)} - \zeta_t^{(1)} \right) + \left(x_t^{(n)} - x_t^{(1)} \right) \right)$$

Holding return

Recall the definition of the expected hold-until-maturity return of a n -period futures contract as a series of one-period expected holding returns of the futures plus a settlement return at maturity.

$$E_t \left[r_{f_t^{(n)}} \right] = E_t \left[\left(s_{t+n} - f_{t+n-1}^{(1)} \right) + \left(f_{t+n-1}^{(1)} - f_{t+n-2}^{(2)} \right) + \dots + \left(f_{t+1}^{(n-1)} - f_t^{(n)} \right) \right]$$

Using the definition of the term premium, we see that the one-period holding return of the futures maturing in n periods is the sum of the one-period spot risk premium and the n -periods term premium.

$$E_t \left[\left(f_{t+1}^{(n-1)} - f_t^{(n)} \right) \right] = \pi_{s,t}^{(1)} + \pi_{y,t}^{(n)}$$

Let's now turn to the settlement return which is defined as the one-period spot risk premium plus the seasonal risk premium attached to the season when the futures contract expires.

$$E_t \left[\left(s_{t+n} - f_{t+n-1}^{(1)} \right) \right] = E_t \left[\pi_{s,t+n-1}^{(1)} \right]$$

The hold-until-maturity return can then be rewritten as a sum of spot and term premia plus a seasonal risk premium specific to the maturity of the futures.

$$E_t \left[r_{f_t^{(n)}} \right] = \sum_{j=0}^{n-1} E_t \left[\pi_{s,t+j}^{(1)} \right] + \sum_{j=0}^{n-1} E_t \left[\pi_{y,t+j}^{(n-j)} \right]$$

Calendar spread return

The curve or calendar spread is a relative return trade along the futures curve that takes a long position in a far contract maturing in n periods and shorts a near contract maturing in $n-j$ periods and thus benefits from a steepening in the futures curve. Note that the term premium definition is the specific case where j is set at $n-1$.

$$E_t \left[r_{f_t^{(n)}} - r_{f_t^{(n-j)}} \right] = E_t \left[\left(f_{t+1}^{(n-1)} - f_t^{(n)} \right) - \left(f_{t+1}^{(n-j-1)} - f_t^{(n-j)} \right) \right] = \left(\pi_{y,t}^{(n)} - \pi_{y,t}^{(n-j)} \right) + \left(\zeta_t^{(n)} - \gamma_t^{(n-j)} \right) - \left(\zeta_t^{(n)} - \zeta_t^{(n-j)} \right) - \left(x_t^{(n)} - x_t^{(n-j)} \right)$$

The main factors driving the curve trade are thus the relative term premium between the contracts above and beyond what is priced in the futures curve (i.e. their relative long-run cost-of-carry), as well as the difference in seasonal and scarcity premia. The seasonal premium differential between two consecutive seasons corresponds to the expected basis change

resulting from the dynamics of supply and demand imbalances across seasons and is driven by the change in inventory levels. The scarcity premium differential captures the expected normalization of unexpected supply or demand shocks over that period. Appendix E investigates the influence of the various factors driving the basis on the calendar spread return.

Appendix E: Calendar spread trade

Table E1 provides information on the calendar spread premium across sectors and through seasons in the spirit of Section 5. None of the results are significant, which is consistent with the those obtained for the spot premium or the term premium. We can observe large seasonal variations in the calendar spread returns both at the market and at the sector levels.

Table E1: Calendar Spread Trade

Sector	All	Season											
		1	2	3	4	5	6	7	8	9	10	11	12
Annualized Mean Returns (%)													
Market	0.08	0.14	0.14	0.14	0.34	0.26	0.09	0.05	0.04	-0.08	-0.13	0.07	0.19
Energy	-0.03	0.14	0.11	0.01	1.09	-0.07	0.00	0.11	0.03	-0.14	-0.15	-0.03	0.14
Grains	0.19	0.19	0.11	0.03	0.10	0.26	0.17	0.27	0.37	0.28	0.13	0.32	0.41
IndustrialMetals	0.02	0.03	-0.20	-0.13	0.17	0.19	0.08	0.21	0.10	0.07	0.02	0.09	0.11
Livestock	0.01	0.52	0.58	0.21	0.35	1.38	0.49	0.13	-1.28	-1.35	-1.12	-0.63	-0.56
PreciousMetals	-0.01	-0.03	-0.01	0.02	-0.02	-0.08	0.04	-0.03	0.02	0.05	-0.04	0.02	0.02
Softs	0.01	0.13	0.40	0.43	0.14	0.15	-0.07	-0.13	-0.22	-0.25	-0.20	-0.03	0.23
Annualized Standard Deviations (%)													
Market	0.37	0.86	0.83	0.84	1.04	1.09	1.53	1.05	0.98	1.20	1.48	1.08	1.27
Energy	1.46	3.33	2.05	1.93	3.21	2.16	2.10	1.67	1.44	1.41	1.89	2.18	2.18
Grains	0.68	1.05	0.85	0.83	1.29	1.21	1.37	1.24	1.93	2.33	1.76	0.94	0.91
IndustrialMetals	0.73	1.02	2.25	1.88	1.20	0.96	0.89	0.93	0.90	0.95	0.86	1.14	1.05
Livestock	1.44	3.34	2.79	2.80	3.45	5.03	4.02	4.66	5.22	2.88	3.01	2.86	2.85
PreciousMetals	0.32	2.20	1.19	1.27	1.04	1.14	2.78	0.85	1.17	1.46	0.95	0.58	0.66
Softs	0.62	0.99	1.20	1.15	1.55	1.69	1.56	1.52	1.64	1.73	2.00	1.74	1.66
t-Statistics													
Market	1.30	0.95	1.02	1.00	1.90	1.43	0.35	0.27	0.24	-0.40	-0.51	0.37	0.88
Energy	-0.13	0.23	0.29	0.03	1.81	-0.17	0.01	0.34	0.09	-0.55	-0.43	-0.08	0.34
Grains	1.70	1.00	0.76	0.18	0.43	1.18	0.69	1.22	1.07	0.69	0.40	1.91	2.64
IndustrialMetals	0.18	0.16	-0.43	-0.34	0.70	0.99	0.43	1.11	0.55	0.36	0.12	0.41	0.53
Livestock	0.06	0.75	1.03	0.36	0.49	1.33	0.60	0.14	-1.21	-2.27	-1.86	-1.07	-0.96
PreciousMetals	-0.14	-0.07	-0.05	0.10	-0.09	-0.42	0.08	-0.19	0.08	0.19	-0.25	0.17	0.21
Softs	0.08	0.76	1.94	2.21	0.50	0.50	-0.25	-0.50	-0.76	-0.83	-0.60	-0.11	0.81

Note: Following Harvey et al. (2016), t-statistics in bold satisfy the statistical significance threshold of 3.0. For those we can safely reject the null hypothesis of returns not being significantly different from zero.

Table E2 reports the result of a set of predictive regressions as in Section 5. The calendar spread forward returns are regressed against the factors identified in Appendix D that are driving this curve trade as well as other control variables. The results are similar to those on the term premium presented in Table 8.

Table E2: Regression calendar spread on basis factors differential.

	Basis	Basis Factors	Control Sectors	Control Seasons
Constant	0.00001 <i>t</i> = 9.604***	0.00001 <i>t</i> = 9.381***		
MarketEW	-0.011 <i>t</i> = -40.310***	-0.015 <i>t</i> = -43.514***	-0.015 <i>t</i> = -57.402***	-0.015 <i>t</i> = -57.508***
BasisDiff	-0.154 <i>t</i> = -12.808***			
TermPremiumDiff		-0.090 <i>t</i> = -12.134***	-0.090 <i>t</i> = -13.157***	-0.090 <i>t</i> = -13.161***
ForegoneInterestDiff		-3.334 <i>t</i> = -3.828***	-3.245 <i>t</i> = -3.760***	-3.505 <i>t</i> = -4.034***
SeasonPremiumDiff		-0.010 <i>t</i> = -0.971	-0.010 <i>t</i> = -0.987	-0.009 <i>t</i> = -0.910
ScarcityPremiumDiff		0.265 <i>t</i> = 11.104***	0.264 <i>t</i> = 11.352***	0.267 <i>t</i> = 11.437***
Sector_Energy			0.00003 <i>t</i> = 7.325***	
Sector_Grains			0.00001 <i>t</i> = 4.309***	
Sector_IndustrialMetals			0.00001 <i>t</i> = 8.502***	
Sector_Livestock			0.00002 <i>t</i> = 3.154***	
Sector_PreciousMetals			0.00000 <i>t</i> = 4.117***	
Sector_Softs			0.00001 <i>t</i> = 4.235***	
Season_1				0.00002 <i>t</i> = 7.412***
Season_2				0.00002 <i>t</i> = 5.832***
Season_3				0.00002 <i>t</i> = 6.022***
Season_4				0.00002 <i>t</i> = 3.121***
Season_5				0.00002 <i>t</i> = 5.411***
Season_6				-0.00000 <i>t</i> = -0.355
Season_7				0.00002 <i>t</i> = 5.124***
Season_8				0.00001 <i>t</i> = 1.635
Season_9				-0.00000 <i>t</i> = -0.918
Season_10				0.00000 <i>t</i> = 0.817
Season_11				0.00002 <i>t</i> = 5.098***
Season_12				0.00002 <i>t</i> = 4.895***
Adjusted R ²	0.007	0.020	0.020	0.020
Residual Std. Error	0.001 (df = 1486596)	0.001 (df = 1265355)	0.001 (df = 1265350)	0.001 (df = 1265344)

Note:

* p<0.1; ** p<0.05; *** p<0.01

References

Bakshi, G., Gao, X. and Rossi, A. G. (2017), 'Understanding the sources of risk underlying the cross section of commodity returns', *Management Science* **65**(2), 619–641. 5.1, 5.1

Borovkova, S. and Geman, H. (2006), 'Seasonal and stochastic effects in commodity forward curves', *Review of Derivatives Research* **9**(2), 167–186. 1.2, 6

Brennan, M. J. (1958), 'The supply of storage', *The American Economic Review* pp. 50–72. 1.2

Brooks, C., Prokopcuk, M. and Wu, Y. (2013), 'Commodity futures prices: More evidence on forecast power, risk premia and the theory of storage', *The Quarterly Review of Economics and Finance* **53**(1), 73–85. 1.2

Campbell, J. Y. and Shiller, R. J. (1988), 'Stock prices, earnings, and expected dividends', *The Journal of Finance* **43**(3), 661–676. 2.2

Chordia, T., Goyal, A. and Saretto, A. (2017), 'p-hacking: Evidence from two million trading strategies', *Swiss Finance Institute Research Paper No. 17-37* . 10

Deaton, A. and Laroque, G. (1992), 'On the behaviour of commodity prices', *The Review of Economic Studies* **59**(1), 1–23. 1.1, 1.2

Fama, E. F. and French, K. R. (1987), 'Commodity futures prices: Some evidence on forecast power, premiums, and the theory of storage', *Journal of Business* pp. 55–73. 1.1, 1.2, 2

French, K. R. (1986), 'Detecting spot price forecasts in futures prices', *Journal of Business* pp. S39–S54. 1.2

Geman, H. and Nguyen, V.-N. (2005), 'Soybean inventory and forward curve dynamics', *Management Science* **51**(7), 1076–1091. 1.2

Gorton, G. B., Hayashi, F. and Rouwenhorst, K. G. (2013), 'The fundamentals of commodity futures returns', *Review of Finance* **17**(1), 35–105. 1.1, 4, 3.3, 9, 4.1, 6.3

Hamilton, J. D. (2017), 'Why you should never use the hodrick-prescott filter', *The Review of Economics and Statistics* **100**(5), 831–843. 3.3

Hannan, E. J., Terrell, R. and Tuckwell, N. (1970), 'The seasonal adjustment of economic time series', *International Economic Review* **11**(1), 24–52. 1.2, 6

Harvey, C. R. (2017), 'Presidential address: The scientific outlook in financial economics', *The Journal of Finance* **72**(4), 1399–1440. 10

Harvey, C. R. and Liu, Y. (2015), 'Backtesting', *The Journal of Portfolio Management* **42**(1), 13–28. 10

Harvey, C. R., Liu, Y. and Zhu, H. (2016), ‘... and the cross-section of expected returns’, *The Review of Financial Studies* **29**(1), 5–68. 10, 3.5, 5.1, 6.3

Hausman, J. A. (1978), ‘Specification tests in econometrics’, *Econometrica* pp. 1251–1271. 4.1

Hevia, C., Petrella, I. and Sola, M. (2018), ‘Risk premia and seasonality in commodity futures’, *Journal of Applied Econometrics* **33**(6), 853–873. 1.1, 1.2, 6, 6.3

Hirshleifer, D. (1991), ‘Seasonal patterns of futures hedging and the resolution of output uncertainty’, *Journal of economic theory* **53**(2), 304–327. 1.2, 16

Hirshleifer, D. (2001), ‘Investor Psychology and Asset Pricing’, *The Journal of Finance* **56**(4), 1533–1597. 13

Hong, H., Lim, T. and Stein, J. C. (2000), ‘Bad news travels slowly: Size, analyst coverage, and the profitability of momentum strategies’, *The Journal of Finance* **55**(1), 265–295. 12, 15

Hong, H. and Stein, J. C. (1999), ‘A unified theory of underreaction, momentum trading, and overreaction in asset markets’, *The Journal of Finance* **54**(6), 2143–2184. 12, 15

Hou, K., Xue, C. and Zhang, L. (2015), ‘Digested anomalies: An investment approach’, *The Review of Financial Studies* **28**(3), 650–705. 10

Hou, K., Xue, C. and Zhang, L. (2017), ‘Replicating anomalies’, *Fisher College of Business working paper No. 2017-03-010* (23394). 10

Jegadeesh, N. and Titman, S. (1993), ‘Returns to buying winners and selling losers: Implications for stock market efficiency’, *The Journal of Finance* **48**(1), 65–91. 12

Kaldor, N. (1939), ‘Speculation and economic stability’, *The Review of Economic Studies* **7**(1), 1–27. 1.1

Keloharju, M., Linnainmaa, J. T. and Nyberg, P. (2016), ‘Return seasonalities’, *The Journal of Finance* **71**(4), 1557–1590. 1.2, 6.3

Koijen, R. S. J., Moskowitz, T. J., Pedersen, L. H. and Vrugt, E. B. (2018), ‘Carry’, *Journal of Financial Economics* **127**(2), 197–225. 1

Linnainmaa, J. T. and Roberts, M. R. (2018), ‘The history of the cross-section of stock returns’, *The Review of Financial Studies* **31**(7), 2606–2649. 10

Malkiel, B. G. and Fama, E. F. (1970), ‘Efficient capital markets: A review of theory and empirical work’, *The Journal of Finance* **25**(2), 383–417. 11

Newey, W. K. and West, K. (1994), ‘Automatic Lag Selection in Covariance Matrix Estimation’, *Review of Economic Studies* **61**(4), 631–653. 4.1, 5.1

Newey, W. K. and West, K. D. (1987), 'A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix.', *Econometrica* **55**(3), 703. 4.1, 5.1

Routledge, B. R., Seppi, D. J. and Spatt, C. S. (2000), 'Equilibrium forward curves for commodities', *The Journal of Finance* **55**(3), 1297–1338. 1.1, 4

Shleifer, A. and Vishny, R. W. (1997), 'The Limits of Arbitrage', *The Journal of Finance* **52**(1), 35–55. 14

Sørensen, C. (2002), 'Modeling seasonality in agricultural commodity futures', *Journal of Futures Markets* **22**(5), 393–426. 1.2, 6

Szymanska, M., Roon, F., Nijman, T. and Goorbergh, R. (2014), 'An anatomy of commodity futures risk premia', *The Journal of Finance* **69**(1), 453–482. 1.1, 6.3

Telser, L. G. (1958), 'Futures trading and the storage of cotton and wheat', *The Journal of Political Economy* pp. 233–255. 1.2

Working, H. (1948), 'Theory of the inverse carrying charge in futures markets', *Journal of Farm Economics* **30**(1), 1–28. 1.1