An Econometric Analysis of Global Agricultural Commodity Prices

Simon Spavound
BSc Lancaster University

Submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy at the Department of Economics, Lancaster University

January 2019

Declaration

I hereby declare that this thesis is my own work and that it has not been submitted for any other degree.

Simon Spavound

To Dr. Efthymios Pavlidis, Prof. Ivan Paya, Prof. David Peel
and my parents.

Abstract

This thesis analyses global agricultural commodity price dynamics, with an emphasis on the causal factors behind movements and their forecastability. The analysis builds upon recent developments in the areas of time series econometrics, agricultural economics and applied economics to provide an empirical examination of agricultural commodity price movements. The main research questions addressed are the following. First, recursive unit root tests are employed to examine whether global commodity prices experienced explosive sub-periods which cannot be explained by underlying economic fundamental movements. Second, a Bayesian Structural VAR is used to model global wheat prices and decompose the causal factors behind price movements. Finally, an examination of the forecastability of agricultural commodity price series is conducted using recently developed dynamic models.

Acknowledgements

I am grateful to my PhD supervisors; Dr. Efthymios Pavlidis, Prof. Ivan Paya and Prof. David Peel, for their support and guidance throughout this project. Their enthusiasm instilled into me an interest in applied economics which I will be forever thankful for. I am particularly thankful to Dr. Pavlidis for his exceptional support, encouragement and guidance over the past several years.

I am indebted to my parents and Amy, for their unconditional support and motivation along this long road. I want to thank them for keeping on the path when I could have so often strayed away. For all that they have done and continue to do, this thesis is dedicated to my supervisors and my parents.

I thank my colleagues across the Management School for their inspirational discussions and presentations which have provided a truly motivating environment. I would particularly like to thank Oliver Schaer, members of the Management Science department, as well as the many others I have met in my years at Lancaster University. I am indebted to the Economic and Social Research Council for their financial and logistical support allowing this PhD to be completed. Finally, I express my gratitude to my fellow PhD students in the Department of Economics for providing a stimulating research environment, as well as the faculty who have inspired me to complete this journey.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Rapid price increases in the first decade of the 21st century have led to a renewed interest in understanding the causal economic factors behind agricultural price dynamics. In the period up to 2008, corn and wheat prices tripled, and rice traded at five times its previous level (Von Braun 2008), leading commentators around the world to re-examine food prices which had been comparatively stable for many years. Particular concern was placed on the impact of speculation by index investors, as epitomized by the stark testimony of Michael Masters to a United States Senate Committee convened to figure out what was causing high food prices:

> "If immediate action is not taken, food and energy prices will rise higher still. This could have catastrophic economic effects on millions of already stressed U.S. consumers. It literally could mean starvation for millions of the world's poor."
>
> Michael W. Masters, May 2008[1]

Although agriculture is a relatively small part of developed world economies, accounting for approximately 4% of global GDP, for low income countries this rises to a substantial 26% for Low Income countries (The World Bank 2018). Given the importance that food prices have on food security and recent evidence that high food price levels have been documented as leading to riots in 14 African countries (Berazneva & Lee 2013) as well as Europe, Asia and the Americas (Bellemare 2015).

With this potential for large negative economic and social effects, understanding the dynamics of agricultural commodity prices and the factors which cause such movements

[1] Masters (2008*a*)

is vital. This thesis is built upon three research questions. First, have recent explosive dynamics in agricultural commodity prices been driven by economic fundamentals, or have non-fundamental factors such as rational speculative bubbles been driving food price dynamics? Second, which are the structural economic factors which can explain these, particularly for global wheat prices? And finally, is there a parsimonious time series model which can accurately forecast global food prices?

Considering the first question, there has long been an interest in the processes behind asset prices which diverge from their underlying fundamental value. Blanchard (1979) demonstrates that such periods can be consistent with a rational expectations framework. Blanchard & Watson (1982) define such periods as rational bubbles, which occur whenever prices rationally deviate from that implied by their fundamental value. Diba & Grossman (1988) demonstrate that the presence of a rational bubble induces explosive dynamics to prices, as asset holders must be compensated for the risk of a crash by increasing prices.

This feature of rational speculative bubbles leads to a direct test for them; if prices are explosive, and the underlying fundamental price is stationary or integrated of order one (i.e., they are not explosive), then prices are being driven by a rational bubble. As a result, there have been several recent developments in the literature for detecting explosive sub-periods (e.g., Phillips et al. (2011), Phillips et al. (2015*a*)). However, explosive dynamics in prices alone is insufficient, as the underlying fundamentals driving prices could also grow exponentially. That is, observed explosive dynamics in prices is a necessary, but not sufficient, condition for rational speculative bubbles to have been present in food markets. Consequently, a substantial literature has emerged which investigates agricultural commodity markets in this way (e.g., Gutierrez 2015, Etienne et al. 2017).

However, determining the fundamental price of agricultural commodity markets is a challenging task as prices are determined globally by supply and demand conditions. Previous literature utilizes proxies for the fundamental price as it is not directly observable, leading to a joint hypothesis problem, that both the specified fundamentals are correct and that bubbles are present. Chapter 2 approaches this difficulty and addresses the first research question by utilizing the recently developed test of Pavlidis et al. (2017) which takes into consideration agents expectations of the path of prices from futures prices. This leads to the construction of a time series which mitigates the fact that the fundamental may be unobserved and as such we are unable to directly test whether it is I(1). This chapter then empirically applies the recently developed recursive unit root tests of Phillips et al. (2011) and Phillips et al. (2015*a*) to this fundamentals adjusted series as well as the price. These tests are particularly well-suited to this task as they can both detect the presence of explosive behavior in a series, whilst also providing a method for date-stamping which identifies the start and end of any detected explosive behavior. Four major agricultural commodities; wheat, corn, soybeans and oats, are examined using this strategy. The results suggest that although wheat, corn and oat prices have undergone explosive sub-periods, there is no evidence of rational bubbles being present in these markets, when we control for fundamentals.

Chapter 3 addresses the second research question by conducting a modeling exercise to understand the relative importance of causal economic factors which may have been driving agricultural prices. It builds upon the recent development of a Bayesian Structural Vector Autoregressive (SVAR) model by Baumeister & Hamilton (2015*a*) for examining the oil market. This model is adapted and applied to the global wheat market to examine the impact that important causal factors have had on the global wheat price. This model allows for the decomposition of the importance of selected economic factors by analyzing their relative contribution to wheat price movements over time. It has advantages over other Structural VAR models in that it requires less onerous restrictions for estimation and adopts a Bayesian framework which allows limited prior information to be transparently imposed upon parameter estimates. Specifically, these models require explicitly defined Bayesian priors for identification, rather than implicitly applying informative priors which are never specified although they may affect the results (Baumeister & Hamilton 2015*a*). Overall we find that the majority of wheat price movements over the period under investigation are due to wheat market specific factors rather than, for instance, increased demand from India and China pushing price upwards. However, during the period around 2008, shocks to precautionary demand explain some of the witnessed price rises, suggestive that agents may have had increased incentives to hold stocks of wheat until supply conditions are resolved with new crop entering the market.

Chapters 2 and 3 are predominately in-sample examinations of agricultural commodity price dynamics. An investigation into the out-of-sample forecast performance of a variety of models for wheat price series is presented in Chapter 4. This third research question is motivated by the importance that understanding future movements of food prices has for policymakers. This chapter analyses which models performed best at forecasting out-of-sample wheat prices in a period which covers the recent crisis (1990M1:2016M12). To address this question a battery of parsimonious econometric tests are evaluated for forecasting performance. In addition to these standard tests, dynamic models recently proposed by Raftery et al. (2010), and introduced to the economic literature by Koop & Korobilis (2012), are empirically applied. These models have several advantages over previously implemented forecasting models. They present time-varying model specifications, allowing different models to be considered at each point in time. They also indicate the probability of parameter inclusion at each time period, which provides additional analysis of the time-varying importance of factors which drive wheat price dynamics. The results from the forecasting exercise suggest that models which are based on futures prices perform poorly during this period. However, the dynamic model averaging approach provides particularly impressive forecast performance, besting the random walk benchmark at all horizons and provides a significant improvement when predicting the direction of the price changes.

These three empirical chapters contribute to the literature by applying new time series econometric methods to the analysis of the behavior of agricultural commodity price series, particularly during the recent crisis. The first chapter examines whether speculative bubbles have been present in agricultural commodity prices and led to a large increase in observed prices, before subsequently crashing. We find no evidence of rational bubbles playing such a role in any of the four commodities we investigate, suggesting that agricultural commodity prices have been driven by fundamentals. To better understand what these fundamentals are chapter 3 examines a series of factors which have been suggested are driving prices using a Bayesian Structural VAR. Having found that the majority of the price changes are down to wheat market specific supply and demand factors, rather than one specific external factor, we turn to forecast performance to give additional information to agents within these markets to better inform expectations about the future path they may take. Overall, we provide an analysis of agricultural commodities which is important for policymakers, who wish to form a coherent policy response which is dependent upon the factor driving it, for instance controls on futures market speculation in the case of speculative bubbles.

The final chapter summarizes the findings and discusses the conclusions presented in the thesis.

Chapter 2

Bubbles or Fundamentals? New Evidence for Agricultural Commodities from Futures Prices

2.1 Introduction

Global agricultural commodity prices have increased rapidly since 2006, with particularly marked run-ups peaking in 2008 and once again from 2011 onwards (see Figure 2.1). These rapid price increases disproportionately negatively impact consumers living in developing countries, as they are more directly exposed to price fluctuations in commodities in comparison to their developed world counterparts. This effect is channeled through the consumption effect as increased costs of buying a set bundle of goods reduces consumer welfare and agricultural commodities make up a large share of the consumption bundle for developing world consumers. This negative impact is moderated by the increased income that such consumers receive from higher food prices, however the available evidence suggests that the net effect is negative (De Hoyos & Medvedev 2009). For instance, high food prices have been correlated with increases in poverty in developing countries, and also with increased food riots and other social upheavals (Ivanic & Martin 2008, Ivanic et al. 2012, Bellemare 2015). The World Bank has estimated that high food prices diminished the ability of countries to meet the Millennium Development Goals most closely related to food and nutrition, with the 2007-2008 crisis resulting in approximately 105 million people being kept below the poverty line[1].

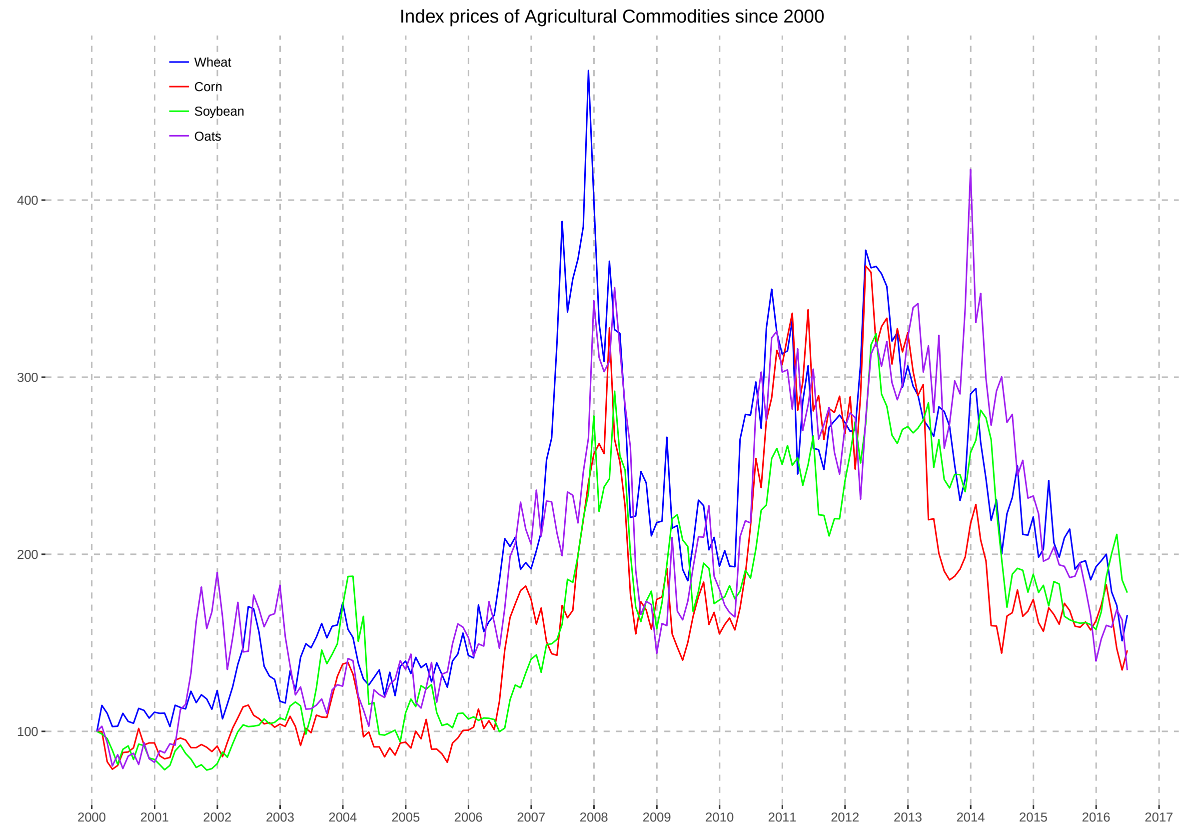

Figure 2.1: Index price comparison of major agricultural commodity markets from 2000 onwards, (January 2000 = 100).

Given the importance of agricultural markets for social welfare, it is not surprising that a large literature has developed trying to explain their behavior. This literature has mainly focused on whether high prices are the result of underlying market fundamentals, or whether there is some non-fundamental process generating higher prices. With respect to the latter factor, two key strands of literature have emerged, both of which look for speculative bubbles in commodity prices. The first searches directly for the impact of speculators on the price of agricultural commodities. This strand is motivated by the so-called 'Masters Hypothesis.'[2] Its primary concern is that financial index investors (who historically had held smaller positions in commodities) entered markets and placed excessive upward pressure on

[1]'Global Monitoring Report 2012: Food Prices, Nutrition, and the Millennium Development Goals.' World Bank, 2012.

[2]So named due to the testimony of Michael Masters, founder of Masters Capital Management, to the U.S. Senate in May 2008.

the prices of commodity futures, beyond anything that was intrinsic to the market itself. Index traders are investors who wish to gain broad exposure to agricultural markets without necessarily investing in the underlying commodities directly. They invest in instruments such as the S&P GSCI (previously, Goldman Sachs Commodity Index) which seek to replicate the returns of commodity markets. Masters (2008*b*) suggests that this increased pressure lead to the creation of large asset bubbles. This hypothesis has been investigated by a series of papers (e.g., Sanders & Irwin (2010), Irwin & Sanders (2011), Irwin (2013), Sanders & Irwin (2017)) as well as Hamilton & Wu (2015). Overall, there is little evidence to support the idea that the positions of index investors impacted upon commodity markets in the way proposed by the 'Masters Hypothesis'.

The evidence against specific groups of investors putting upwards pressure on prices does not preclude the existence of bubbles in agricultural markets due to other factors, such as self-fulfilling beliefs of future price rises. The second strand of literature investigates the presence of such speculative bubble episodes by utilizing recent advances in time series econometrics for non-stationary processes. Among the various methods proposed in the literature, the recursive unit root tests derived by Phillips et al. (2011), Phillips & Yu (2011) and further developed by Phillips et al. (2015*a*) are the most widely used. These tests are particularly appealing as they display high power in detecting episodes of explosive dynamics, and can date the origin and collapse of any bubble periods detected.

Recent studies that employ these tests to investigate whether rational bubbles appear to be present in agricultural commodity times series include Gutierrez (2013), Gilbert (2010), Etienne et al. (2014, 2015) and Etienne et al. (2017). The first two studies employ cash prices or near futures prices of agricultural commodities. Gilbert (2010), in particular, utilizes the test of Phillips et al. (2011) and finds evidence of explosive behavior in soybeans, but little evidence of such dynamics in wheat, and none for corn. Gutierrez (2013) uses the sieve bootstrap version of Phillips et al. (2011) and finds evidence for bubbles in wheat, corn and rice, but minimal evidence for soybeans. Rather than continuous near futures or spot prices, Etienne et al. (2014) utilize daily prices from individual futures contracts for a broad cross section of 12 commodities over a long time span (1970-2011). By applying the wild bootstrap version of the recursive unit root test of Phillips et al. (2015*a*), they find that all twelve investigated markets have experienced brief periods of explosive dynamics. When investigating the explosive properties of agricultural commodities from 2004 to 2015, Etienne et al. (2017) find that all of the commodities had undergone explosive behavior, however as in Etienne et al. (2014) these periods are short lived.

Although the above studies are informative about the presence of explosive dynamics in agricultural prices they cannot by themselves provide conclusive evidence on the existence of bubbles. The reason is that explosive dynamics may be induced, rather than from bubbles, from economic fundamentals that may not be observed by the researcher: the so-called joint hypothesis problem. For agricultural markets in particular, Bobenrieth et al. (2014) construct a model for storable agricultural commodities in which the presence of supply and demand shocks can make price series for storable commodities behave in a bubble-like manner, even though this is entirely the result of underlying fundamentals. Therefore, the existence of periods of explosive dynamics is a necessary, but not sufficient, condition for bubbles.

In this Chapter, we adopt a methodology recently proposed by Pavlidis et al. (2017) that mitigates this joint hypothesis problem by using futures prices to control for fundamentals. This methodology is applied to a basket of the most widely traded commodities over a period longer than previously studied in the literature, including the recent price dynamics. We find that although prices of wheat, corn and oats display explosive sub-periods there is no evidence of rational bubbles when controlling for market fundamentals. We find little evidence of explosiveness in soybeans, which conclusively rules out the presence of a rational bubble in this series. These findings illustrate that in the absence of a complete set of fundamentals, caution should be taken when interpreting the time series properties of price series alone.

We note here that these results cannot comment on the presence of any potential irrational speculative behavior which have been suggested may have influenced agent actions. Irwin (2013) explicitly suggests that the Master's hypothesis may represent irrational increases in prices across agricultural commodities due to the flows from index investors, rather than rational investment behavior. In this way there is the potential analogue of investor herding behavior or similar, with previous periods of irrational speculative bubbles such as the infamous South Sea Bubble, as detailed in Dale et al. (2005). Similarly, recent experience of so-called crypto-currencies, such as Bitcoin, which have been argued to have fundamental values of zero, experiencing explosive behavor have been pointed towards irrational behavior taking root (see the analysis of Cheah & Fry (2015) and Hafner (2018), for example). Due to the difficulties in evaluating such periods, this chapter restricts itself to analysis of the presence of rational bubble behavior, as many other authors have done (Blanchard & Watson 1982).

The rest of this Chapter is structured as follows. The next section outlines the theoretical framework for detecting rational bubbles in commodity prices. Section 3 presents the futures market dataset utilized in our empirical application and how it is prepared. Section 4 outlines the recursive unit root tests and dating strategy proposed by Phillips et al. (2015*a*), and Section 5 presents the results. The final section offers a few concluding remarks.

2.2 Rational Bubbles in Commodity Prices

To analyze whether bubbles occur in storable agricultural commodities, we adopt the popular discounted present value model for commodity pricing (Pindyck 1993), which has been widely applied (e.g., Gutierrez 2013). In this model the log spot price of a storable commodity, s_t, is equal to the discounted expected future spot price in $t+1$, and the benefits to the investor of holding the asset, known as the convenience yield:

$$s_t = E_t \left[\frac{s_{t+1} + \psi_{t+1}}{1+r} \right], \tag{2.1}$$

where E_t denotes the expectations operator, ψ_t represents the convenience yield. $r > 0$ is a commodity specific, one period discount factor which represents the rate of return an investor in the commodity would require for holding it (Pindyck 1993). Along with being positive, this discount factor is also assumed to be constant within this analysis, as in other literature examining for the presence of bubbles (Engsted & Nielsen 2012). The convenience yield is analogous to dividends within a financial context, in the sense that it represents any benefits the holder of the storable commodity receives from possession of the asset, such as the ability to avoid stock-outs (see Kaldor 1939, 1940, Working 1949). Using the law of iterated expectations and recursively substituting k periods into the future:

$$s_t = E_t\left[\sum_{i=1}^{k} \frac{\psi_{t+i}}{(1+r)^i}\right] + E_t\left[\frac{s_{t+k}}{(1+r)^k}\right]. \tag{2.2}$$

Provided that the transversality condition holds, the second term of the right hand side in Equation (2.2) converges to 0 as the value of k increases, i.e., $\lim_{k\to\infty} E_t\left[\frac{s_{t+k}}{(1+r)^k}\right] = 0$, and the current spot price is given by the discounted stream of economic fundamentals, x_t:

$$s_t = x_t = E_t\left[\sum_{i=1}^{\infty} \frac{\psi_{t+i}}{(1+r)^i}\right]. \tag{2.3}$$

In this case where $s_t = x_t$, spot prices are related to the convenience yield directly, in an analogous way to the long run relationship between stock prices and their underlying dividends, and house prices and the related to rental yields (Kivedal 2013). If the transversality condition is not imposed, there are infinitely many solutions to Equation (2.2) of the form:

$$s_t = x_t + b_t, \tag{2.4}$$

where b_t is a rational bubble term which displays the property:

$$E_t[b_{t+1}] = (1+r)b_t, \tag{2.5}$$

as shown by Diba & Grossman (1988). One of the simplest models for rational bubbles that periodically collapse which satisfies Equation (2.5) is proposed by Blanchard (1979). This model has two states. In the first state, the bubble grows exponentially; in the second, the bubble collapses to a white noise process:

$$b_{t+1} = \begin{cases} \frac{(1+r)b_t}{\pi} + \varepsilon_{t+1} & \text{with prob. } \pi, \\ \varepsilon_{t+1} & \text{with prob. } 1-\pi, \end{cases} \tag{2.6}$$

where π is the probability of being in the first state and $\varepsilon_t \sim IID(0, \sigma_\varepsilon^2)$. Several other bubble processes have been proposed in the literature. The procedure we undertake is robust to more complex bubble models, such as the periodically collapsing bubble of Evans

(1991), as long as they are rational.

Since the structural parameter r is positive, rational bubbles are in expectation explosive. On this basis, many researchers have tested for speculative bubbles in agricultural commodity markets by running right-tailed unit root tests on agricultural commodity prices. However, such tests are not conclusive. To illustrate this point most simply, suppose that economic fundamentals follow an autoregressive process of order 1:

$$x_t = \phi x_{t-1} + \mu_t, \tag{2.7}$$

where $\mu_t \sim IID(0, \sigma_\mu^2)$. If this process is explosive, i.e., $\phi > 1$, then the spot price may display explosive dynamics, even in the absence of bubbles. That is, explosive pricing is a necessary, but not sufficient, condition for bubbles.

To deal with this shortcoming, some researchers have applied unit root tests to observable variables which serve as proxies for fundamentals, for instance storage costs, alongside commodity prices. For instance, Gutierrez (2013) constructs a series for the convenience yield utilizing the difference between the nearby future contract and the settlement price of the next-to-expire futures price, net of storage costs for the difference in maturity dates.

Specifically the approach adopted by Gutierrez (2013) models convenience yields as follows:

$$\Psi_t = P_t - PF_{t,T}e^{-(i(t,T)(T-t))/365},$$

where Ψ_t is the the convenience yield and P_t, the settlement price of a nearby futures contract. The second term on the right hand side of the above equation is the settlement price of the next-to-expire futures contract, net of the cost of storage. This storage cost is computed as the CBOT daily storage cost multiplied by the number of days between the expiring contract and the next-to-expire futures contract.

If prices are found to be explosive but fundamentals are not, it follows from Equation (2.4) that bubbles must be present in the market. The problem with this approach is that researchers do not know the true model for the fundamentals, x_t, and therefore cannot correctly specify an appropriate proxy variable. As argued by several authors, this misspecification may lead to spurious inference in favor of bubbles. To mitigate this issue, we follow Pavlidis et al. (2017) and utilize information on investors expectations about economic fundamentals from the futures market.

Under general conditions of risk neutrality, the log futures rate at time t, for delivery after n periods, $f_{t,n}$, is given by agents expectations of the spot price n periods into the future, s_{t+n}:

$$f_{t,n} = E_t[s_{t+n}] = E_t[x_{t+n}] + E_t[b_{t+n}]. \tag{2.8}$$

From Equations (2.5) and (2.7), Equation (2.8) can be rewritten:

$$f_{t,n} = \phi^n x_t + (1+r)^n b_t, \tag{2.9}$$

and bringing Equation (2.4) forward n periods into the future, the spot rate at time $t+n$ can be expressed:

$$s_{t+n} = x_{t+n} + b_{t+n}. \tag{2.10}$$

Assuming that the bubble is still extant, recursive substitution of Equations (2.6) and (2.7) into Equation (2.10) leads to:

$$s_{t+n} = \phi^n x_t + \frac{(1+r)^n}{\pi^n} b_t + \epsilon_{t+n}, \tag{2.11}$$

where ϵ_{t+n} is a combination of two moving average processes, $\epsilon_{t+n} = \sum_{i=1}^{n} \phi^{n-i} \mu_{t+i} + (1+r/\pi)^{n-i} \varepsilon_{t+i}$. When examining Equations (2.9) and (2.11), it can be seen that the future spot price (s_{t+n}), is greater (in expectation) than the forward rate $f_{t,n}$. This is due to rational agents attaching a non-zero weight on the possibility that the bubble bursts at some future point in time. Consequently the actual growth rate, $(1+r)^n/\pi^n$, exceeds the expected growth rate, $(1+r)^n$. By subtracting Equation (2.9) from Equation (2.11) it follows that:

$$
\begin{aligned}
s_{t+n} - f_{t,n} &= \left(\phi^n x_t + \frac{(1+r)^n}{\pi^n} b_t + \epsilon_{t+n}\right) - \left(\phi^n x_t + (1+r)^n b_t\right) \\
&= (1+r)^n \left(\frac{1}{\pi^n} - 1\right) b_t + \epsilon_{t+n}.
\end{aligned} \tag{2.12}
$$

Equation (2.12) shows that the difference between the spot price n periods into the future and its corresponding futures rate, $s_{t+n} - f_{t,n}$, or equivalently the current spot price minus the associated n period futures rate, $s_t - f_{t-n,n}$, is a linear combination of the rational bubble and two stationary moving average terms. Thus the series is explosive when a bubble is ongoing. This suggests an empirical strategy where an agricultural commodity market is tested by running recursive unit root tests on this series directly. As Equation (2.12) does not feature any fundamental component, x_t, this test has a major advantage over traditional tests on the spot rate alone by mitigating the potential misspecification of the underlying fundamentals.

2.3 Data

For our empirical application, we use the monthly futures prices of wheat, corn, soybeans and oats, traded on the Chicago Mercantile Exchange (CME)[3]. These commodities have been widely investigated in the literature as grain markets are believed to have been most susceptible to the effects of index traders and therefore are of the most concern (Etienne et al. 2017). Prices for each commodity are available at a monthly frequency over a long time span, from January 1960 for wheat and corn, and January 1970 for soybeans and oats, until June 2016. This results in 678 observations for both wheat and corn, and 558 for soybeans and oats, approximately 56 and 46 years, respectively. This time span is much longer than that has previously been investigated, and includes the most recent pricing episodes as well as historic dynamics, and utilizes all available data. Table 2.1 details summary statistics for the four commodities post-2000, showing the rapid increase in prices which has occurred since the turn of the century. For instance from 2000 prices increased by a factor of 8 for corn and oats from their post-2000 lows. All commodities underwent well-known comovement (Pindyck & Rotemberg 1990) during this period and shared similar dynamics until 2008,

[3] All data were downloaded utilizing the Quandl database API.

as shown in Figure 2.1. Each commodity experiences multiple peaks and troughs after this time.

	Obs.	Annualized Growth Rate	Standard Deviation
Wheat	678	7.87	189.98
Corn	678	8.50	168.15
Soybean	558	7.90	350.71
Oats	558	9.30	91.96

Table 2.1: Summary Statistics for the four globally traded commodity monthly dataset. Annualized growth rates calculated since 2000, full analysis conducted on entire dataset.

For spot prices of agricultural commodities, we follow previous studies and utilize near futures prices as a proxy for the spot prices of agricultural commodities. This overcomes two potential confounding issues which affect these markets. First, many authors have suggested that spot markets for agricultural commodity contracts are not true spot prices, due to delayed delivery (Fama & French 1987). Second, it avoids any complications arising due to the recent phenomenon of non-convergence of futures prices to spot prices (Garcia et al. 2015). One proposed reason for non-convergence of futures prices to spot prices is the liquidity causes by underlying market structures which affected the conversion of futures prices into real grain (Adjemian et al. 2013). As this hampers the information aggregation function of futures markets, spot prices during this period may not fully reflect the full information set of market participants. This period of non-convergence occurs during the period of interest when many commentators suggest bubble like behavior has occurred, therefore a consistent spot price proxy is required throughout the period. The use of near futures prices as a proxy for the spot price in agricultural commodity markets is widespread, for instance Gutierrez (2013), apply it to a similar study. We therefore investigate the time series characteristics of the CME near futures price series for wheat, corn, soybeans and oats.

We construct a series of constant maturity (3-month) futures prices for each commodity, capturing agents expectations about the future spot price. A continuous series is required for the testing methodology of Phillips et al. (2015*a*) to be implemented. Due to agricultural commodity futures contracts having fixed maturities spaced unequally throughout the calendar year (e.g., wheat futures contracts mature in five months; March, May, July, September and December), we must interpolate between adjacent contracts to obtain a single continuous time series with constant maturity which captures agents evolving expectations. This is accomplished by taking the price series for every individual contract (280 separate contracts have been traded since 1960 in the case of wheat and corn) available and utilizing linear interpolation[4] to generate expectations of prices for months when no contract matures.

For instance, to generate a 3 month futures price for a hypothetical wheat contract maturing in April, agents expectations from the two adjacent contracts are used (i.e., those maturing in March and May of that calendar year). The prices from each neighboring contract are taken in January and used to produce a futures price which captures agents expectations of wheat prices in April. From this a series of spot prices minus agents expectations of that spot price can be generated, and on this we can conduct our analysis. The three series, s_t, f_t and $s_t - f_{t-3,3}$ for each of the four commodities is shown in Figure 2.2. Close inspection of the first two series reveal similar patterns, exhibiting relative stability with prices roughly doubling between peaks and troughs over multiple years. Then around 2005 all series begin to experience a rapid increase in prices (more than tripling) to around 2008 followed by a sudden drop, and much higher prices than in the earlier period subsequently. This latter period has motivated increased scrutiny into the functioning of these markets.

2.4 Recursive Unit Root Testing

A variety of unit root tests have been employed in the literature to test a time series for evidence of explosive behavior. In this Chapter, we adopt the recursive unit root methodology proposed by Phillips et al. (2015*a*) as it has two important features which make it well suited for our purpose. First, in contrast to the standard augmented Dickey-Fuller (ADF) (Dickey & Fuller 1979) test and others which allow for a single change in persistence, this methodology is consistent and displays good power properties in the presence of multiple bubble episodes. Second, it provides an associated date-stamping strategy which allows *ex post* identification of the origin and termination dates of explosive dynamics within a time

[4] The procedure undertaken here is the same as detailed in Alexander (2012), which details a procedure using linear interpolation for generating constant maturity futures series in the energy market.

series.

The strategy of Phillips et al. (2015a) is based on a sequence of standard ADF regression equations:

$$\Delta y_t = a_{r_1,r_2} + \gamma_{r1,r2} y_{t-1} + \sum_{i=1}^{k} \psi^i_{r_1,r_2} \Delta y_{t-i} + \epsilon_t, \tag{2.13}$$

where y_t denotes a generic time series process, $\epsilon_t \overset{iid}{\sim} N(0, \sigma^2_{r_1,r_2})$, and r_1 and r_2 are fractions of the sample size specifying the start and end points of a sub-sample. The null hypothesis of interest is non-explosive behavior (either stationarity, I(0) or a unit root, I(1)) in y_t, $H_0 : \gamma_{r_1,r_2} = 0$, against the alternative of an explosive process, $H_1 : \gamma_{r_1,r_2} > 0$. The corresponding test statistic is given by:

$$ADF^{r_2}_{r_1} = \widehat{\gamma}_{r1,r2} / \text{s.e.}(\widehat{\gamma}_{r1,r2}). \tag{2.14}$$

To allow for multiple switches between $I(1)$ and explosive dynamics, Phillips et al. (2015a) propose a recursive procedure which is based on the estimation of the ADF regression equation on sub-samples of the available data. In particular, the authors propose estimating the test statistic given by Equation (2.14) for all possible sub-samples of data, for a given minimum window size r_0. The maximum of all estimated test statistics, called the Generalized Supremum ADF ($GSADF$) test statistic:

$$GSADF(r_0) = \sup_{r_2 \in [r_0,1], r_1 \in [0, r_2 - r_0]} ADF^{r_2}_{r_1}, \tag{2.15}$$

is compared to the right-tailed critical value from its (non-standard) limit distribution under the null:

$$\sup_{r_2 \in [r_0,1], r_1 \in [0, r_2 - r_0]} \left\{ \frac{\frac{1/2}{r_w}[W(r_2)^2 - W(r_1)^2 - r_w] - \int_{r_1}^{r_2} W(r)dr[W(r_2) - W(r_1)]}{r_w^{1/2}\{r_w \int_{r_1}^{r_2} W(r)^2 dr - [\int_{r_1}^{r_2} W(r)dr]^2\}^{1/2}} \right\}, \tag{2.16}$$

where $r_w = r_2 - r_1$.

If the null hypothesis of a unit root is rejected, an associated dating strategy can be implemented to date the start and end of periods of explosive behavior. This strategy is based on the following Backward Supremum ADF ($BSADF$) statistic:

$$BSADF_{r_2}(r_0) = \sup_{r_1 \in [0, r_2 - r_0]} ADF_{r_1}^{r_2}. \tag{2.17}$$

Phillips et al. (2015a) define the origin of the explosive episode as the first instance that the $BSADF$ test statistic sequence exceeds its critical value:

$$\widehat{r}_e = \inf_{r_2 \in [r_0, 1]} \{r_2 : BSADF_{r_2}(r_0) > scu_{r_2}^{\beta_T}\}, \tag{2.18}$$

and the termination date as the first observation that the $BSADF$ test statistic sequence falls below its critical value:

$$\widehat{r}_f = \inf_{r_2 \in [r_0, 1]} \{r_2 : BSADF_{r_2}(r_0) < scu_{r_2}^{\beta_T}\}, \tag{2.19}$$

where β_T is the chosen significance level, and $scu_{r_2}^{\beta_T}$ denotes the $100 - \beta_T\%$ critical value.

2.4.1 Wild Bootstrap GSADF Test

Many studies provide evidence that suggests movements in commodity prices display time-varying volatility. To draw inference in the presence of (potentially) heteroskedastic errors in Equation (2.13), we adopt a wild bootstrap version of the $GSADF$ test, following several other studies in the literature (e.g., Etienne et al. (2014), Harvey et al. (2016)). The wild bootstrap procedure deployed consists of five steps:

1. Impose the null of a unit root process, and fit the ADF regression Equation (2.13) to the actual data.

2. Save the estimated coefficients, $\hat{\psi}_i$, and the fitted residuals, $\hat{\epsilon}_t$.

3. Generate a new series, y^*, according to the data generating process suggested by the null hypothesis:

$$y_t^* = y_{t-1}^* + \sum_{i=1}^{k} \hat{\psi}_i \Delta y_{t-i}^* + \epsilon_t^*, \tag{2.20}$$

where ϵ_t^* is constructed by multiplying $\hat{\epsilon}_t$ by a random variable, v_t, which follows the Rademacher distribution:

$$v_t = \begin{cases} +1 & \text{with prob. } \pi = 0.5, \\ -1 & \text{with prob. } 1 - \pi. \end{cases} \tag{2.21}$$

Note that, due to the properties of the Rademacher distribution (i.e., $E_t[v_t] = 0$; $E_t[v_t]^2 = 1$; $E_t[v_t^3] = 0$; and $E_t[v_t^4] = 1$), any heteroskedasticity or symmetric non-normality in the fitted residuals, $\hat{\epsilon}_t$, is preserved in ϵ_t^*. The Rademacher distribution is utilized here as it is a convenient way to produce random draws which maintain the properties of the fitted residuals, other distributions could have been utilized, for instance, a draw from a standard normal distribution.

4. Compute the $GSADF$ test statistic for y_t^*.

5. Repeat steps three and four B times, where B is the desired number of bootstrap replications, to obtain the bootstrapped distribution.

6. Compare the $GSADF$ test statistic for the actual data to the $100 - \beta_T\%$ critical value of the computed bootstrapped distribution.

The implementation of the above procedure requires the selection of a minimum window size, r_0, and lag length, k. For the former, we follow Phillips et al. (2015*a*)[5] and follow the rule which determines optimal window size, $r_0 = 0.01 + 1.8/\sqrt{T}$. This yields a minimum window size of 54 monthly observations for wheat and corn, and 48 observations for soybeans and oats[6]. With regards to the lag length, we set $k = 2$, the maximum lag length, to deal with the potential presence of serial correlation in the error term due to the overlapping nature of the data. This k therefore captures the serial correlation potentially induced by construction as each traded contract overlaps sequential contracts by two months.

[5] For further technical details see the associated technical supplement (Phillips et al. 2015*b*).
[6] For all series this corresponds to approximately 8% of the total available observations.

2.5 Empirical Results

Our empirical analysis starts with an examination of whether there are explosive subperiods in the levels of the spot and futures price series for each of the agricultural commodity markets[7]. The first row in Table 2.2 reports standard ADF unit root test statistics for each commodity, the third the $GSADF$ test statistics. Overall, for wheat, corn and oats the results are in line with those of previous studies; first, test statistics for the spot and forward series are almost identical, indicating that in each case the two series display very similar persistence properties. Second, for spot and futures prices for these three commodities, the $GSADF$ test rejects the null hypothesis of a unit root in favor of explosive dynamics, but the ADF fails (by a substantial margin) to do so. This discrepancy between tests is expected since the standard ADF performs poorly in the presence of exponential price increases followed by a price collapse, such as the ones displayed in Figure 2.1; while the $GSADF$ has good power properties. Third, a comparison of the last two rows of the table reveals that the inference procedure may be important. Although the null hypothesis of a unit root is rejected by the $GSADF$ test in all cases, the fact that the wild bootstrap yields substantially larger critical values implies that caution should be taken when drawing inference based on the assumption of homoskedasticity.

The ADF test statistic for soybeans follows a similar pattern to the other three commodities, however the $GSADF$ test statistic for soybeans in columns 7 and 8 does not exceed the associated 95% critical value, therefore there is little evidence that there were non-stationary episodes for soybeans. This finding is contrary to the previous results of Gilbert (2010) and Gutierrez (2013), both of which found evidence of explosiveness, possibly due to the differences in frequency and sample size used in these studies. Both used higher frequency data with shorter sample sizes, which may explain the differences in the time series properties between this study and those.

Having rejected the null of a unit root process in three of the commodities, we can proceed with the identification of the exact periods of explosive dynamics in the markets. Figures 2.3 and 2.4 show the series of $BSADF$ test statistics and the associated sequence

[7]Similar results for the $s_t - f_{t-3,3}$ series are found if the logarithm of each series is taken.

	Wheat[a]			Corn[a]			Soybean[b]			Oats[b]		
	s_t	f_t	$s_t - f_{t-3,3}$	s_t	f_t	$s_t - f_{t-3,3}$	s_t	f_t	$s_t - f_{t-3,3}$	s_t	f_t	$s_t - f_{t-3,3}$
ADF	-3.42	-3.52	-6.77	-3.33	-3.37	-8.12	-2.96	-3.02	-8.78	-3.15	-3.17	-8.14
SADF	2.37	2.21	0.26	4.05	4.24	-0.65	-0.94	-0.91	-3.81	0.58	2.09	-3.12
GSADF	4.34	4.75	1.07	4.05	4.27	0.33	2.15	2.03	0.03	3.23	2.94	1.60
Critical Values												
ADF	-0.08	-0.08	-0.08	-0.08	-0.08	-0.08	-0.08	-0.08	-0.08	-0.08	-0.08	-0.08
GSADF	2.29	2.29	2.29	2.25	2.25	2.25	2.22	2.22	2.22	2.22	2.22	2.22
WB GSADF[c]	3.36	3.48	5.46	3.29	3.36	5.22	3.22	3.30	3.56	2.89	2.91	4.01

Table 2.2: ADF, $SADF$ and $GSADF$ test statistics for spot, s_t, future, f_t, and spot minus future, $s_t - f_{t-3,3}$ prices. Critical values reported at the 95% significance level. $GSADF$ critical values are based on 2000 Monte Carlo replications. Wild bootstrap $GSADF$ critical values are based on 2000 bootstrap replications.

[a] Minimum window size, r_0, is 54 observations.
[b] Minimum window size, r_0, is 48 observations.
[c] Lag length, k, set to 1.

of critical values for the spot and futures prices, respectively. It is evident from these figures that each series exhibits periods of explosive behavior, defined as any period when the $BSADF$ test statistic exceeds its critical value sequence. The timing of explosive pricing periods is heterogeneous, both wheat and corn exhibit explosiveness around the time of previous food crisis from 1973-1974, a period which commentators have compared to recent price movements (Headey 2011). The more recent period of concern is also captured here from 2007-2008, particularly for wheat. The oat price has one marked period of explosiveness beginning in the second half of the 1990s, and the $BSADF$ test does not identify any explosive dynamics since then. The finding of explosive dynamics in spot prices is a result which has led previous studies to conclude that speculative bubbles exist in agricultural markets. However, as shown in Section 2, in the absence of knowledge of the true fundamentals process, these explosive dynamics in the price series cannot be interpreted as conclusive evidence for rational bubbles. To accomplish this, analysis must be carried out on the difference between spot and futures price series, $s_t - f_{t-3,3}$.

The third column for each commodity in Table 2.2 shows the unit root test results for the $s_t - f_{t-3,3}$ series. Both the ADF and $GSADF$[8] test statistics are now below the 95% critical values, which suggests that there are no sub-periods of explosive dynamics. Since s_t and f_t are affected by market fundamentals, whereas by construction the $s_t - f_{t-3,3}$ series is not, the documented rapid price rises and falls of agricultural commodities may be attributed to fundamentals rather than rational bubbles. There are a number of explanatory factors proposed in the literature that may offer insight into the recent turbulence in commodity markets, for instance the impact of increased biofuel production for grains (Ajanovic 2011, Carter et al. 2016). Overall, the above results complement those studies that find little evidence of bubbles when examining the relationship between index investors positions and the price of global agricultural commodities.

[8]For completeness even though identified as stationary by the $GSADF$ test the $BSADF$ test results for these series are included in Figure 2.5.

2.6 Conclusion

Recent global food commodity price spikes have been highly disruptive, with a particularly negative impact on consumers in developing countries. This has led to questions about the efficiency of global agricultural markets. Many studies have found that agricultural commodity prices experienced periods of explosive pricing during the past decade and concluded in favor of speculative bubbles. However, due to the unobservable nature of the fundamentals driving agricultural commodity prices, their findings may be subject to a joint hypothesis problem. This arises due to the fact that any rejection of the null hypothesis may occur due to either the presence of bubbles, or the incorrect specification of the underlying fundamental process.

In this Chapter, we adopt a recently developed methodology which mitigates the joint hypothesis problem by using information agents expectations of the future. We apply this methodology to a panel of four globally traded agriculturally commodity prices over a time span longer than previously examined in the literature. Our findings suggest that although the spot and futures price series for wheat, corn and oats exhibit episodes of explosive behavior over the past few decades, there is no evidence for speculative bubbles. In the case of soybeans, contrary to previous studies we find no evidence of explosive behavior in the time period examined. This result is indicative of the importance of correctly accounting for movements in fundamentals when explaining the evolution of commodity prices.

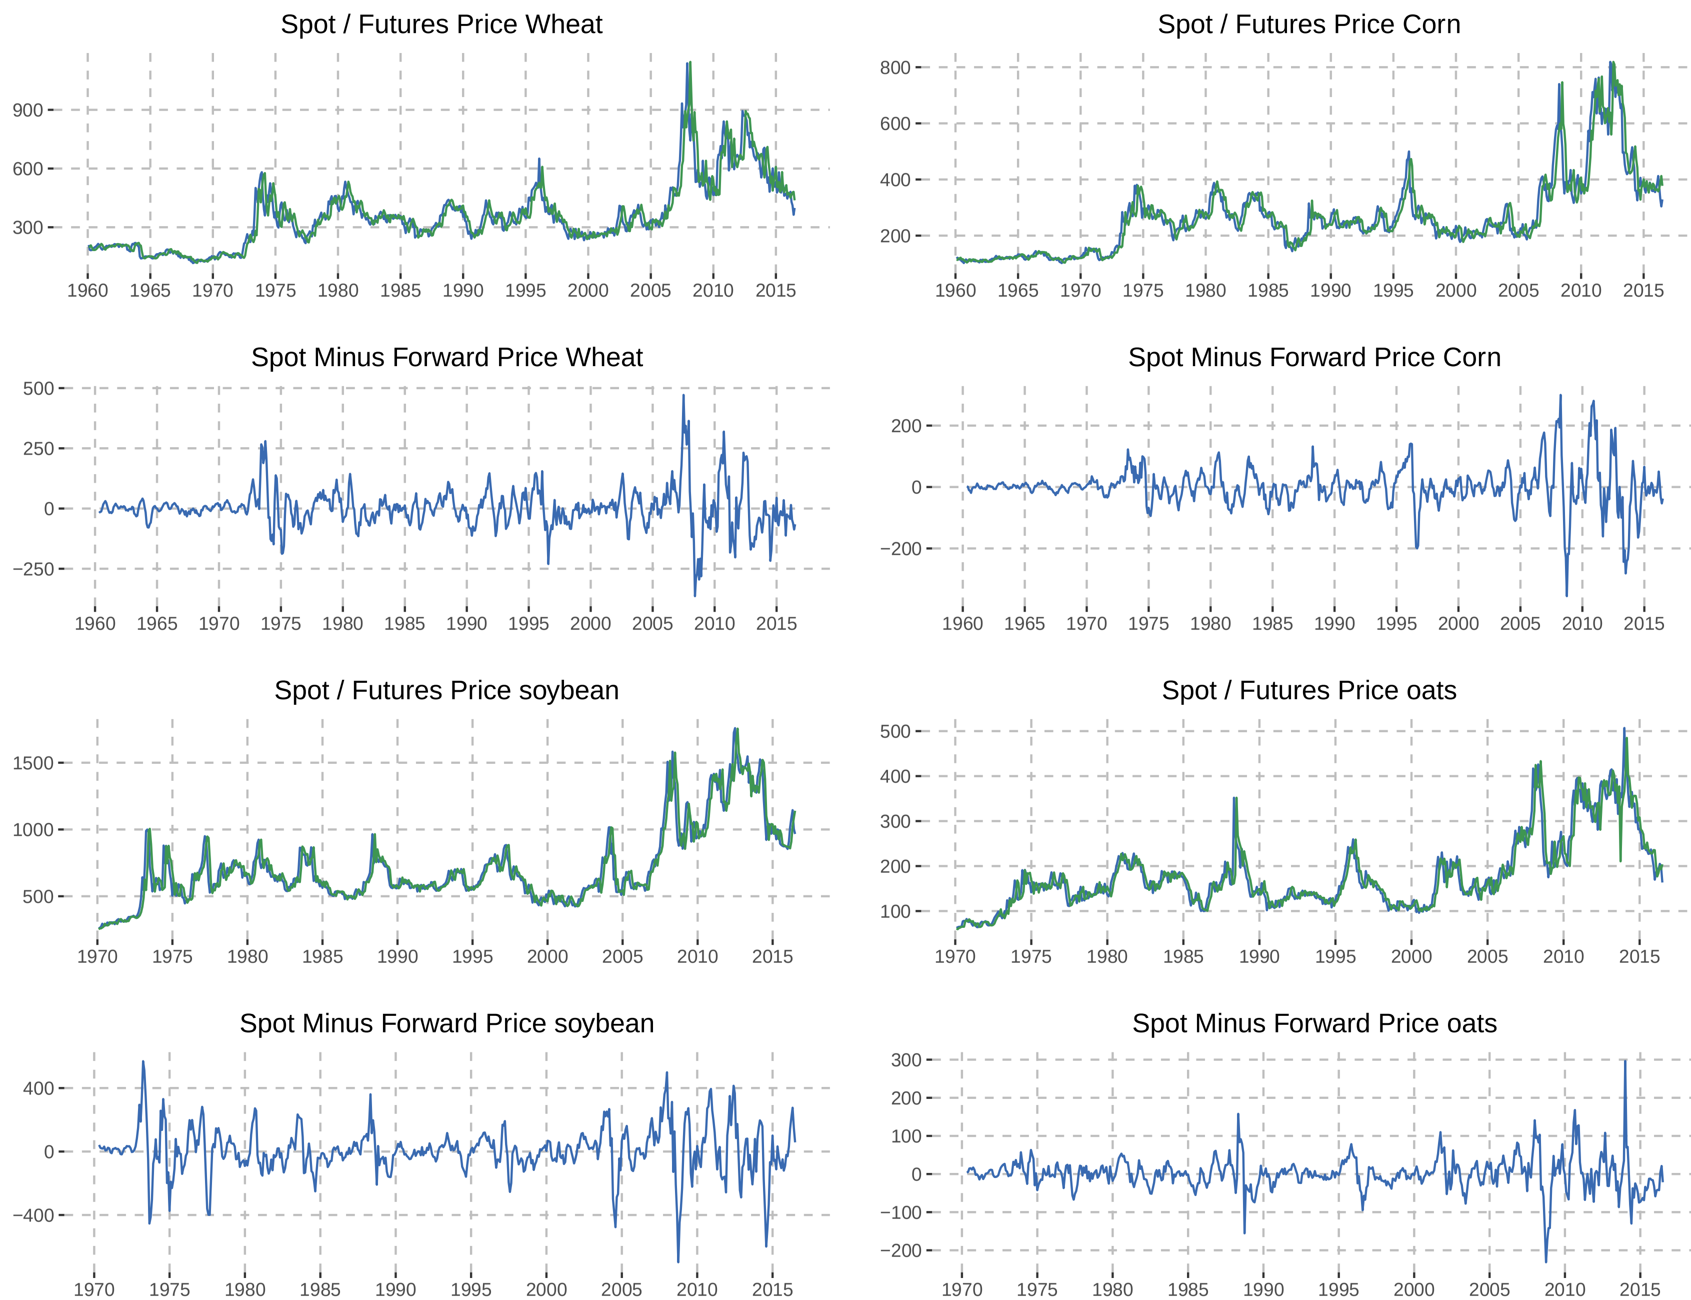

Figure 2.2: Spot, s_t, 3-month future, $f_{t-3,3}$, and spot minus futures prices, $s_t - f_{t-3,3}$ for all four commodity series.

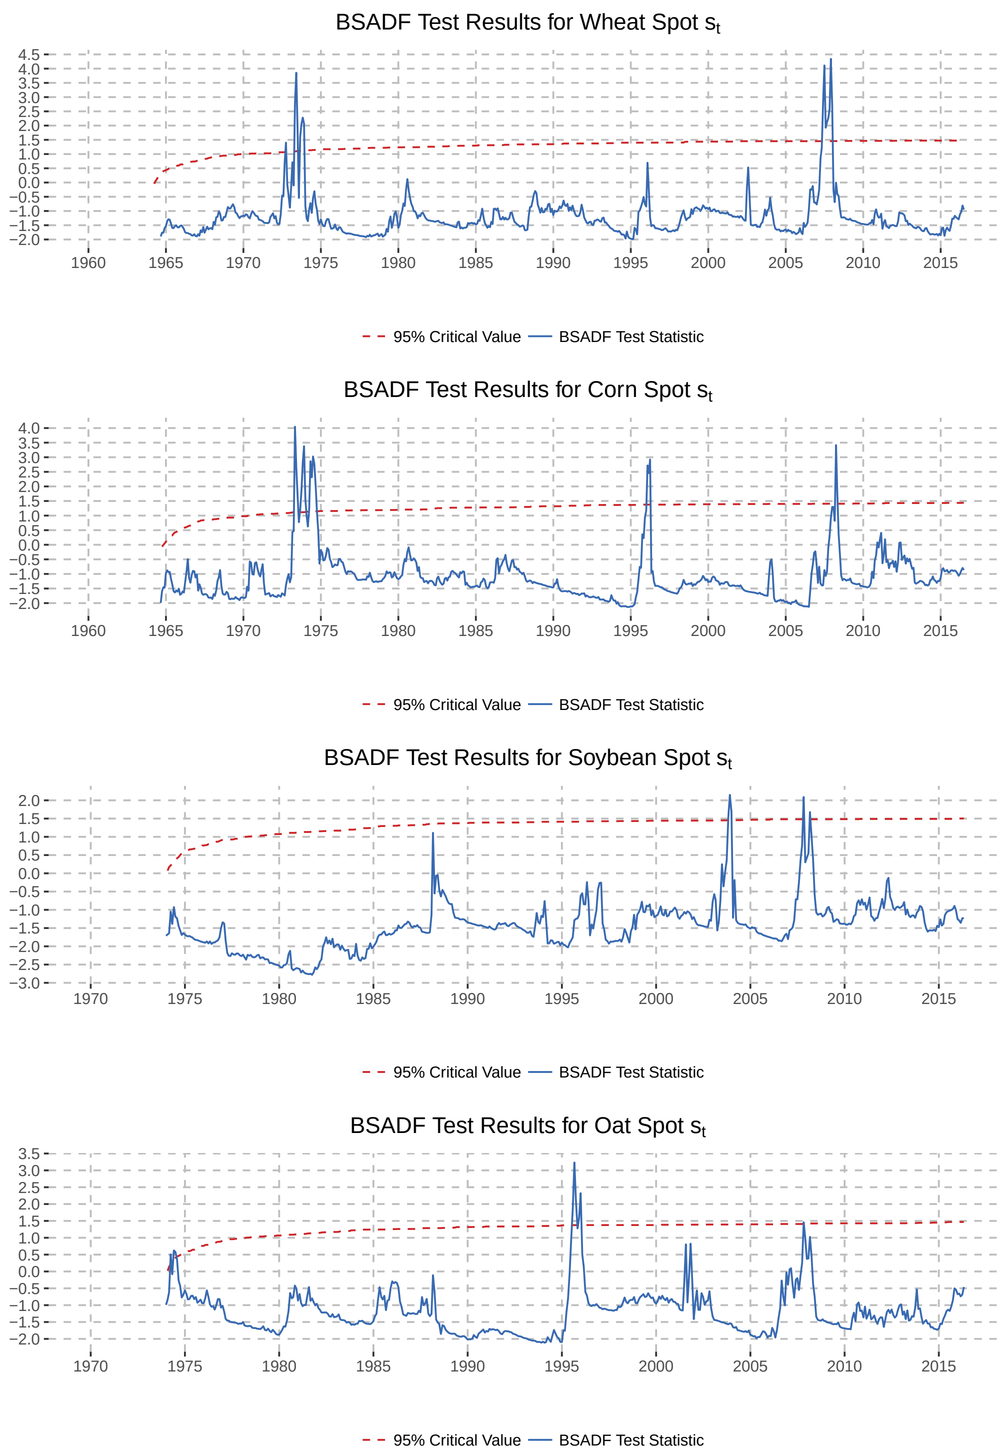

Figure 2.3: $BSADF$ test results for spot prices, s_t.

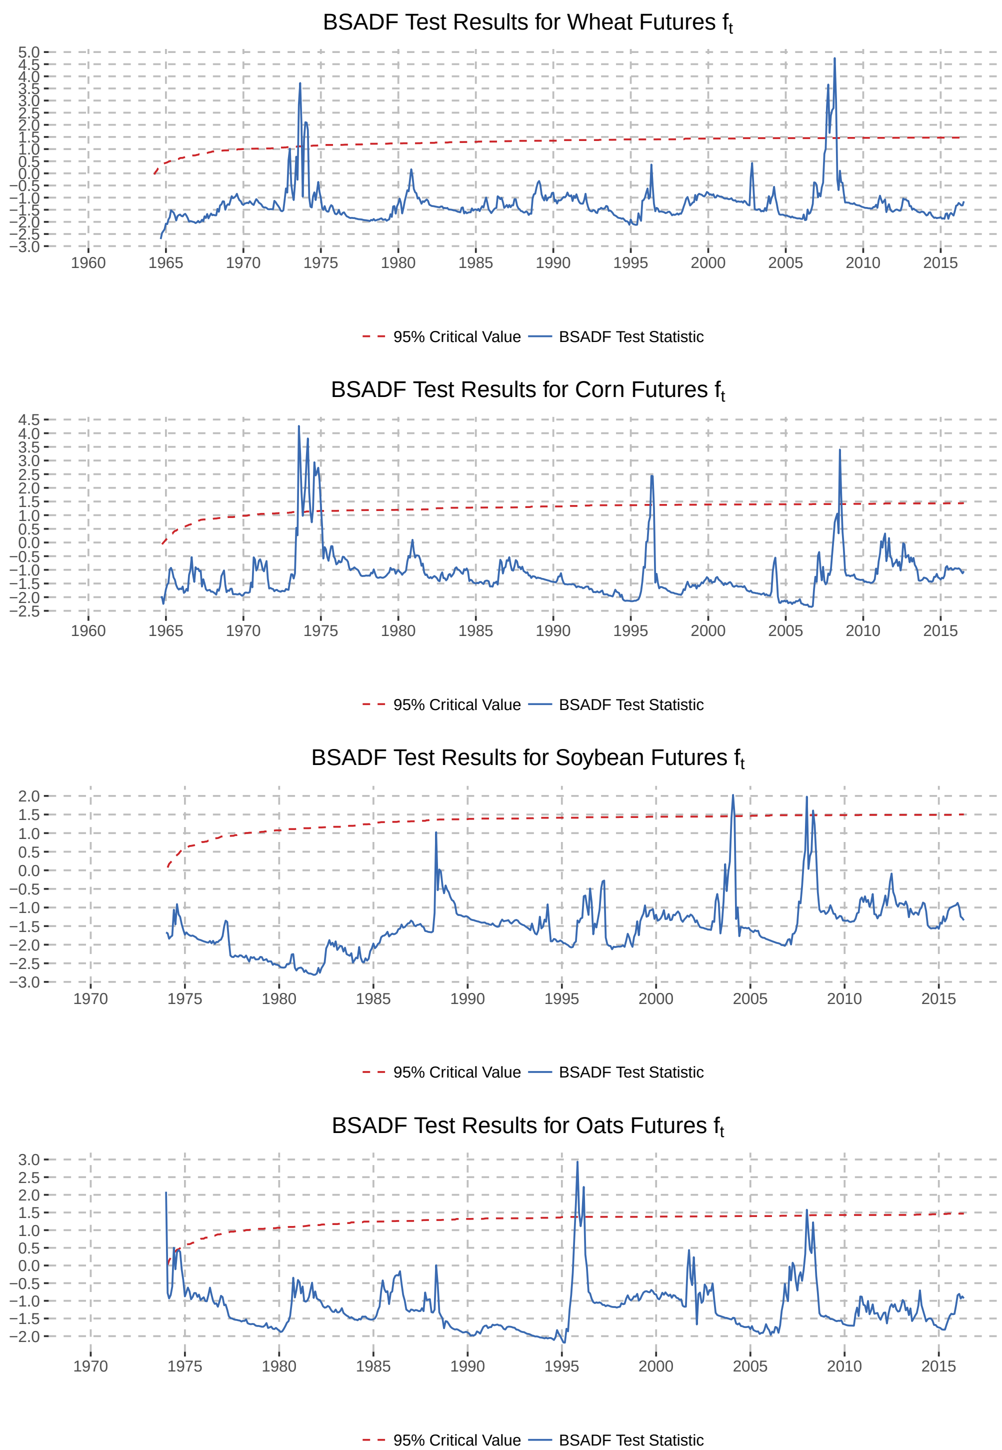

Figure 2.4: $BSADF$ test results for futures prices, f_t.

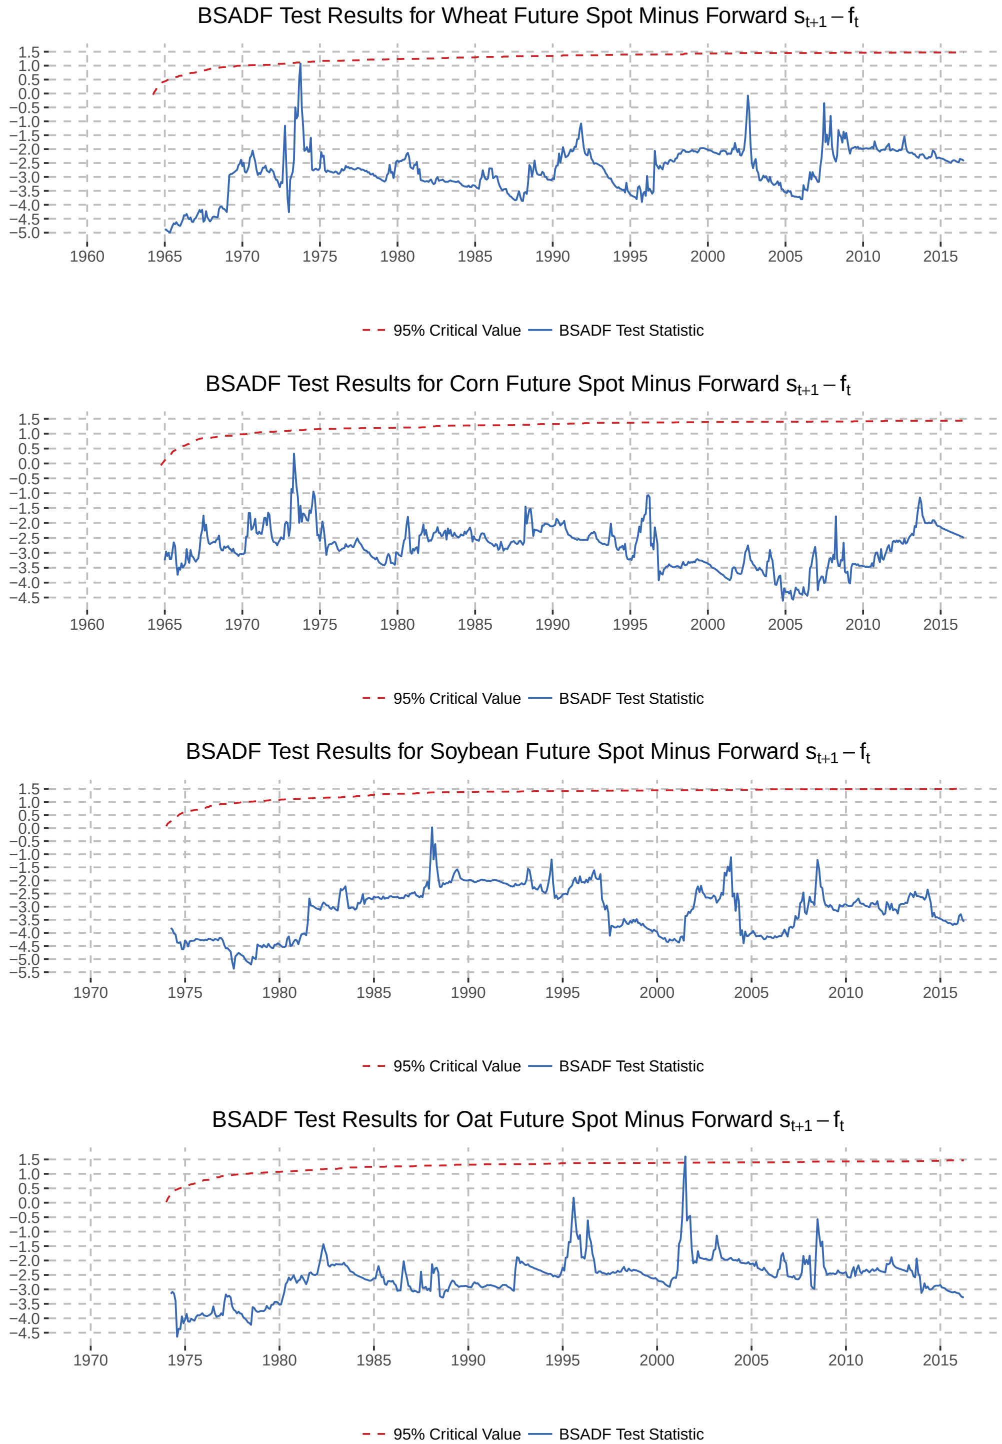

Figure 2.5: $BSADF$ test results on the spot minus 3 month futures price, $s_t - f_{t-3,3}$.

Chapter 3

Decomposing Shocks to Global Wheat Prices: Evidence from a Structural Bayesian VAR

3.1 Introduction

The High Level Panel of Experts on Food Security and Nutrition (Daviron et al. 2011), convened by the United Nations in the wake of the recent crisis, detailed the disruptive impact that food price spikes and the associated increased volatility have had on nutritional outcomes and food security, particularly in the developing world. Sudden and persistent movements in the price of agricultural commodity prices, such as those witnessed in the early 2000s (see Figure 3.1), have led to renewed interest in understanding the drivers behind commodity price movements. Many studies have since investigated different potential causal economic factors as key drivers of increased food prices (Headey & Fan 2008, 2010). However, less well understood is the relative contribution each potential factor has had on global agricultural commodity prices and how this has changed over time. As policymakers must choose an appropriate response understanding the relative importance of each proposed factor is crucial to identifying an appropriate policy response.

In this chapter, we assess the relative contribution to wheat price movements of key economic factors which have been proposed as driving agricultural commodity prices. Wheat is a staple crop which has the highest land area devoted to its cultivation worldwide (Janzen et al. 2014) and provides a significant fraction of calories to consumers with (increasingly) westernized diets (Shewry & Hey 2015). The impact of large-scale commodity price rises has had a relatively minor effect on developed country consumers due to the highly manufactured

Figure 3.1: Global Near Futures Wheat Price (cents/bushel) since 2000.

nature of wheat-based products they consume. However, the impact of these price changes disrupts the economics of producers worldwide, and can also be felt acutely by consumers in developing countries who are directly exposed to global commodity price movements. As well as being an important foodstuff, wheat is an essential input into the farming industry as animal feed. Therefore, understanding the underlying dynamics of wheat prices is key to explaining agricultural commodity movements as a whole, and correctly identifying the appropriate response to avert adverse effects globally.

This chapter contributes to an extensive literature which seeks to investigate an extensive array of factors which can help to explain recent price dynamics. Causal economic factors which have been considered by the literature include: i) the influence of non-fundamental factors, such as rational speculative bubbles or explosive dynamics caused by the sudden entry of index investors (the so-called 'Masters Hypothesis') (Sanders & Irwin 2010, Irwin & Sanders 2011, Irwin 2013, Gutierrez 2013, Etienne et al. 2014), ii) Changes in global economic activity caused by macroeconomic shifts, for instance increased demand from developing countries such as India and China, iii) Increased comovement with other commodities (including oil due to stronger linkages from increased bioethanol usage) (Ajanovic 2011, Tang & Xiong 2012) and iv) Trade shocks (Headey 2011).

Typically these factors are looked at in isolation, and we capture the relative impact the most widely mentioned of these have had on wheat price movements within a structural model. Specifically, we adopt a Bayesian Structural Vector Autoregressive approach to model the structural relationships between the economic factors under investigation. The SVAR approach has been widely adopted across many different fields of applied macroeconomics since first proposed by Sims (1980). Due to the well-known parameter identification problem, when constructing a structural VAR model a series of restrictions are required to ensure that the model is correctly identified to form valid inference from impulse response functions. However, typical exclusion restrictions may not accurately capture the uncertainties surrounding structural relationships and are therefore potentially unappealing from a theoretical perspective (Kilian 2015).

More recently, a series of models have been proposed using a part-complete set of identification restrictions, sign restrictions, which were originally applied to the problem of disentangling the effect of monetary supply shocks to output (Uhlig 2005). This methodology, along with recent updates to the procedure (see Rubio-Ramírez et al. (2010, 2014)), has become a standard tool in the empirical macroeconomics literature. Empirical applications have included examining; monetary policy (Vargas-Silva 2008, Neri et al. 2017), technology shocks, (Dedola & Neri 2007) and, most similar to this study, oil markets (e.g., Baumeister & Peersman (2010), Van Robays (2012), Kilian & Murphy (2012, 2014)).

This sign-restricted approach requires sign restrictions to be imposed on the matrix of contemporaneous impacts between variables. These restrictions should be based on prior information or theoretical guidance regarding the expected direction (not magnitude) of the contemporaneous response. This approach allows prior economic information regarding the direction of a response to be used for identification. For instance, in a model of supply and demand a supply shock will tend to increase quantity and decrease price, whereas shocks to demand will tend to increase both. This information from economic theory relates to the direction of the impact on prices and quantities and can be used to distinguish between shocks to supply and demand due to the differential response.

Until recently this approach has been seen as a neutral way of overcoming the parameter identification problem. However, in a recent contribution, Baumeister & Hamilton (2015*a*), show that this methodology is fundamentally Bayesian (as some prior information is utilized to justify the chosen restrictions), although the implicit uncertainty surrounding the chosen sign restrictions is not acknowledged. Furthermore, Baumeister & Hamilton (2015*a*) demonstrate that the procedure as typically implemented implicitly imposes an informative Cauchy prior on the model. This informative prior is not recognized even though it can have a large impacted on estimated coefficients. They note that under standard conditions in a simple supply and demand market VAR that; 'even if one has available an infinite sample of data, any inference about the demand elasticity is coming exclusively from the prior distribution.'

Additional problems arise with ensuring inference is valid if only prior information concerning sign restrictions is used for identification as this results in a model which is only *set identified* (Kilian 2015)In this case models which rely entirely on sign restrictions produce results which are only bound within a set of possible solutions, and the point estimates for impulse response functions are not unique. Even in the case where an infinite amount of data were available to estimate even the simplest models (such as the supply and demand model discussed previously), no point estimates can be produced (Kilian & Lütkepohl 2017). This difficulty arises because there are potentially many models which satisfy the chosen sign restrictions. Without additional information or assumptions each model within the set is equally plausible and there is nothing which indicates which model is most likely to be correct.

Baumeister & Hamilton (2015*a*) propose a Bayesian SVAR approach overcoming this limitation and which explicitly accounts for the uncertainty arising from the underlying structural relationships, as well as the limited observed data. There are several benefits to adopting this methodology; first, compared to alternative approaches, it is explicit in how prior information is imposed on the impulse response functions and how valid inference is formed. It therefore overcomes the previously mentioned difficulties that the sign restricted methods as proposed by Uhlig (2005) may encounter, where informative priors are implicitly imposed and their effect is unacknowledged. Second, it provides valid inference even in the case where a system would traditionally be under-identified, allowing the impact of structural shocks to be historically decomposed. The impulse response functions arising from the model are also correctly identified, i.e., they are not only *set identified*. In a forthcoming paper, Baumeister & Hamilton (2018), apply a Bayesian SVAR approach similar to that presented here, to disentangle the relative effects of supply and demand shocks on global oil prices. This chapter is the first to apply this methodology to global agricultural commodity prices, with an application to a long series of wheat prices.

We find that the majority of observed dynamics can be explained by wheat market specific idiosyncratic supply and demand shocks. Although widely discussed, there is little evidence that increasing demand from developing countries such as India and China are driving global food prices to unprecedented heights. Similarly, much attention has been focused on the potential impact of a new class of speculators, commodity index traders, who have attracted widespread concern that their entry into agricultural markets has led to increased 'financialization' and rapid price rises in global food prices. If financialization were behind recent movements in agricultural commodity price movements, spillovers from external commodities would be expected (Tang & Xiong 2012). Within our estimated results we find no evidence of such spillovers which suggests that this was not a causal factor behind recent commodity price movements.âĂİ During 2008 shocks to precautionary demand help to explain that rapid spike that wheat prices witnessed, perhaps as a result of agents utilizing wheat futures markets to deal with expected reductions in future production.

The rest of this chapter is structured as follows. Section 3.2 provides a summary of the Bayesian approach to SVAR analysis and the identification strategy of Baumeister & Hamilton (2015*a*). Section 3.3 presents an overview of recent literature concerning agricultural commodity markets and the key factors our model wishes to capture. A description of the data sources utilized is presented in Section 3.4. The following section details how we combine prior information about the impact of factors which influence wheat prices and outlines our empirical identification strategy and Section 3.6 presents the empirical results of the model. A few concluding remarks are offered in the Section 3.7.

3.2 A Bayesian Approach to SVAR models

To analyze the influence a variety of structural factors have had on global wheat prices we adopt the standard SVAR functional form. This model captures underlying structural dynamics in a set of n endogenous structural equations:

$$\boldsymbol{A}\boldsymbol{y}_t = \boldsymbol{B}\boldsymbol{x}_{t-1} + \boldsymbol{u}_t, \tag{3.1}$$

where $\boldsymbol{y}_t$ is an $(n \times 1)$ vector of variables which are observed, $\boldsymbol{A}$ is an $(n \times n)$ matrix of coefficients capturing their contemporaneous impacts, $\boldsymbol{x}_{t-1}$ is a $(k \times 1)$ vector consisting of a constant and m lags of $\boldsymbol{y}_t$. We follow Baumeister & Hamilton (2015a) and define $\boldsymbol{u}_t$ as an $(n \times 1)$ vector of normally distributed structural shocks with variance given by the matrix $\boldsymbol{D}$. $\boldsymbol{D}$ is a diagonal matrix such that all cross-correlations equal 0, given that by definition structural shocks are uncorrelated, a feature well established in the empirical literature.

The values in the matrix of the contemporaneous impact coefficients, $\boldsymbol{A}$, and the associated matrices containing coefficients on the lagged variables, $\boldsymbol{B}_1, \ldots, \boldsymbol{B}_m$, are the parameters of interest. Once these are estimated full insight into the structural dynamics of the estimated model is available. In the empirical example presented these represent the impact that individual economic factors have on wheat prices.

In addition to the observed data, prior information exists which should be incorporated into the model to allow identification and captures expectations about the size of the coefficients of interest. In this model, prior information includes the expected size of contemporaneous impacts between the included variables, i.e., the values within the matrix of contemporaneous coefficients, $\boldsymbol{A}$, as well as the lag structure and variance of the structural shocks. These prior beliefs are represented by probability distribution functions:

$$p(\boldsymbol{A}, \boldsymbol{D}, \boldsymbol{B}) = p(\boldsymbol{A})p(\boldsymbol{D} \mid A)p(\boldsymbol{B} \mid \boldsymbol{A}, \boldsymbol{D}), \tag{3.2}$$

where $p(A)$ is the arbitrary prior probability density over the individual values of elements in the $(n \times n)$ matrix of contemporaneous coefficients, $\boldsymbol{A}$. By specifying this probability density, prior information can be transparently imposed on the coefficients. This information can take many forms, for instance, that specific contemporaneous responses are 0, fall within certain bounds, or have a nonzero probability of taking any real value. Put simply, larger values of $p(\boldsymbol{A})$ are associated with values for $\boldsymbol{A}$ which are more likely *a priori*.

Standard solutions for solving the parameter identification problem such as exclusion restrictions would be represented in this framework as a dogmatic prior. As has often been noted in the literature, exclusion restrictions are rarely theoretically justified and may be difficult to validate, and such models do not acknowledge any uncertainty surrounding this imposed information. In comparison, the methodology proposed by Baumeister & Hamilton (2015*a*) that we present here explicitly acknowledges the uncertainty arising from any prior information imposed upon the model as it is given by the prior probability distribution imposed on parameters of interest.

Prior information which takes the form of sign-restrictions on contemporaneous coefficients are also compatible with this framework. A positive or negative sign restriction would be associated with a prior, $p(\boldsymbol{A})$, with a distribution which places a probability of 0 on the specified contemporaneous response taking either a negative or positive value, respectively. Standard sign restricted VARs appear to be more general than traditional VAR models based on exclusion restrictions. The imposed restrictions can typically be motivated directly by economic theory, which is usually more straightforward than with exclusion restrictions. Kilian (2015) notes that this appearance of generality (and therefore credibility) is a misperception. Further, he states that sign restricted VAR models do not nest those based upon exclusion restrictions; by construction they are less restrictive in some dimensions and more restrictive in others. Additional difficulties arise when ensuring that inference from such models is valid. This disadvantage of the sign restriction approach occurs as a potentially large set of models satisfy the given sign restrictions. This *set identified* approach leads to a wide variety of possible conclusions all of which are equally plausible given the observed data. Without additional identifying restrictions, sign restriction priors are not enough to identify which model in the set is most probable.

The insight of Baumeister & Hamilton (2015*a*) is that if prior beliefs about $\boldsymbol{B}$ and $\boldsymbol{D}$ can be expressed using certain natural conjugate distributions then closed form analytic solutions exist which allow the formation of Bayesian inference. These natural conjugate priors have a functional form which arises naturally if they are informed by previously conducted analyses, a feature which makes them particularly appealing for empirical application. Specifically, the distributions of the prior beliefs for the matrices of lagged structural coefficients, $\boldsymbol{B}$, are given by conditional Gaussian distributions, $b_i \mid \boldsymbol{A}, \boldsymbol{D} \sim \mathcal{N}(\boldsymbol{m}_i, d_{ii}\boldsymbol{M}_i)$:

$$p(\boldsymbol{B} \mid \boldsymbol{A}, \boldsymbol{D}) = \prod_{i=1}^{n} p(b_i \mid \boldsymbol{D}, \boldsymbol{A}), \tag{3.3}$$

and

$$p(\boldsymbol{b}_i \mid \boldsymbol{D}, \boldsymbol{A}) = \frac{1}{(2\pi)^{k/2} \mid d_{ii}\boldsymbol{M}_i \mid^{1/2}} \exp[-(1/2)(\boldsymbol{b}_i - \boldsymbol{m}_i)'(d_{ii}\boldsymbol{M}_i)^{-1}(\boldsymbol{b}_i - \boldsymbol{m}_i)], \tag{3.4}$$

where $\boldsymbol{m}_i$ and $\boldsymbol{M}_i$ summarize prior beliefs surrounding the lagged coefficients in the ith structural equation. $\boldsymbol{m}_i$ is a vector containing the most likely values of the columns of the matrices of lagged structural impacts, $\boldsymbol{B}$, and $\boldsymbol{M}_i$ is a matrix that captures the certainty which surrounds these prior beliefs with larger values of $\boldsymbol{M}_i$ indicative of less certainty around the prior beliefs.

For the matrix capturing the variance of the structural disturbances, $\boldsymbol{D}$, an independent gamma distribution, $\Gamma(\kappa_i, \tau_i)$, is required for the reciprocal of the variance of each of the i structural equations:

$$p(\boldsymbol{D} \mid \boldsymbol{A}) = \prod_{i=1}^{n} p(d_{ii} \mid \boldsymbol{A}), \tag{3.5}$$

where

$$p(d_{ii}^{-1} \mid \boldsymbol{A}) = \begin{cases} \frac{\tau_i^{\kappa_i}}{\Gamma(\kappa_i)}(d_{ii}^{-1})^{\kappa_i - 1}\exp(-\tau_i d_{ii}^{-1}) & \text{for } d_{ii}^{-1} \geq 0, \\ 0 & \text{otherwise.} \end{cases} \tag{3.6}$$

d_{ii} is the element of $\boldsymbol{D}$ appearing in the ith row and ith column. κ_i and τ_i are parameters which capture the prior information surrounding the variances of the structural relationships.[1] κ_i/τ_i represents the prior expectation of d_{ii}^{-1}, with a distribution having a variance

[1] In general terms nothing prevents κ_i and τ_i being functions of $\boldsymbol{A}$, however the resulting posterior

given by κ_i/τ_i^2. Larger values of κ_i and τ_i are indicative of higher confidence in the prior beliefs represented by these priors.

Baumeister & Hamilton (2015*a*) provide a methodology which allows the generation of draws from the posterior distribution, $p(\boldsymbol{A}, \boldsymbol{D}, \boldsymbol{B} \mid \boldsymbol{Y}_t)$, using Markov Chain Monte Carlo methods. Specifically, a form of the Metropolis-Hastings algorithm is used to select the next multi-dimensional sample. Any remaining uncertainty surrounding the parameters given the information provided by the data, $\boldsymbol{Y}_t = (\boldsymbol{y}_1', \boldsymbol{y}_2', \ldots, \boldsymbol{y}_T)'$, will be revealed, providing a posterior distribution for the parameters of interest. In addition, they provide recommendations for values of κ_i, $\tau_i(\boldsymbol{A})$, $\boldsymbol{m}_i(\boldsymbol{A})$ as well as the uncertainty matrix, $\boldsymbol{M}_i$, which we follow in our empirical application. Applying this methodology to the global wheat market allows valid inference to be gained from impulse response functions and the associated historical decompositions of interest. Before discussing the precise parameterization of the model, we now review the potential structural factors in agricultural commodity markets and the mechanism through which they impact wheat prices.

distributions become more analytically complex.

3.3 Review of Structural Factors

Many different economic factors have been proposed as underlying recent wheat price dynamics. In the wake of the recent spike increased attention has been focused on the correct policy response to counteract the negative outcomes of rapid price rises. The model we adopt attempts to parsimoniously include those factors which are most important and widely discussed in the literature. The historical contribution that each has had on the global wheat price can then be decomposed, revealing the relative importance that each has had to wheat price movements over time.

A similar structural analysis was carried out by Janzen et al. (2014), who identify global economic activity, external commodity comovement and precautionary demand as critical structural economic factors that should be included within the model. These economic factors capture the impact of the key drivers thought to be behind recent agricultural price movements. The influence of increased demand by India and China will be captured by global economic activity, potential speculation by external commodity comovement. In addition a precautionary demand shock captures demand arising from market participants expectations of future supply and demand conditions - a vital function of the price discovery process. Finally, wheat market specific supply and demand shocks are captured by the residual variation left unexplained by other factors. This final factor effectively captures net supply shocks to the wheat market. In common with many agricultural commodities, once planted various supply shocks, for instance weather shocks like droughts or excess precipitation, can reduce wheat supply and cause price rises.

Although the model captures the same four structural shocks as Janzen et al. (2014), we utilize a modified set of variables which extend over a longer timespan. Additionally, a methodology which does not rely on the change of volatility of wheat prices for identification purposes is adopted. The remainder of this section reviews each of the shocks that the model captures and the mechanism through which they impact wheat prices.

3.3.1 Global Economic Activity

Many commentators have suggested that global economic activity may have impacted on agricultural commodity prices. Increased demand from developing economies due to their economic growth is often suggested to have pushed up global demand for agricultural products and led to higher prices. The transmission mechanism for this movement comes from the following directions. First, economic growth drives up input prices, such as energy and other raw materials, leading to higher costs of production and transportation for agricultural commodities (Janzen et al. 2014).

Second, the transmission mechanism most commonly discussed, is that increases in real economic activity increase global incomes which raises demand for basic foodstuffs, particularly in developing countries (Rosegrant 2008). A further, higher order, effect of increasing wages is changing dietary patterns towards a more 'westernized' diet which has a higher proportion of meat in it (Shewry & Hey 2015). As global meat consumption increases, primary crops, such as wheat, are shifted from human consumption into animal feed. For an equivalent calorific portion of meat, much more wheat is used, resulting in an increased overall demand for wheat and other commodities used as feed. In the long run there is some evidence that this trend will be reversed and that consumers who reach increasingly high levels of income begin to reduce their meat consumption due to health or environmental concerns (see Vranken et al. (2014) for a discussion of this relationship).

There is relatively sparse evidence on the actual influence that global economic activity has had on agricultural commodity pricing. McPhail et al. (2012) find that global demand is a significant contributing factor to corn price variability, peaking at 16% after six months in variance decompositions. However, Headey & Fan (2008) find little evidence that global demand based factors have pushed up food prices. A further difficulty with this proposed causal factor is that both India and China are largely self-sufficient in wheat consumption (Timmer 2008), so the mechanism for their impact on global wheat prices cannot be through increased upward pressure on prices caused by increasing imports. Although this does not rule out increasing demand from sub-Saharan Africa or the changes in dietary patterns in these countries having an effect on wheat prices.

3.3.2 External Commodity Comovement

Critics of agricultural commodity market speculation contend that agents who are not concerned solely with price discovery (the price which is justified by underlying fundamentals) are disrupting the correct functioning of markets. Over recent years, increasing numbers of traders have entered agricultural commodity markets, leading to a process of 'financialization' (Cheng & Xiong 2014), mainly in the form of institutional index investors wanting to gain exposure to a basket of agricultural commodities. Therefore, a large amount of attention has been paid to the influence of Commodity Index Traders (CITs), and any impact they have had on wheat prices, as well as agricultural commodity markets in general. The potential for these CITs to have unduly increased prices has led to increased scrutiny of food markets and led to several proposed policy responses.

Excess comovement amongst agricultural commodities is a well-known phenomenon (see, for instance, the seminal paper of Pindyck & Rotemberg (1990)). A more recent phenomenon is the increased comovement with external commodities, such as industrial commodities like oil. Tang & Xiong (2012) propose that this is a mechanism by which the impact of CITs can be captured, as the correlation between agricultural commodities and other external commodities increased as CITs became established. The effect found for those included in agricultural indices, the instrument investors use to expose themselves to agricultural commodity markets, was larger than for those not listed, similar to the well-known equity index inclusion effect (Barberis et al. 2005).

This external commodity comovement shock captures the extent to which wheat prices comove with an external commodity. To capture this shock a global oil price series (see Section 3.4 for precise details) is included as one of the variables within the SVAR. This variable captures excess comovement which cannot be explained by the other economic factors, such as demand shocks from changes in global economic activity. If there is evidence of increased comovement with an external non-agricultural commodity, it would suggest that CITs had impacted upon wheat prices.

3.3.3 Precautionary Demand

Although speculation is typically seen as having a negative influence on agricultural commodity markets, with a destabilizing effect on commodity prices, it has an important role to play in price determination. Storable agricultural commodities allow for investments in futures markets which can reduce the risk for participants and reduce volatility by minimizing the impact that fundamental shocks have over time. In agricultural commodities, these fundamental shocks are normally shocks to supply from weather and other related phenomena which reduce yields. Including a factor which captures the precautionary demand motive allows for the disentanglement of demand shocks for purely speculative purposes (which are potentially destabilizing) and investments in anticipation of future fundamental conditions which will drive food prices in the future.

The mechanism that standard models of storable commodities in competitive markets allow market participants expectations of future supply and demand conditions to influence prices today is through stockholding behavior (Janzen et al. 2014). If agents expect that discounted future prices are above the current cash price they increase the size of their holding. Conversely, if discounted future prices are below the current price inventories are sold. Buying and selling based on expected future market conditions lead to a source of demand that is entirely focused on the future, hence 'precautionary demand.'

The most straightforward way of capturing precautionary demand would be through the current holdings of wheat. However, this data is not available at either the scale (i.e., global wheat stocks) or the frequency (data is unavailable at a monthly level) to be suitable for use in this analysis. Instead, we rely on the well known link between futures spreads and inventories known as the 'Working Curve' (Working 1933), which relies on the observation that if futures prices at time t for delivery in n periods is an unbiased estimator of the spot price $t + n$ then the spread is positively related to inventories. This 'Working Curve' relationship allows for a variable to be constructed which can distinguish shocks to precautionary demand.

3.3.4 Net Supply

The final shock the model captures is wheat market specific supply factors. Wheat, like most agricultural commodities, has well defined growing seasons so that during any particular crop year supply is effectively perfectly inelastic. Any production shocks will manifest themselves as increases in wheat prices due to the inability of farmers to quickly react and produce additional crop. In the empirical example which follows production shocks are modeled as *reductions* in net supply, which lead to *higher* prices. The majority of wheat production is based in a few specific regions of the world such that local weather variations can have an outsized impact on global commodity prices. Changing weather patterns have meant that yield increases have stalled in major producing countries such as Australia (Hochman et al. 2017), and are likely to continue to be impacted into the future.

During the period before and immediately after 2008 there was a sequence of noteworthy events which are likely to have impacted global wheat supply. Australia, having had static yields for the best part of two decades began the 2007 planting season experiencing the second year of severe drought. Subsequently, there was a reduction in wheat yields of 57% for 2007, even though planted area fell by just 5% (Australian Bureau of Statistics 2008). Similarly, in Russia, another major wheat exporter, production fell in 2010 due to drought, leading to the imposition of an export ban. This resulted in a reduction of wheat exports from 18.6 to 4 million metrics tons from 2009/10 to 2010/11 (Vocke 2012). These primarily weather based factors, which are intrinsic to the volatility of the global wheat market, are captured within the idiosyncratic net supply shock. This shock consists of the residual variation which is not explained by the other included economic factors and is captured by the structural shock associated with wheat prices, the final included variable.

3.4 Data

As we wish to identify the impacts of four identified structural shocks, the model requires a system of four simultaneous equations in the form of Equation 3.1, and four observed variables. As previously mentioned we wish to identify the impact of the three structural shocks from economic activity, comovement and precautionary demand. The final variable, wheat prices, captures residual variation in the form of idiosyncratic net supply shocks.

As the shocks to real economic activity which we wish to capture are principally focused on developing economies, any variable used to capture these shocks must incorporate global changes in real economic activity, rather than any proxy measure based on a particular region (e.g., macroeconomic conditions in the United States). Due to the inappropriateness of using such variables, and the lack of any reliable direct measure of global GDP, proxy measures are required. The Index of Real Economic Activity as proposed in Kilian (2009) is utilized for the analysis. This index measure is based upon the empirically established link between industrial production and ocean-going freight rates. As international shipping prices react to underlying changes in global economic conditions, they are not biased to any one particular region (or even direction of trade). This measure has been successfully used in several studies looking primarily at the oil market (e.g., Baumeister & Kilian (2012)). For a discussion of the benefits of this measure, particularly for modeling commodity market demand, see Kilian & Zhou (2018), which concludes that the Index of Real Economic Activity is the most suitable for this purpose when compared to similar measures which are based on steel production and other related metrics.

If fears of 'financialization' are correct, wheat prices will have been increasingly driven by prices of non-agricultural commodities over time. As Tang & Xiong (2012) show, if this hypothesis is correct, prices of food crops should increasingly correlate and comove with oil prices towards the end of the sample period under investigation. To capture excess external commodity comovement the oil price for West Texas Intermediate (WTI) is included as one of the variables. This is collected from the Federal Reserve Economic Data (FRED) database and deflated by the U.S. consumer price index, this series has been widely used (see, for example, Baumeister & Hamilton (2018)).

To capture market participants incentives to hold physical inventories of wheat due to precautionary demand we must make use of a proxy due to the lack of suitable direct measures. The proxy measure is the spread between the fifth and nearby futures contracts from the Chicago Board of Trade (CBOT). The fifth futures contract is the contract that has the maturity date which has the fifth smallest time to maturity out of all traded contracts (as opposed to the nearby contract which is the closest to maturity). This is based on the well-known 'Working Curve' (Working 1933) which captures the robust relationship between spreads and storage. This relationship has been recently demonstrated as still being an accurate characterization of wheat market structure (Joseph et al. 2016). It captures the incentive that market participants have to hold onto physical inventories until future supply enters as new crops are harvested. The wheat market spread captures shocks to expectations about wheat market supply and demand conditions in a future cropping period.

For the final variable of interest, we include the CBOT near futures wheat price. As is typical for analysis of agricultural commodity market price series, the nearby futures price is used as a proxy for spot prices due to particular issues with spot price analysis. Firstly, there has been increased non-convergence between futures prices at expiration and the cash price during the period of most interest (Garcia et al. 2015). To avoid potential complications arising from the use of cash prices we follow other analyses (e.g., Gutierrez (2013)) and utilize the near futures price series. This final variable captures any remaining variation not captured by the previous three shocks and therefore represents idiosyncratic net supply shocks.

The data are collected at a monthly frequency beginning in January 1978, the first date at which all data are available, until December 2014, leading to a total of 433 monthly observations for each variable. This time span is sufficiently long enough to capture recent changes in the underlying structural relationships between variables, especially the period leading up to the crisis in 2008. Both oil and wheat price series enter as the model as first (log) differences.

3.5 Empirical Strategy

3.5.1 Model Description

We adopt a similar structural framework to Janzen et al. (2014) with one key modification: we allow shocks to the external commodity variable to impact global economic activity. This allows us to analyse the same structural relationships and shocks, whilst not constraining shocks to the external commodity to having no effect real economic activity. This alteration is made due to using oil as an external factor, of which there is an extensive literature which identify the size of the effect that oil price supply and demand shocks have on the real economic activity (see for instance Kilian (2009), Baumeister & Hamilton (2018). Any external commodity can be expected to impact on global real economic activity when subject to unanticipated shocks. Figure 3.2 depicts the underlying structural dynamics that we wish to capture, with the additional inclusion of external commodity market to real economic activity spillovers marked in red.

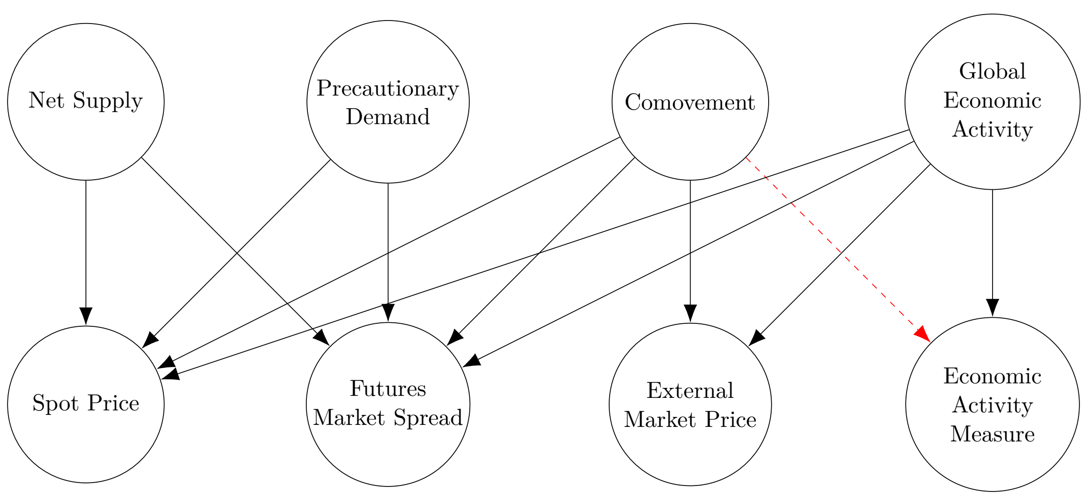

Figure 3.2: Diagrammatic representation of how unanticipated structural shocks (top row) are expected to impact on observed variables included in structural model (bottom row). Additional transmission mechanism from oil to real economic activity is highlighted in red. Adapted from the US agricultural department report of Janzen et al. (2014).

We now apply the methodology outlined in Section 3.2 to the global wheat market. We implement a parsimonious four variable structure. The four monthly variables are represented in the in the vector $\boldsymbol{y}_t$: (i) real economic activity, rea; (ii) the real price of the external commodity, ext; (iii) the spread between the fifth futures contract and the near for wheat, spr; and (iv) the near price of wheat futures, pri. The system therefore consists of a system of four reduced form equations:

$$y_t^{rea} = \alpha_{11} y_t^{rea} + \alpha_{12} y_t^{ext} + \alpha_{13} y_t^{spr} + \alpha_{14} y_t^{pri} + \boldsymbol{\beta}' \boldsymbol{x}_{t-1} + u_t^{rea} \tag{3.7}$$

$$y_t^{ext} = \alpha_{21} y_t^{rea} + \alpha_{22} y_t^{ext} + \alpha_{23} y_t^{spr} + \alpha_{24} y_t^{pri} + \boldsymbol{\beta}' \boldsymbol{x}_{t-1} + u_t^{ext} \tag{3.8}$$

$$y_t^{spr} = \alpha_{31} y_t^{rea} + \alpha_{32} y_t^{ext} + \alpha_{33} y_t^{spr} + \alpha_{34} y_t^{pri} + \boldsymbol{\beta}' \boldsymbol{x}_{t-1} + u_t^{spr} \tag{3.9}$$

$$y_t^{pri} = \alpha_{41} y_t^{rea} + \alpha_{42} y_t^{ext} + \alpha_{43} y_t^{spr} + \alpha_{44} y_t^{pri} + \boldsymbol{\beta}' \boldsymbol{x}_{t-1} + u_t^{pri} \tag{3.10}$$

Which can be written in the familiar structural system:

$$\boldsymbol{A} \boldsymbol{y}_t = \boldsymbol{B} \boldsymbol{x}_{t-1} + \boldsymbol{u}_t \tag{3.11}$$

where $\boldsymbol{x}_{t-1} = (\boldsymbol{y}'_{t-1}, \boldsymbol{y}'_{t-2}, \ldots, \boldsymbol{y}'_{t-m}, 1)'$ and $\boldsymbol{u}_t$ is a vector of the four structural shocks $(\boldsymbol{u}_t^{rea}, \boldsymbol{u}_t^{ext}, \boldsymbol{u}_t^{spr}, \boldsymbol{u}_t^{pri})$. The number of lags included in the system, m, is set as 3 months for comparability with similar analyses. Other studies have used similar lag specifications, e.g., Janzen et al. (2014) utilize $m = 3$, Baumeister & Hamilton (2018) utilize $m = 4$, with $m = 4$ for this study the results are qualitatively the same. As the logarithm of the price series enter the structural model elements in the matrix of contemporaneous coefficients, $\boldsymbol{A}$ (denoted a_{ij} for the element in the ith row and jth column), can be interpreted as elasticities.

3.5.2 Model Parameters

Many researchers will have strong prior beliefs about several parameters in the system characterized by Equations (3.7)-(3.10) and typically will assume values when conducting analysis. We first normalize the elements in the matrix of contemporaneous coefficients, $\boldsymbol{A}$, by setting each element in the diagonal equal to 1, capturing the relative effects of structural shocks to each of the observed variables. This normalization explicitly results in all shocks producing increases in wheat prices, meaning that a net supply shock (one which impacts the wheat price variable) are, by construction, restrictions in the supply of wheat.

Following previous studies such as Janzen et al. (2014), which utilize ordering restrictions or exclusion restrictions in their identification strategy, we also set several elements of $\boldsymbol{A}$ equal to 0. Instead of forming an exclusion restriction, within this Bayesian framework this is explicitly representative of a strong prior belief that there is no short-run contemporaneous effect of certain variables upon each other. The short run impact here is within the observable frequency of the data, in this case a contemporaneous response is one which occurs within a single month. For the wheat market model the upper right quadrant is set equal to 0, corresponding to a strong prior belief that neither precautionary demand shocks or wheat market specific supply and demand shocks have any observable short-run impact upon either real economic activity or global oil prices. This has the additional benefit of reducing the parameter space that must be estimated by the Metropolis-Hasting Random Walk algorithm. The remaining elements of $\boldsymbol{A}$ that we do not have dogmatic prior beliefs about (a_{ij}) form our parameters of interest:

$$\boldsymbol{A} = \begin{bmatrix} 1 & a_{12} & 0 & 0 \\ a_{21} & 1 & 0 & 0 \\ a_{31} & a_{32} & 1 & a_{34} \\ a_{41} & a_{42} & a_{43} & 1 \end{bmatrix} \tag{3.12}$$

We represent prior information around each of the elements of $\boldsymbol{A}$ using a Student t distribution, with mode c_a, scale σ_a, and ν_a degrees of freedom. The magnitude of oil price shock impacts on real economic activity (α_{12}) is expected to be small, due to low ratio of oil expenditures to total GDP (Hamilton 2013). We follow Baumeister & Hamilton (2018) and represent this with a negatively truncated Student t prior distribution with mode -0.05, scale parameter $\sigma = 0.1$ and 3 degrees of freedom.

Although there has been a large body of work looking at the impacts of a wide variety of causes behind the movements in wheat prices, there is surprisingly little consensus on the quantitative magnitude of these impacts. This makes the specification of the priors in the matrix of contemporaneous coefficients more difficult. For instance, a large body of work has investigated the impact of oil prices on wheat which should influence our choice of prior capturing the effect of external commodity comovement on wheat (a_{42}). However, very few studies report elasticities which we can be used to guide the choice of a prior distribution. Using annual data, Baffes & Dennis (2013), find statistically significant estimates of the elasticity of wheat prices with respect to oil prices ranging from 0.22 to 0.28. In a more recent study Fernandez-Perez et al. (2016) find an effect of 0.0795, however when they try and find the difference between high oil price regimes, and low ones, the effects are statistically insignificant. Similarly, the income elasticity for international wheat prices has been estimated as -0.14 and -0.12, albeit without statistical significance, by Baffes & Dennis (2013), 0.001 by Ai et al. (2006) and also insignificantly estimated as 0.03 by Frankel & Rose (2010).

Therefore, for all of the other potential contemporaneous effects, for which there is apparently very little prior information available in the literature, we adopt a neutral view on the contemporaneous impacts between the variables. We utilize prior distributions which capture our reasonable prior belief that as these coefficients capture elasticities; they should be relatively small and centered at 0, we represent these prior beliefs with t distributions with mode 0, scale parameter $\sigma = 0.3$ and 2 degrees of freedom. Table 3.1 details the specification of each element in $\boldsymbol{A}$ which is used for the rest of our empirical investigation. These priors have the standard characteristics of a t distribution; they have a finite mean, infinite variance and the majority of the probability mass is focused around 0. Each column details the specified mode, c_a, the standard deviation, σ_a, the degrees of freedom , ν_a. The final column details whether the t distribution is truncated such that its values are negative, (-), positive (+) or has no truncation ( ). Each of the prior distributions are represented as solid red lines in Figure 3.3.

For priors for the coefficients capturing the impact of lagged values, $b_i \mid \boldsymbol{A}, \boldsymbol{D} \sim \mathcal{N}(\boldsymbol{m}_i, d_{ii}\boldsymbol{M}_i)$ we use the specification of Baumeister & Hamilton (2015a), who implement a prior adapting the well known Minnesota prior of Doan et al. (1984) with a modification of the specification proposed by Sims & Zha (1998). Effectively our prior expectation is that each of the variables follows a random walk process, achieved by setting the elements of $\boldsymbol{m}_i$ corresponding to the first lag equal to the value of the ith row of $\boldsymbol{A}$ and all other elements to 0. This prior expectation can be justified on the basis that if markets are (at

	c_a	σ_a	ν_a	Sign Restriction
a_{12}	-0.05	0.1	3	(-)
a_{21}	0.0	0.3	2	()
a_{31}	0.0	0.3	2	( )
a_{32}	0.0	0.3	2	( )
a_{34}	0.0	0.3	2	( )
a_{41}	0.0	0.3	2	( )
a_{42}	0.2	0.3	2	( )
a_{43}	0.0	0.3	2	( )

Note: for each variable we specify a prior mode, c_a, variance, σ_a, degrees of freedom in the Student t distribution, ν_a and whether the distribution is truncated to be positive (+), negative (-) or no sign restriction is imposed ( ).

Table 3.1: Specification of prior distributions for unknown elements in $\boldsymbol{A}$

least weakly) efficient, then no other additional factor can help predict future movements, leading to variables that follow a random walk process. The diagonal matrix which captures the variance, $\boldsymbol{M}_i$, has progressively smaller values applied to elements corresponding to lags of a higher order. The choice of the value of τ_i is dependent upon the scale of the data, and particularly the prior mode of $p(\boldsymbol{A})$. This captures the well known stylized fact that the time series properties of commodities exhibit a great deal of persistence (Deaton & Laroque 1992).

For the final parameters required for implementation of the Metropolis-Hastings random walk we follow the recommendations of Baumeister & Hamilton (2015*a*,*b*) and call on standardized values. We set λ_0, which captures the overall confidence in the prior equal to 0.2. λ_1 captures the confidence on lags of higher order than 1, represents how quickly our expectation that very distant lags will no longer have an effect on the current observation (i.e., the rate at which the prior for these lagged coefficients tends towards 0 as the lag order increases). For reference, setting $\lambda_1 = 0$, would mean giving all m lags in the model an equal weighting. Finally, we set $\lambda_3 = 100$ and make the prior on the constant term in the four structural equations irrelevant.

3.6 Empirical Results

We now turn to the results of the empirical application of the model. All results are generated using 200000 draws, with a burn-in sample of the first half of the draws, following the derivation of the random walk Metropolis-Hastings algorithm to generate draws from $p(\boldsymbol{A} \mid \boldsymbol{Y_t})$ presented in Baumeister & Hamilton (2015a). One beneficial aspect of this Bayesian implementation of SVAR analysis is that graphical representations of the estimated coefficients are available, and demonstrate to what extent parameters are being driven by the imposed prior distributions. Posterior distributions for the eight contemporaneous coefficients are plotted as solid blue histograms in Figure 3.3, with prior distributions represented as solid red curves. For seven of the eight parameters of interest the data are highly informative (as evidenced by relatively tight posterior distributions) but relatively uninformative about the impact of wheat prices on wheat spreads (a_{34}). Figure 3.3 demonstrates that our prior beliefs about the coefficients can be revised substantially in this context and coefficients (and therefore the associated IRFs and historical decompositions) are not being driven by the imposed prior distributions. After the SVAR has been estimated, simulations are conducted to generate valid impulse response functions and historical decompositions from the fitted lag structure.

Figure 3.4 plots impulse-response functions. Median values from the posterior distribution are plotted as solid lines for each time horizon. Shaded regions represent 95% posterior credibility regions which include the uncertainty generated from having access to only a limited set of data and that arising from underlying uncertainty about the true structural model. Each subplot summarizes the response of each variable to one of the four shocks and traces the response path out over 20 months. Due to the nature of the normalization used a net supply shock causes wheat prices to rise, and is therefore indicative of a reduction in the overall supply of wheat. This captures the empirical fact that shocks to wheat supply are reductions in supply, and therefore tend to increase the price of wheat. These graphs allow us to ascertain what the response of each variable is to each shock and providing a check on the validity of the utilized prior, indicated by the response to shocks occurring in the correct direction.

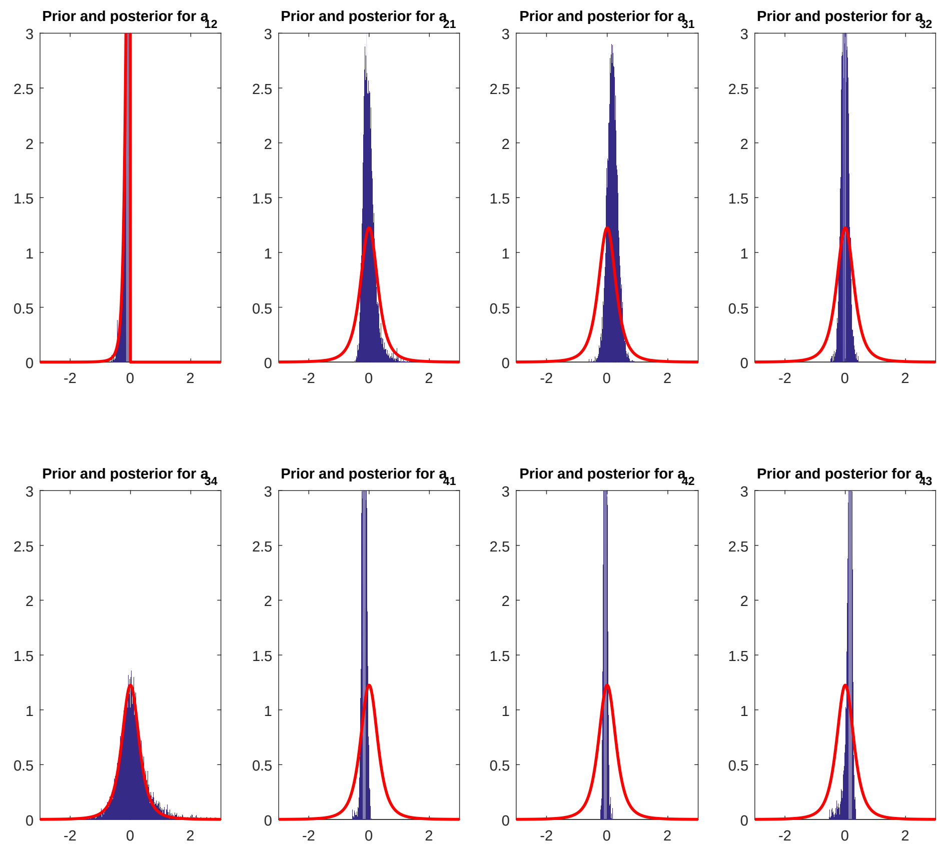

Figure 3.3: Prior and Posterior Distributions of elements of $\boldsymbol{A}$

The final row of the panel presents the response of wheat prices to shocks to global economic activity, external commodity price (comovement), precautionary demand and net supply. Contrary to many expectations there is a relatively muted response to each of the first three shocks and a highly persistent response to wheat net supply shocks. Shocks to global economic activity initially cause wheat prices to rise marginally; however, they quickly return to baseline levels. This suggests that the principal factor behind the majority of observed wheat dynamics during our sample period, which includes the rapid rise up to 2008, are wheat market specific, idiosyncratic shocks. Historical decompositions allow us to examine whether there have been any periods when any of the factors became more or less important in wheat price dynamics.

Figure 3.5 decomposes wheat price movements into contributions by different structural shocks from 1983 until 2014. Red dashed lines in the panel denote the realized deviation of wheat price from its mean over the entire sample.[2] The solid blue line in each subplot is the calculated impact of the real economic activity, comovement, precautionary demand and net supply shocks, respectively. 95% posterior credibility regions are represented by the shaded areas - Baumeister & Hamilton (2015*b*) note that these are rarely presented in the partially-identified VAR literature - giving additional insight into the impact and importance of each factor. The vast majority of wheat price dynamics are explained as shocks to net supply, with neither real economic activity or external commodity comovement driving price movements. This is particularly evident during the 2008 price rise where neither of these two factors appears to be playing a role in driving prices higher. However, there is some evidence from the third subplot that precautionary demand shocks may have contributed to the rapid price rise seen during this period. The majority of historical price movements over the past 30 years have been driven by idiosyncratic wheat market specific supply and demand factors.

Average contributions from individual shocks are computed using variance decompositions. Table 3.2 presents the fraction of the mean-squared error (MSE) of a 12 month ahead forecast that shocks to real economic activity, the external commodity price, precautionary demand and net supply contribute to each of the four observed variables. 95% credibility intervals are reported below. The fourth row presents the results for wheat prices. 72% of the variance of wheat prices in our sample period is explained by net supply shocks, and the majority of the remaining variation is contributed by shocks to precautionary demand. These results are qualitatively similar to the findings of other studies. For instance, variance decompositions during a forecasting exercise conducted by Delle Chiaie et al. (2017) also find that 72% of wheat price variability is idiosyncratic. McPhail et al. (2012) find 64% of variability for corn market prices after six months is explained by corn market specific shocks. Both real economic activity and excess commodity comovement are significantly less important in the variability of wheat prices, explaining 5% and 2% respectively. This second result suggests that external factors were not driving prices. This is especially important given the widely held fear that 'financialization' had led to speculation on the part

[2] Appendix 3.8 presents equivalent diagrams decomposing the remaining observed variables.

of external commodity index funds which drove prices higher. If such activity were driving global wheat prices, it would be expected that a larger portion of wheat price variability would be explained by shocks to the external commodity price.

Although not the focus of this chapter, results are also reported for the decompositions of real economic activity, the external commodity and the wheat price spread. The decomposition of the external commodity, in this application oil prices, are of interest as there is a large literature devoted to the deconstruction of shocks that impact world oil prices. Of particular note is the relatively small portion of the WTI oil price that appears to be driven by shocks to global economic activity, seen as demand shocks, which within this model explain just 5% of the variability of oil price, suggesting factors other than heightened demand from real economic activity drive oil prices.

This result is contrary to previous evidence suggesting that shocks to global economic activity drive real oil prices (see for instance, Kilian (2009)). Specifically, the period around 2003, which Kilian (2009), ascribe oil price rises as a response to unanticipated global demand shocks, does not occur significantly within Figure 3.7, suggestive that alternative explanations may be behind such price changes, for instance oil supply, which is not captured by this model. Additionally, contrary to previous findings where real economic activity shocks result in large, significant and long-lasting effects on oil prices, the response within this model is muted and does not persist. This result supports recent evidence from Baumeister & Hamilton (2018), which uses a similar Bayesian SVAR as that used here and find qualitiatively similar result: shocks to demand appearing to be relatively less important than shocks to supply.

Tables 3.3 and 3.4 present variance decompositions for 6 month and 18 month horizons, respectively. These results allow further examination into the relative importance of the investigated structural shocks to wheat prices and whether these relationships are robust over different horizons. We can see that the estimated contributions made by each of the variables are relatively stable over time. Examining the fourth row in Tables 3.2 - 3.4, shows that there is some variation, for instance, shocks to global economic activity appear to strengthen marginally over time. However, all of the results lie within overlapping credibility intervals indicating that the variance contributed by each shock is the same at different

Table 3.2: Decomposition of variance of 1 year ahead forecast errors

	Global Economic Activity	Comovement	Precautionary Demand	Net Supply
Economic Activity Measure	**0.9350** (0.7686, 0.9917)	**0.0458** (0.0065, 0.1458)	**0.0125** (0.0011, 0.0550)	**0.0067** (0.0007, 0.0306)
External Market Price	**0.0415** (0.0073, 0.0963)	**0.9110** (0.7822, 0.9848)	**0.0154** (0.0024, 0.0496)	**0.0321** (0.0055, 0.0719)
Futures Market Spread	**0.0149** (0.0023, 0.0488)	**0.0110** (0.0019, 0.0350)	**0.9469** (0.7453, 0.9899)	**0.0272** (0.0059, 0.1710)
Nearby Price	**0.0455** (0.0161, 0.0790)	**0.0179** (0.0051, 0.0483)	**0.2086** (0.0452, 0.2765)	**0.7280** (0.5962, 0.9336)

Table 3.2 gives the estimated fraction of the 1 year ahead mean squared forecast error contributed by each shock to each variable in bold. Values in parentheses present 95% credibility intervals.

horizons.

Overall the results from the model estimated demonstrate that the majority of wheat market price movements during the recent past can be explained by supply and demand factors specific to the wheat market. There is little evidence that two factors that were subject to a great deal of speculation, increasing demand from developing economies and speculation by commodity index traders entering markets, have had any notable impact on wheat prices over recent times. The impacts of global economic activity, external commodity comovement and precautionary demand shocks are dwarfed by the impact that this factor has - although during 2008 there is some evidence that shocks to precautionary demand may have resulted in wheat prices reaching their crisis levels.

Table 3.3: Decomposition of variance of 6 month ahead forecast errors

	Global Economic Activity	Comovement	Precautionary Demand	Net Supply
Economic Activity Measure	**0.9355** (0.7854, 0.9938)	**0.0532** (0.0054, 0.1659)	**0.0072** (0.0004, 0.0298)	**0.0041** (0.0004, 0.0190)
External Market Price	**0.0375** (0.0049, 0.1007)	**0.9438** (0.8421, 0.9922)	**0.0064** (0.0007, 0.0234)	**0.0123** (0.0021, 0.0338)
Futures Market Spread	**0.0072** (0.0012, 0.0234)	**0.0070** (0.0011, 0.0245)	**0.9626** (0.7564, 0.9932)	**0.0232** (0.0045, 0.1975)
Nearby Price	**0.0379** (0.0099, 0.0690)	**0.0122** (0.0035, 0.0385)	**0.1937** (0.0237, 0.2739)	**0.7562** (0.6186, 0.9629)

Table 3.3 gives the estimated fraction of the 6 month ahead mean squared forecast error contributed by each shock to each variable in bold. Values in parentheses present 95% credibility intervals.

Table 3.4: Decomposition of variance of 18 month ahead forecast errors

	Global Economic Activity	Comovement	Precautionary Demand	Net Supply
Economic Activity Measure	**0.9317** (0.7606, 0.9910)	**0.0438** (0.0065, 0.1316)	**0.0171** (0.0016, 0.0758)	**0.0074** (0.0009, 0.0320)
External Market Price	**0.0432** (0.0080, 0.0993)	**0.9046** (0.7743, 0.9824)	**0.0185** (0.0034, 0.0515)	**0.0337** (0.0061, 0.0748)
Futures Market Spread	**0.0227** (0.0030, 0.0764)	**0.0115** (0.0022, 0.0332)	**0.9383** (0.7321, 0.9885)	**0.0275** (0.0063, 0.1582)
Nearby Price	**0.0495** (0.0180, 0.0875)	**0.0185** (0.0055, 0.0490)	**0.2106** (0.0483, 0.2756)	**0.7213** (0.5879, 0.9281)

Table 3.4 gives the estimated fraction of the 18 month ahead mean squared forecast error contributed by each shock to each variable in bold. Values in parentheses present 95% credibility intervals.

3.7 Conclusion

In this chapter, we present an evaluation of the relative magnitude of the impact of the most widely discussed economic factors which drive global wheat price dynamics. To do so, we construct a structural model including four of the most important economic factors; global economic activity, speculation pressures through external commodity comovement, precautionary demand and net supply shocks. We utilize a recently developed Bayesian structural econometric model to provide valid inference. We apply this model to a long dataset which encompasses the recent volatile price movements leading up to 2008 and the subsequent period of raised prices and higher volatility.

Forming inference from structural models is impossible without imposing some prior information onto the model, which we may not know with any certainty. This chapter is the first to our knowledge to implement a Bayesian SVAR model with incomplete prior information and apply it to agricultural commodity markets. This methodology has two key advantages; first, any prior information which is informative to our results has been explicitly acknowledged, and the uncertainty surrounding such information accounted for, second, we can form valid inference from this scenario which is traditionally seen as under-identified. For our application to agricultural markets, this strategy allows more insight to be gained into the relative impact that economic factors have had on prices, in a transparent fashion.

We find that the majority of wheat price movements over the past three decades can be explained by wheat market specific net supply shocks. During 2008, a period of particular concern, there is some evidence that precautionary demand shocks may have contributed to wheat price rises, perhaps as a result of agents utilization of the futures markets to deal with expected future supply tightness. Although widely discussed, given that shocks to global real economic activity appear to have minimal contributions to global food price rises, there is little evidence that increasing demand from developing countries such as India and China are driving global food prices to unprecedented heights. Similarly, much attention has been focused on the potential impact of a new class of speculators, commodity index traders, who have attracted widespread concern that their entry into agricultural markets has led to increased ‘financialization’ and rapid price rises in global food prices. The spillovers from an external commodity market which would indicate such are not present within our estimated results.

3.8 Appendix A

This appendix contains the historical decompositions of the other three variables contained within the system, global real economic activity, external commodity price and wheat market spread.

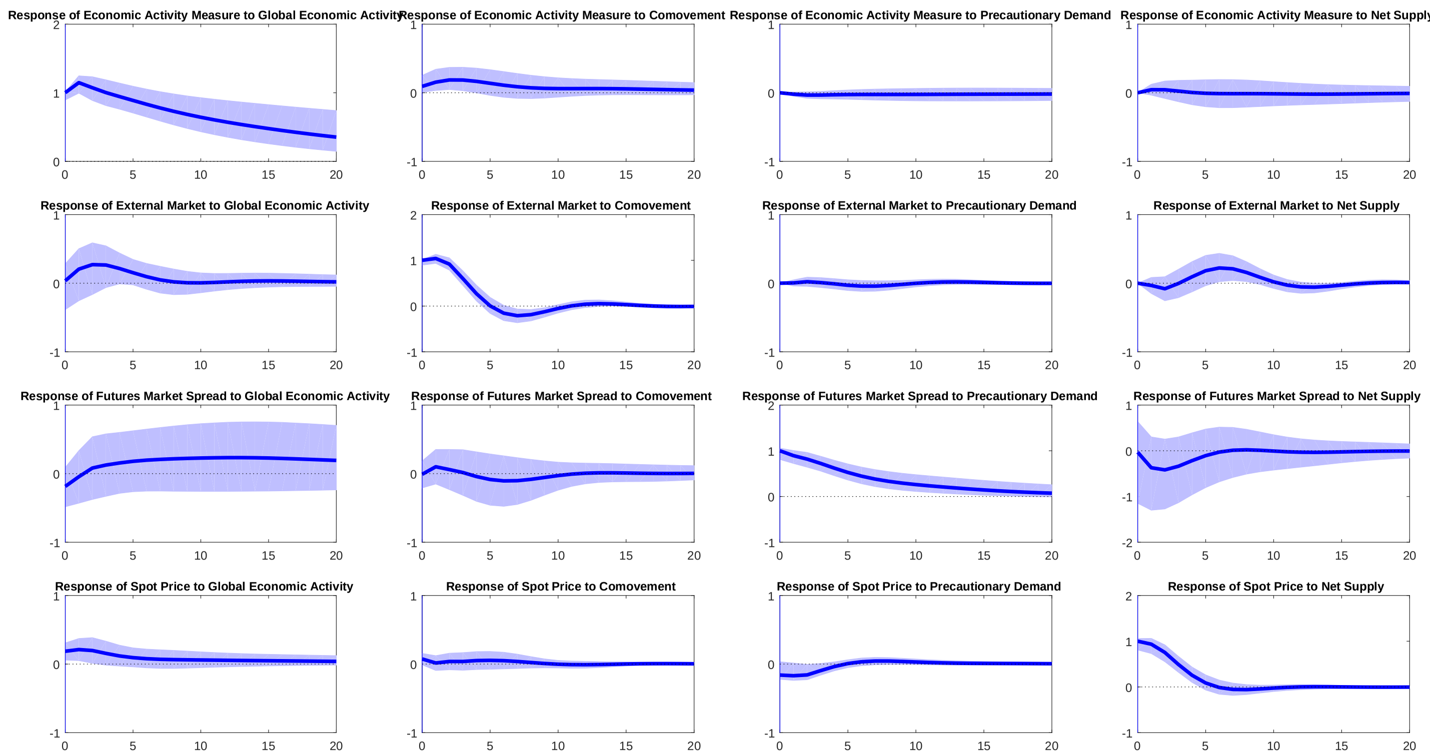

Figure 3.4: Impulse Response Functions

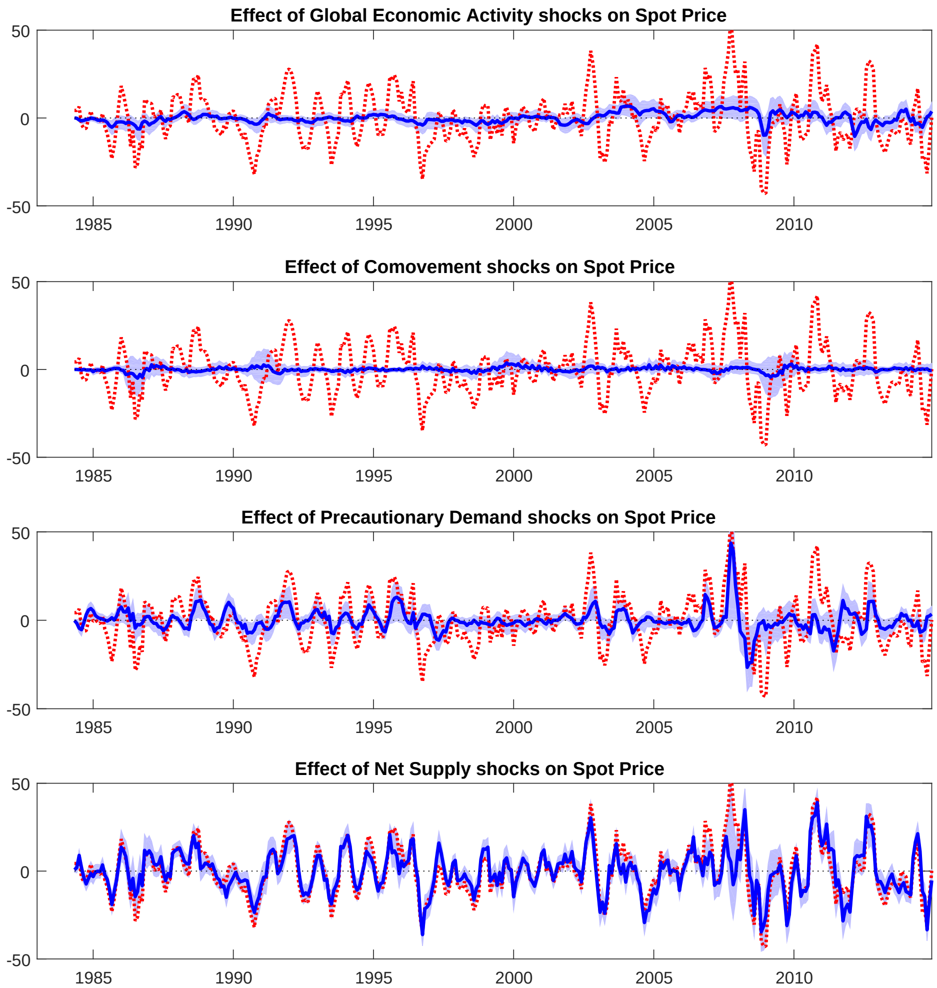

Figure 3.5: Actual changes in wheat prices (red dashed lines) and historical contribution of separate structural shocks with 95% posterior credibility regions (blue and shaded).

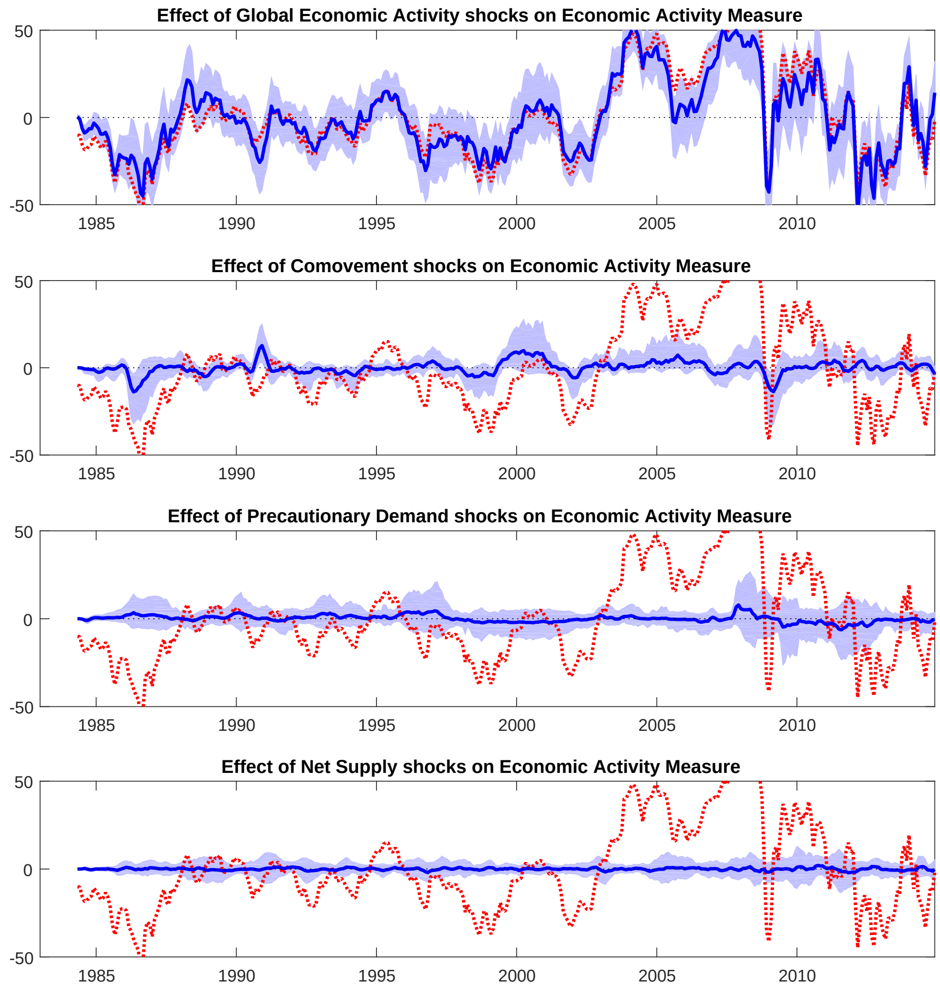

Figure 3.6: Actual changes in global real economic activity (red dashed lines) and historical contribution of separate structural shocks with 95% posterior credibility regions (blue and shaded).

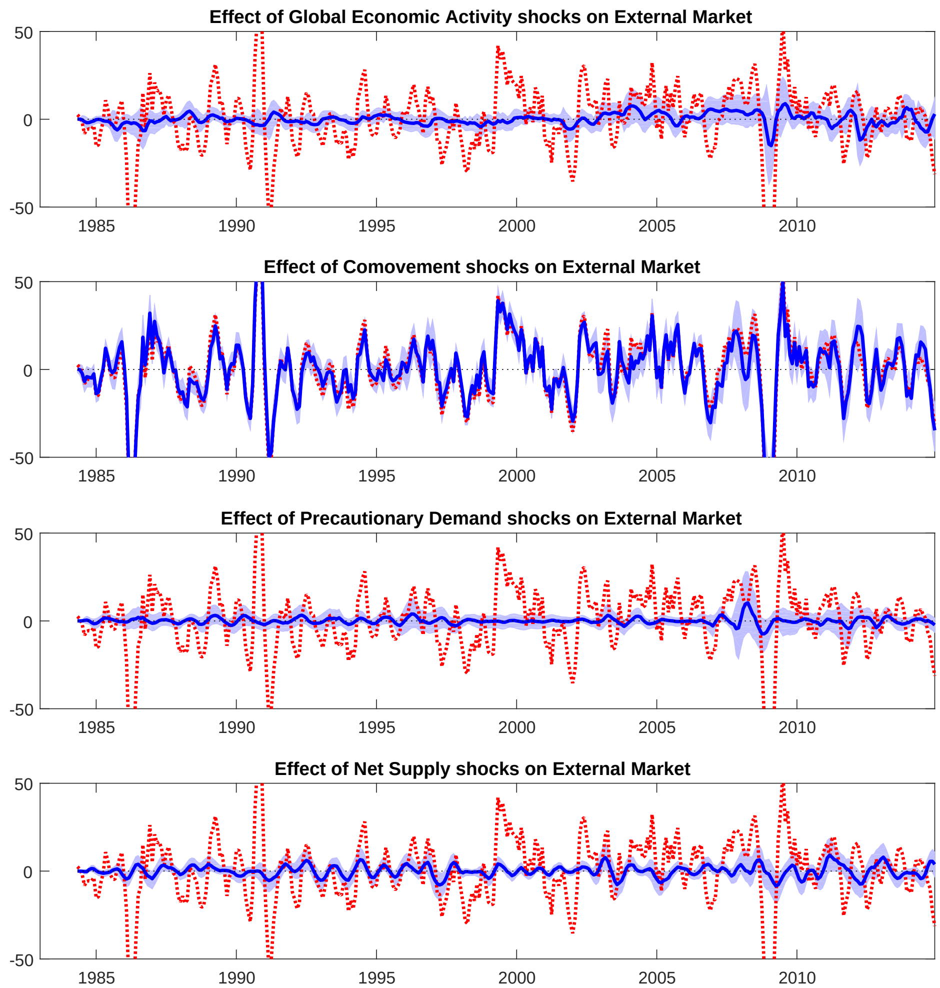

Figure 3.7: Actual changes in the WTI oil price (red dashed lines) and historical contribution of separate structural shocks with 95% posterior credibility regions (blue and shaded).

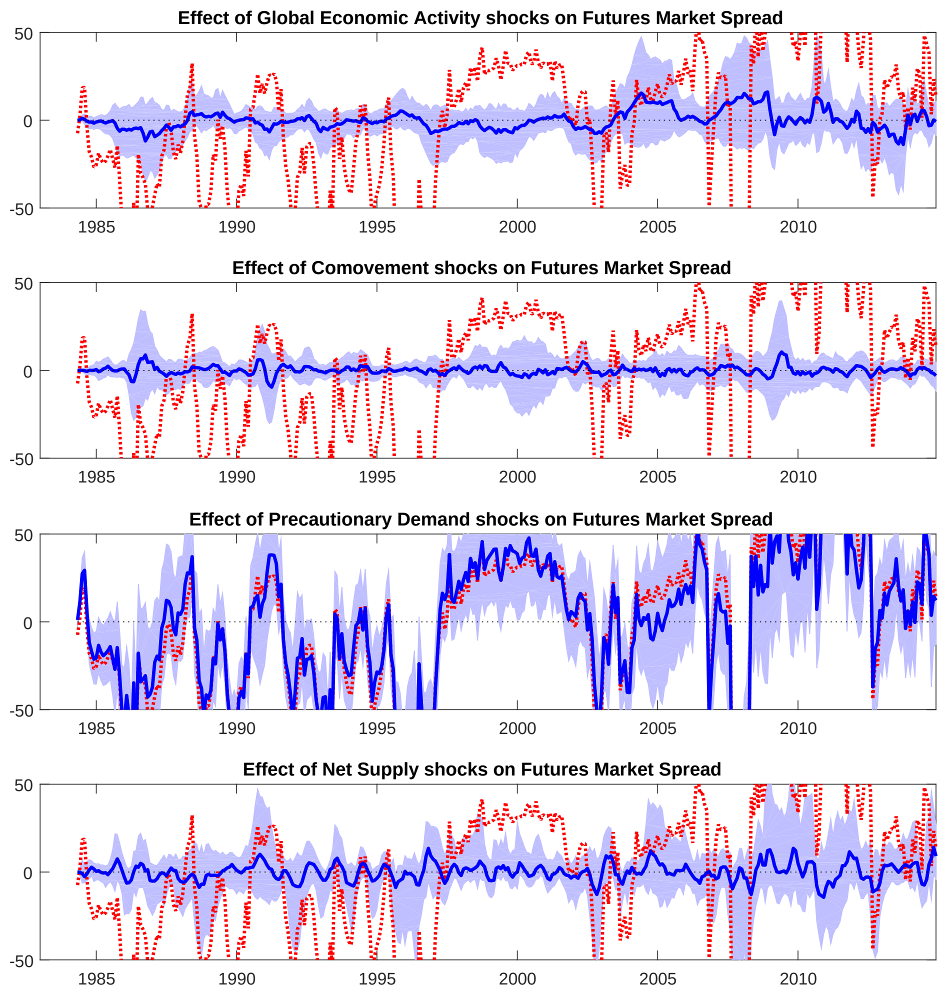

Figure 3.8: Actual changes in wheat futures market spreads (red dashed lines) and historical contribution of separate structural shocks with 95% posterior credibility regions (blue and shaded).

Chapter 4

Forecasting Global Wheat Prices

4.1 Introduction

Given the negative social impact that rapid changes in agricultural commodities have, predicting future movements of such price series is essential. Given the long lead times associated with production decisions and realized prices, accurate forecasts are vitally important to inform both producers and economic policies. There has been a large literature examining the performance of futures prices at predicting spot prices (see, for example, Fama & French (1987), Reichsfeld & Roache (2011), Ahumada & Cornejo (2016*b*)), and we contribute to this literature and extend it by examining the comparative performance of models based on futures prices and those of more complex dynamic models.

Given the importance of food prices to general welfare, forecasting agricultural commodities has been a topic of interest since the earliest formal attempts at forecasting (see, for example, the early forecasts of Working (1927)). Many of the initial analyses of the performance of futures markets were made on agricultural commodity markets, as these were some of the first to have these traded. More recently analysis has focussed on the ability of futures markets to accurately forecast agricultural commodity markets compared to other methodologies, with mixed results Fama & French (1987). This has led to another large literature which examines the predictive content of futures prices in agricultural commodity markets (see for example, Just & Rausser (1981), Tomek (1997)). This analysis investigates the preditive power or forecasts based not only on futures prices themselves but a variety of other forecasts derived from them, such as spread (Alquist et al. 2013).

In this chapter we conduct a pseudo out-of-sample forecasting exercise (Stock & Watson

2003), with the evaluation period of 1990M1:2016M12. The forecasting evaluation evaluates the performance of models which are fitted on an initial training set and evaluated on observations which are held back from the model fitting. Conducting the forecast evaluation in this way ensures that models are only using current and past observations of variables to form forecasts, as would be done by practitioners forming forecasts of the future path of commodity prices. This hold out sample encompasses both the initial period of relative stability throughout the 1990s as well as subsequent volatile price movements. The target series considered include both the nominal and real price of wheat. We present the performance of a variety of parsimonious econometric models, including futures price performance, following Alquist et al. (2013) who examine the oil market. Additionally we present the results for Dynamic Model Averaging (DMA) and Dynamic Model Selection (DMS), two recently developed dynamic models which allow for changes in both model variables and parameter estimates (Koop & Korobilis 2012).

Our forecasting exercise yields two key findings: first, parsimonious econometric models do not perform well at forecasting wheat prices during the sample period under investigation when compared to a benchmark of a random walk. Second, dynamic models which capture the volatile nature of wheat prices perform very well, with DMA in particular providing the best forecast model, exceeding the performance of the benchmark at all horizons under consideration. DMA is also the only model able to consistently predict the direction of the future price change at all horizons, with most other forecast strategies performing no better than random chance.

Additionally, DMA analysis also reveals there is no single variable that leads to performance that is superior to the no-change forecast, and gives insight into the relative importance of the included variables at predicting future wheat prices. Consistently with the results of the previous chapter, there is no evidence that real economic activity (a proxy for increased demand from developing economies) has been driving food prices higher. However, during the rapid increase of wheat prices, grain prices are increasingly likely to be included in predictive models, suggesting that comovement has increased, at least during periods of heightened activity in the wheat market.

The rest of this chapter is structured as follows. Section 4.2 discusses the data used and variables included in the forecasting exercises. An overview of the various testing methodologies and how they are to be applied is presented in Section 4.3. The following section examines the forecasting power of each of the proposed forecasting models, and examines the number of variables dynamic models select to include at each point in time, and the selected drivers of agricultural commodity price movements over the past two decades. A few concluding remarks are offered in Section 4.5.

4.2 Data

4.2.1 Correct Specification of Futures Prices

It is important that the futures price series which are used to compute their predictive power are correctly specified. This is so because misspecification could lead to a bias due to inclusion of information which was not available to market participants at the time they were forming expectations about the future. The empirical analysis presented here is based on monthly prices of wheat futures traded on CME (Chicago Mercantile Exchange) from the data provider Quandl. Near futures prices obtained from the wheat futures (W1) contract continuous series is used as a proxy for wheat spot prices. Therefore we restrict our analysis to forecasts with a horizon of at least $h = 3$ months. Contracts are standardized at 5000 bushels (just over 136 metric tons), of a specified grade. Computing the correct futures price for wheat is more difficult in comparison to other commodities (e.g., crude oil) as wheat contracts are delivered non-sequentially thoughout the year (compared to oil contracts which mature every calendar month). Delivery dates for wheat contracts are March, May, July, September and December (contracts denominated H, K, N, U and Z, respectively). To create a regularly spaced futures price series, we take the prices of all futures contracts traded in a given month and linearly interpolate between the traded horizons to construct any desired h-month ahead contract. This procedure is undertaken as wheat futures with fixed maturities are not traded at all times throughout the year, and due to the irregularly spaced contracts maturities, missing contract prices must be created. The approach taken mitigates the previously mentioned issues and follows Alquist et al. (2013) in creating a consistent end-of-month time series of wheat futures prices of different maturities. This allows uniform matching of computed wheat futures prices of a given maturity with associated spot (or near futures) prices for computing predictive performance.

4.2.2 Variables Included in Dynamic Models

For the DMA analysis, we utilize nine variables that have been proposed in the literature as potential drivers of agricultural commodity prices over the past two decades. Following the findings of Ahumada & Cornejo (2016a) that other commodities are important predictors of agricultural commodities, three of the variables are log-differenced prices of other key grains, namely corn, soybean and oats.

These will capture any changes in the relationship between different agricultural commodity prices, which could indicate whether patterns of comovement have changed over the sample period, due to factors such as the finanicalization of commodity markets. We also include three real effective exchange rate series, motivated by the findings of Chen et al. (2010), who suggest that changes in exchange rates have driven movements in agricultural commodity prices. These are the real effective exchange rate series for USA, Canada and Australia, as compiled by the Bank for International Settlements. Finally, similar to (Ahumada & Cornejo 2016a) two additional macroeconomic indicators, T-Bills and M2 Money supply are included. These U.S. macro variables are included to capture the impact that the prevailing monetary and financial environment, which has been suggested may impact commodity prices. Several commentators have suggested that relatively loose monetary policy may have been an explanatory factor in agricultural price rises and evaluated the impact that US macro variables have had (see for instance, Frankel & Rose (2010) and Ahumada & Cornejo (2015)).

For the construction of real prices, we deflate nominal series by the U.S. CPI. All series are in constant 1984 U.S. dollars. U.S. CPI figures were obtained from the FRED database which are collected by the U.S. Bureau of Labor Statistics.

4.3 Methodology

To explore the forecastability of wheat prices, we follow Alquist et al. (2013) who evaluate the performance of various models for forecasting oil prices and employ a similar structure. We utilize three different target series, nominal wheat prices, real wheat prices and wheat returns (log-differences).

4.3.1 Forecasts Based on Wheat Futures

The first models evaluated investigate the performance of wheat futures prices at forecasting future spot prices. Futures prices are denoted $F_t^{(h)}$, which is the price at a time t for a wheat contract which matures in h periods time. S_t is todays realized spot price, $E_t[S_{t+h}]$ is the expectations of the wheat price at a time $t+h$, given all information available at time t.

To compare forecast performance a benchmark model for comparative purposes is required. The most parsimonious of these is the driftless random walk model, otherwise known as a no-change forecast. This model suggests that no additional information can improve forecast performance (i.e., changes in spot prices are effectively random disturbances). Implicitly this can be interpreted as agricultural commodity markets being efficient (at least weakly), in the sense that they incorporate all available public information, with no additional factor able to improve forecastability. This implies that the best forecast of the of the spot price at any point in the future is the current spot price:

$$\hat{S}_{t+h|t} = S_t. \tag{4.1}$$

The first model which utilizes futures prices suggests that the best available predictor of future spot prices is the futures price for a contract which matures in h periods:

$$\hat{S}_{t+h|t} = F_t^{(h)}. \tag{4.2}$$

A related set of models, widely examined in agricultural commodity markets (see Fama & French (1987)) and other commodities (see Alquist & Kilian (2010)), utilizes the current spread (the difference between spot and futures price) as a predictor of future spot prices. If futures prices correctly encapsulate agents expectations of the spot price in h periods time, then the spread will equal the expected change in the spot price over h periods. The corresponding prediction is given by:

$$\hat{S}_{t+h|t} = S_t(1 + \ln(F_t^{(h)}/S_t)). \tag{4.3}$$

Given that the spread may be a biased predictor, the zero intercept assumption can be relaxed:

$$\hat{S}_{t+h|t} = S_t(1 + \hat{\alpha} + \ln(F_t^{(h)}/S_t)), \tag{4.4}$$

and also the proportionality restriction (i.e., that the slope coefficient is equal to unity):

$$\hat{S}_{t+h|t} = S_t(1 + \hat{\beta}\ln(F_t^{(h)}/S_t)). \tag{4.5}$$

The final forecasting model relaxes both the intercept and slope restrictions:

$$\hat{S}_{t+h|t} = S_t(1 + \ln(F_t^{(h)}/S_t)). \tag{4.6}$$

4.3.2 Parsimonious Econometric Forecasts

There are many parsimonious econometric forecasting techniques which have been employed widely in the literature. We consider each of these models and assess their ability to forecast nominal wheat prices. The first specification under consideration is the double-differenced model of Hendry (2006):

$$\hat{S}_{t+h|t} = S_t(1 + \Delta s_t)^h, \tag{4.7}$$

where Δs_t is the percentage change in spot price since the previous period. This model builds on the insight that when forecasting time series with rare changes in trend, an improved forecast can be made by extrapolating from the most recent growth rate. Put another way, this can be seen as a random walk forecast applied to the growth rate of wheat prices rather than the level. Given the relative frequency of rapid changes in wheat prices in the sample period, this method has the potential to perform well.

A further method similar to that proposed by Hendry (2006) in Equation (4.7) utilizes locally estimated drifts from rolling regressions:

$$\hat{S}_{t+h|t} = S_t(1 + \Delta s_t^{-(h)}), \tag{4.8}$$

where $\Delta s_t^{-(h)}$ is the recent difference over the previous h months for the spot price. The intuition is that agents are backwards looking when they form expectations about future spot prices. In this case the recent growth over the forecast horizon is extrapolated forward to inform the spot price. This model therefore encapsulates the ability of "short-term" forecasts to model localized trends in wheat prices for predictive purposes.

The final parsimonious models we test are based on the findings of Chen et al. (2010), and use recent percentage changes in bilateral exchange rates of selected exporters of wheat:

$$\hat{S}_{t+h|t} = S_t(1 + \Delta e_t^i)^h, \tag{4.9}$$

where Δe_t^i is the percentage change in the bilateral exchange rate for country i. In our empirical study we examine $i \in \{Australia, Canada\}$, which were the third and fourth largest exporters of wheat in 2017. We do not include the second largest (Russia) due to difficulties obtaining comparable exchange rate data.

4.3.3 Unrestricted AR, ARMA and VAR models

To examine the forecastability of real prices of wheat out-of-sample, we utilize standard $AR(p)$ and $ARMA(p, q)$ iterative methods. For comparative purposes we evaluate all $ARMA$ based forecasting models twice, once on log levels, and once on log differences where we in essence impose a unit root. We note that after imposing the unit root the $AR(p)$ lag order is reduced by 1, i.e., a $AR(12)$ model in log-levels corresponds to a $AR(11)$ model in differences. The $ARMA(p, q)$ model is the most general of these and nests the $AR(p)$:

$$y_t = \alpha + \sum_{i=1}^{p} \phi_i y_{t-i} + \sum_{j=1}^{q} \theta_i \epsilon_{t-j} + \epsilon_t, \tag{4.10}$$

where p is the autoregressive lag order and q is the moving average order.

The previously described models are effectively limited to past values of the real price of wheat. For comparison purposes an unrestricted Vector Autoregressive model is constructed containing the same variables as in the previous Chapter. This VAR model includes the log-differenced wheat price, oil price, real economic activity indicator and the futures spread.

All of the iterative models are estimated from 1980M1 and the evaluation period is 1990M1 - 2016M12 to allow direct comparison with the other methods presented. Although these methods produce iterative forecasts, for comparison purposes and forecast error evaluation we utilize $\hat{y}_{t+h}$ with $h \in \{3, 6, 9, 12\}$.

4.3.4 Dynamic Model Averaging and Dynamic Model Selection

We now turn attention to Dynamic Model Averaging (DMA) and Dynamic Model Selection (DMS) and the application of them to the log difference of wheat prices. DMA was initially introduced by Raftery et al. (2010) and introduced to the economic literature by Koop & Korobilis (2012). These methods have since gained broad acceptance in the applied macroeconomic literature and subsequently have been applied widely to a variety of areas of interest. These include forecasting oil (Drachal 2016, Naser 2016), copper (Buncic & Moretto 2015) and gold (Aye et al. 2015) prices.

To examine how DMA operates we begin by examining a single model equation, before building to the multiple forecasting model framework. If y_t is the log-differenced wheat price, with several exogenous predictor variables we are interested in a class of models:

$$y_t = \theta_{1t} + \theta_{2t} y_{t-1} + \theta_{3t} y_{t-2} + \theta_{4t} x_{t-1} + \cdots + \theta_{ht} z_{t-1} + \epsilon_t, \quad \epsilon_t \sim N(0, V_t). \tag{4.11}$$

For tractability we will simplify the above model and write it in state-space form for estimation using the Kalman Filter. With no uncertainty regarding the correct forecasting model for $t = 1, \ldots, T$ the standard specification can be written:

$$y_t = z_t\theta_t + \epsilon_t, \tag{4.12}$$

$$\theta_t = \theta_{t-1} + \eta_t. \tag{4.13}$$

For our empirical example $z_t = [1, z_{t-h}]$, the $1 \times m$ vector of predictive variables for log-differenced wheat price changes and θ_t is a $m \times 1$ vector of coefficients (the states) which can vary over time. The errors, $\epsilon_t \overset{iid}{\sim} N(0, H_t)$ and $\eta_t \overset{iid}{\sim} N(0, Q_t)$, are assumed to be mutually independent of all leads and lags.

The model is estimated recursively using the Kalman Filter. If $\mathcal{F}_{t-1}$ is all available information available at a time $t-1$ then:

$$\theta_{t-1} \mid \mathcal{F}_{t-1} \sim N(\hat{\theta}_{t-1|t-1}, \Sigma_{t-1|t-1}), \tag{4.14}$$

where $\Sigma_{t-1|t-1}$ is the state covariance matrix which updates according to the standard formula:

$$\Sigma_{t|t-1} = \Sigma_{t-1|t-1} + Q_t. \tag{4.15}$$

Prediction is done according to the following expression:

$$\theta_{t-1} \mid \mathcal{F}_{t-1} \sim N(\hat{\theta}_{t|t-1}, \Sigma_{t-1|t-1}). \tag{4.16}$$

However, to complete this procedure, Q_t in Equation (4.15) needs to be specified. Raftery et al. (2010) propose replacing Equation (4.15) by using a forgetting factor. Such forgetting factors have been utilized often throughout the state-space literature. We note here that this simplification means that we no long need to estimate Q_t, resulting in the substantially simpler:

$$\Sigma_{t|t-1} = \frac{1}{\lambda}\Sigma_{t-1|t-1}. \tag{4.17}$$

This implies that the state error covariance matrix Q_t is directly related to the forgetting factor, λ:

$$Q_t = \frac{1-\lambda}{\lambda}\Sigma_{t-1|t-1}. \tag{4.18}$$

Therefore when $\lambda = 1$, $Q_t = 0$, $\Sigma_{t|t-1} = \Sigma_{t-1|t-1}$, and there is no variation in the parameter estimates over time[1]. Specifying values of $\lambda < 1$ introduces time variation in the parameter estimates by controlling the effective sample size under consideration. For example, in a monthly context if $\lambda = 0.99$ then observations 1 year ago receive approximately 88% as much weight as the previous periods observation, whereas $\lambda = 0.95$ places a weight of 54% on an observation 12 months previously. Lower values of λ correspond to significantly increased parameter instability. For this reason it is recommended that values of λ are restricted to near unity (Koop & Korobilis 2012). We consider multiple values of λ in our empirical application as discussed below.

The tth iteration is updated according to the updating equation:

$$\Sigma_t \mid y_t \sim N(\hat{\theta}_{t|t}, \Sigma_{t|t}) \tag{4.19}$$

where:

$$\hat{\theta}_{t|t} = \hat{\theta}_{t|t-1} + \Sigma_{t|t-1}z_t^\top(H_t + z_t\Sigma_{t|t-1}z_t^\top)^{-1}(y_t - z_t\hat{\theta}_{t|t-1}), \tag{4.20}$$

and

$$\Sigma_{t|t} = \Sigma_{t|t-1} - \Sigma_{t|t-1}z_t^\top(H_t + z_t\Sigma_{t|t-1}z_t^\top)^{-1}z_t\Sigma_{t|t-1}. \tag{4.21}$$

This leaves H_t, the error covariance matrix, unspecified. Various proposals for dealing with H_t exist, with Koop & Korobilis (2012) proposing an Exponentially Weighted Moving Average specification. We follow the original suggestion of Raftery et al. (2010), who suggest replacing H_t with a consistent estimate H_t^*:

$$H_t^* = \frac{1}{t}\sum_{r=1}^{T}\left[(y_t - z_t\hat{\theta}_{t|t-1})^2 - z_t^\top\Sigma_{t|t-1}z_t\right]. \tag{4.22}$$

[1] As discussed later, this result is referred to as Bayesian Model Averaging (BMA), see Koop & Korobilis (2012).

This is consistent as $H_t^* \to H_t$ as $t \to \infty$. However as it is not guaranteed that $H_t^* > 0$, Raftery et al. (2010) define a recursive moment estimator:

$$\hat{H}_t = \begin{cases} H_t^*, & \text{if } A_t > 0, \\ \hat{H}_{t-1}, & \text{otherwise,} \end{cases} \tag{4.23}$$

where:

$$A_t = \left(\frac{t-1}{t}\right)\hat{H}_{t-1} + \frac{1}{t}((y_t - z_t\hat{\theta}_{t|t-1})^2 - z_t^\top \Sigma_{t|t-1} z_t)). \tag{4.24}$$

We can now generate one-step-ahead forecasts recursively using the predictive distribution:

$$y_t \mid \mathcal{F}_{t-1} \sim N\left(z_t\hat{\theta}_{t-1},\, H_t + z_t\Sigma_{t|t-1}z_t^\top\right). \tag{4.25}$$

So far we have limited our analysis to a single forecasting model. Parameter estimates have been allowed to vary over time, however the model itself has been fixed. DMA and DMS relax this restriction and allow model specification to vary over time, with $k = 1, \ldots, K$ potential model specifications. With n predictive variables there are $K = 2^n$ potential models, each containing a different subset of the available predictive variables. For our empirical application $K = 512$. If K is restricted to 1 as in Equations (4.12) - (4.25), DMA collapses to the well-known Time-Varying parameter methodology.

We generalize Equations (4.12) and (4.13) to the multiple model case as follows:

$$y_t = z_t^{(k)}\theta_t^{(k)} + \epsilon_t^{(k)}, \tag{4.26}$$

and

$$\theta_t^{(k)} = \theta_{t-1}^{(k)} + \eta_t^{(k)}. \tag{4.27}$$

$\theta_t^{(k)}$ is the vector of parameters for the variables within a particular model k, $\epsilon_t^{(k)} \sim N(0, H_t^{(k)})$ and $\eta_t^{(k)} \sim (0, Q_t^{(k)})$. In the multiple model case, the Kalman Filter Equations (4.14), (4.16) and (4.19) generalize to:

$$\Theta_{t-1} \mid L_{t-1} = k,\, \mathcal{F}_{t-1} \sim N(\hat{\theta}^{(k)}_{t-1|t-1}, \Sigma^{(k)}_{t-1|t-1}), \tag{4.28}$$

$$\Theta_{t} \mid L_{t} = k,\, \mathcal{F}_{t-1} \sim N(\hat{\theta}^{(k)}_{t|t-1}, \Sigma^{(k)}_{t-1|t-1}), \tag{4.29}$$

and

$$\Theta_{t} \mid L_{t} = k,\, \mathcal{F}_{t-1} \sim N(\hat{\theta}^{(k)}_{t|t}, \Sigma^{(k)}_{t|t}). \tag{4.30}$$

Where $L_t = k$, is the particular model specification which holds at time t and Θ_t is the vector of all coefficients (i.e, $\Theta_t = (\theta_t^{(1)\top}, \ldots, \theta_t^{(K)\top})^\top$). $\hat{\theta}^{(k)}_{t|t}$, $\Sigma^{(k)}_{t|t}$ and $\Sigma^{(k)}_{t|t-1}$ are effectively the same as in the single model case, with (k) indicating the particular model k, and are obtained using Equations (4.17), (4.20) and (4.21). The above specification suggests that, conditional on $L_t = k$, information from the prediction and updating equations will not provide information on the entirety of Θ_t, they only provide information on the model specific $\theta_t^{(k)}$. To achieve unconditional results (i.e., not conditional on a particular model, $L_t = k$), a transition matrix which governs the evolution of L_t, $P = (p_{kl})$, where $p_{kl} = P\left[L_t = l \mid L_{t-1} = k\right]$, is required. Computationally calculating P with a large number of models is onerous and empirically undesirable due to the excessive number of parameters (Koop & Korobilis 2012). Analogously with Q_t, Raftery et al. (2010), propose an elegant solution to this issue using a second forgetting factor, α, for the state equation for the models.

To see how this forgetting factor greatly aids computation we need the multi-model analogue of Equation (4.14), using the simplification that $\pi_{t|s,\,l} = p(L_t = l \mid \mathcal{F}_s)$:

$$p(\Theta_{t-1} \mid \mathcal{F}_{t-1}) = \sum_{k=1}^{K} p\left(\theta^{(k)}_{t-1} \mid L_{t-1} = k,\, \mathcal{F}_{t-1}\right) \times \pi_{t-1|t-1,\,k}, \tag{4.31}$$

where $p\left(\theta^{(k)}_{t-1} \mid L_{t-1} = k,\, \mathcal{F}_{t-1}\right)$ is given by Equation (4.28). The probability that a particular model k should be used to predict wheat price changes is:

$$\pi_{t|t-1,\,k} = \sum_{l=1}^{K} \pi_{t-1|t-1,\,l}\, p_{kl}, \tag{4.32}$$

where p_{kl} is an element from the unrestricted transition matrix P. To avoid specification of this transition matrix Raftery et al. (2010) replace Equation (4.32) with:

$$\pi_{t|t-1,\,k} = \frac{\pi^{\alpha}_{t-1|t-1,\,k}}{\sum_{l=1}^{K} \pi_{t-1|t-1,\,l}}, \tag{4.33}$$

where α is a parameter slightly below 1. To allow interpretation of this forgetting parameter we can investigate the weight which DMA places upon a particular model at a point in time:

$$\begin{aligned} \pi_{t|t-1,\,k} &\propto \left[\pi_{t-1|t-2,\,k} p_k(y_{t-1} \mid \mathcal{F}_{t-2})\right]^{\alpha} && (4.34) \\ &= \prod_{i=1}^{t-1} \left[p_k(y_{t-i} \mid \mathcal{F}_{t-i-1})\right]^{\alpha^i}, && (4.35) \end{aligned}$$

where the predictive density for a model l is given by $p_l(y_t \mid \mathcal{F}_{t-1})$ which is given by updating the distribution from Equation (4.25), $N\left(z_t^{(l)}\hat{\theta}_{t-1}^{(l)},\, H_t^{(l)} + z_t^{(l)}\Sigma_{t|t-1}^{(l)} z_t^{(l)\top}\right)$. The weight that a model k will receive at time t therefore depends on whether it has performed well at forecasting in the recent past. The definition of 'recent past' exponentially decays at the rate α^i for observations i periods into the past, the same rate of decay as the previous forgetting factor, λ. For a value of $\alpha = 0.99$, performance one year ago receives around 88% of the weight as forecast performance last period (when using monthly data), whilst if $\alpha = 0.95$ forecast performance one year ago receives approximately 54% as much weight. As with the range of the parameter λ, Koop & Korobilis (2012) suggest limiting the range to $\alpha \in [0.95, 1]$.

Finally, the model probabilities can be updated each period according to:

$$\pi_{t|t,\,k} = \frac{\pi_{t|t-1,\,k} p_k(y_t \mid \mathcal{F}_{t-1})}{\sum_{k=1}^{K} \pi_{t|t-1,\,l} p_l(y_t \mid \mathcal{F}_{t-1})}. \tag{4.36}$$

We implement a two stage procedure for the Kalman Filter. First, we calculate the probability that any particular model, $L_t = k$, should be used for forecasting at a time period t, from Equation (4.33). Then conditional on $L_t = k$, the predicted vector of regression coefficients, $\hat{\theta}_{t|t-1}^{(k)}$ is computed using Equation (4.29). Following this updating Equations (4.30) and (4.36) are used to update the parameter estimates and the model probabilities, respectively.

Recursive forecasts can then be calculated by taking the weighted average of the predicted results for wheat price changes from all models, using $\pi_{t|t-1,\,k}$ as weights:

$$\hat{y}_t^{DMA} = E\left[y_t \mid \mathcal{F}_{t-1}\right] = \sum_{k=1}^{K} \pi_{t|t-1,\,k} z_t^{(k)} \hat{\theta}_{t-1}^{(k)}. \tag{4.37}$$

DMA therefore presents a model which allows us to deal with uncertainty surrounding the underlying population regression function. It essentially deals with both model uncertainty and parameter uncertainty in a cohesive way by using model combinations to produce forecasts. As previously mentioned DMS selects the model with the highest inclusion probability, k^*. Therefore DMS results are given by:

$$\hat{y}_t^{DMS} = z_t^{(k^*)} \hat{\theta}_{t|t-1}^{(k^*)}, \tag{4.38}$$

i.e., the optimal variable set multiplied by the computed parameter estimates. Therefore, DMS presents the results of the model with the highest probability at each point in time, as opposed to the weighted average across all models presented by DMA. Empirically, we evaluate DMA and DMS forecasting with a variety of forgetting parameter values; (i) $\lambda = \alpha = 0.95$, which allows rapid changes in both the model and parameter estimates, (ii) $\lambda = \alpha = 0.99$, which implies less volatile parameter estimates and less rapid model switching and (iii) $\lambda = 1, \alpha = 0.95$, which allows relatively rapid model selection, but fixed parameter estimates. The final case under consideration is BMA and BMS, which are effectively special cases of DMA and DMS, with both forgetting factors set to 1. This implies that there is no updating of the parameter estimates or model selection.

To initialize the model we follow Koop & Korobilis (2012) and assume that in the initial period all models are equally likely ($\pi_{0|0,\,k} = 1/K$ for $k \in 1, \ldots, K$). All estimates are obtained using the R package $eDMA$ (Catania & Nonejad 2017).

4.3.5 Forecast Evaluation Methodology

For comparative purposes, ratios of Mean Squared Predictive Errors (MSPEs) between the model under evaluation and the benchmark no-change forecast are reported. Ratios below one indicate that the model under evaluation has exceeded the performance of the benchmark in predicting spot prices. To test the significance of any out-performance we construct p-values based on the Clark & West (2007) test. This procedure examines the difference between $\hat{\sigma}_1^2$ and $\hat{\sigma}_2^2$, the sample mean squared predictive errors. Specifically, the proposed test uses adjusted MSPEs (i.e., examining $\hat{\sigma}_1^2 - (\hat{\sigma}_2^2 - adj)$, where adj indicates the adjusted MSPE), due to the fact that an upward bias exists in the MSPE produced by estimating parameters which are equal to zero under the null hypothesis (Clark & West 2007).

If y_{t+h} is the realized change in wheat prices h periods into the future and $\hat{y}_{RW,\,t,\,t+h}$ and $\hat{y}_{i,\,t,\,t+h}$ are the predicted values for change in wheat prices at a horizon h for the benchmark random walk forecast and the forecast model under evaluation (i), then we can compute:

$$\hat{f}_{t+h} = (y_{t+h} - \hat{y}_{RW,\,t,\,t+h})^2 - \left[(y_{t+h} - \hat{y}_{i,\,t,\,t+h})^2 - (\hat{y}_{RW,\,t,\,t+h} - \hat{y}_{i,\,t,\,t+h})^2)\right]. \tag{4.39}$$

$\hat{\sigma}_1^2 - (\hat{\sigma}_2^2 - adj)$ is the average of $\hat{f}_{t+h}$ so a test statistic can be generated by regressing $\hat{f}_{t+h}$ on a constant. Testing that this coefficient is equal to zero leads to a test which is normally distributed (Clark & West 2007). Standard tests for parameter estimates can be used to perform hypothesis testing and we reject the null of equal performance if the computed statistic exceeds 1.645 at the 5% level. p-values are constructed from the test statistics and shown below relative MSPEs in all forecasting exercises. The null hypothesis of the resulting test is that of equal MSPEs, compared to an alternative that a particular model is able to outperform the benchmark.

As an additional test we also examine the performance of the models at correctly predicting the *direction* of the change in wheat prices. This is for two related reasons: first the assessment of the performance of the forecasting model is based upon the loss function of the user of the forecasts (see Elliott & Timmermann (2008)) and as such squared forecast errors may not always be appropriate. Several alternative measures of forecast performance have been tested, showing standard performance indicators may be sub-optimal for calculating economic evaluations in commodity markets (Gerlow et al. 1993). Second, economic profits from forecasts have been shown to have a significantly stronger relationship with directional accuracy than other measures of forecast performance (Leitch & Tanner 1991). Pesaran & Timmermann (1992), provide a simple non-parametric test of directional accuracy which we perform for all forecast models under consideration.

The realized series of interest is given by y_t, with a forecast given by $\hat{y}_t$, $t = 1, \ldots, T$ and we define $\pi_1 = p(y_t > 0)$ and $\pi_2 = p(\hat{y}_t > 0)$. p_1 and p_2 are the sample proportions for the times that the realized value of y_t is positive, and the forecast $\hat{y}_t$ is positive, respectively. Under a null hypothesis that the forecast and realized values are independently distributed (i.e., there is no predictive power from the forecast values), the number of correct sign predictions follows a binomial distribution, with T trials. The success probability is given by:

$$\pi^* = \pi_1\pi_2 + (1 - \pi_1)(1 - \pi_2). \tag{4.40}$$

The sample proportions can be used to estimate π_1 and π_2:

$$p^* = p_1p_2 + (1 - p_1)(1 - p_2), \tag{4.41}$$

where p^* represents the expected proportion of direction forecasts accurately predicted by the forecast in the sample under the null hypothesis. Finally, let p be the actual number of times the sign of y_t is predicted. When p is presented as a percentage it is the Hit Ratio. The test statistic is given by:

$$S_n = (p - p^*) \left[v\hat{a}r(p) - v\hat{a}r(p^*)\right]^{0.5}, \tag{4.42}$$

where $v\hat{a}r(p)$ and $v\hat{a}r(p^*)$ are given by:

$$v\hat{a}r(p) = T^{-1}p^*(1 - p^*), \tag{4.43}$$

and

$$v\hat{a}r(p^*) = T^{-1}(2p_1-1)^2p_2(1-p_2)+T^{-1}(2p_2-1)^2p_1(1-p_1)+T^{-2}4p_1p_2(1-p_1)(1-p_2). \tag{4.44}$$

Pesaran & Timmermann (1992) show that the test statistic in this scenario converges to a normal distribution under the null that the forecast series is unable to predict the series of interest. From this appropriate p-values can be constructed.

4.4 Results

4.4.1 Comparison of Forecasts based on Futures Prices

We start our empirical analysis with the evaluation of nominal wheat price forecasts generated by the variety of econometric models that utilize monthly future prices. The forecast evaluation period under consideration is 1991M1-2016M12. Table 4.1 presents relative ratios of the MSPE from the proposed model to the MSPE to the random walk without drift benchmark model. Additionally, the hit ratio, the relative frequency with which any model correctly identifies the direction of the change in the price, is reported. p-values, based upon Clark & West (2007) test statistics, are reported in parentheses below the respective relative MSPE. In addition p-values are reported for the hit ratio test using test statistics from the Pesaran & Timmermann (1992) test. In cases where the null that a given model is equally accurate as the benchmark no-change forecast is rejected at the 5% level they are highlighted in bold. Similarly, cases when the null that the model is unable to forecast the target series is rejected at the 5% level when directional forecasting is analyzed are also highlighted in bold.

The results of the second row of Table 4.1 demonstrate the forecast ability of the model with futures as a single predictor and demonstrates that futures prices are unable to significantly outperform the benchmark except at a forecasting horizon of three months. Futures forecasts do not provide additional information about the expected change in direction of wheat prices either, while horizons of 3, 6 and 9 months all have relative frequency's worse than would be expected by random chance.

Turning to the result for spread based forecasts, Rows (3) to (6), we observe an improvement with respect to futures prices. Specifically, the simplest model (Row (6)) shows predictive gains of between 2 and 10% over the no-change forecast. However, we do not find significant gains in forecasting the direction of movements in nominal wheat prices.

Overall, the above results suggest relatively modest forecasting gains from using forecasts based on futures prices, with the best performing model utilizing spreads showing improvements of 10% at the longest horizon under investigation. Forecasters particularly

Table 4.1: Forecast Error Diagnostics for the Nominal Price of Wheat from Parsimonious Econometric Models

$\hat{S}_{t+h\|t}$	$h=3$		$h=6$		$h=9$		$h=12$	
	MSPE	HR	MSPE	HR	MSPE	HR	MSPE	HR
S_t	5819.32	N.A.	10822.3	N.A.	16820.66	N.A.	21532.9	N.A.
$F_t^{(h)}$	**0.98**	0.47	0.99	0.46	1.00	0.49	1.00	0.55
	(0.04)	(0.84)	(0.14)	(0.61)	(0.41)	(0.75)	(0.44)	(0.15)
$S_t(1+\hat{\alpha}+\hat{\beta}\ln(F_t^{(h)}/S_t))$	0.99	0.48	**0.96**	0.52	**0.97**	0.48	**0.92**	0.53
	(0.06)	(0.74)	(0.01)	(0.26)	(0.01)	(0.71)	(0.02)	(0.14)
$S_t(1+\hat{\alpha}+\ln(F_t^{(h)}/S_t))$	1.49	0.50	2.61	**0.57**	2.04	0.52	2.19	0.48
	(0.07)	(0.39)	(0.05)	(0.01)	(0.01)	(0.23)	(0.01)	(0.70)
$S_t(1+\hat{\beta}\ln(F_t^{(h)}/S_t))$	1.00	0.47	1.01	0.49	1.03	0.49	**0.96**	0.52
	(0.21)	(0.82)	(0.07)	(0.61)	(0.04)	(0.62)	(0.02)	(0.18)
$S_t(1+\ln(F_t^{(h)}/S_t))$	**0.98**	0.47	**0.96**	0.49	**0.97**	0.49	**0.90**	0.53
	(0.04)	(0.84)	(0.01)	(0.66)	(0.00)	(0.66)	(0.01)	(0.15)
$S_t(1+\Delta s_t)^h$	6.69	0.51	20.98	0.44	67.29	0.51	254.99	0.53
	(0.90)	(0.28)	(0.74)	(0.98)	(0.91)	(0.27)	(0.94)	(0.12)
$S_t(1+\Delta s_t^{-h})$	2.39	0.51	2.69	0.50	3.67	0.49	4.00	0.50
	(0.79)	(0.37)	(0.69)	(0.53)	(0.98)	(0.64)	(0.99)	(0.48)
$S_t(1+\Delta e_t^{AUS})^h$	1.58	0.48	2.06	0.47	2.79	0.50	3.84	0.47
	(0.80)	(0.76)	(0.50)	(0.81)	(0.70)	(0.42)	(0.78)	(0.85)
$S_t(1+\Delta e_t^{CAN})^h$	1.22	0.51	1.77	0.48	2.16	0.47	3.01	0.54
	(0.35)	(0.33)	(0.93)	(0.77)	(0.79)	(0.82)	(0.77)	(0.05)

Note: MSPE results are presented as ratios relative to the MSPE of the no-change forecast model, for which the level of the MSPE is reported. The evaluation period under consideration is 1990M1:2016M12. The training set for initial estimation is 1980M1:1989M12. p-values are presented in parentheses with a null of equal predictive accuracy as the benchmark forecast from Clark & West (2007) tests. The Hit Ratio, HR, is defined as the percentage of forecasts which correctly identify the direction of movement in the price of wheat. p-values from Pesaran & Timmermann (1992) are provided beneath.

interested in the future direction of nominal wheat prices would be best served by utilizing the no change model, as in many instances these models performed worse than random chance. The above results are consistent with findings for wheat from previous studies such as Fama & French (1987) and agricultural commodity markets more generally where performance of futures prices has been mixed (Just & Rausser 1981). Next we examine the performance of other parsimonious econometric models under investigation.

4.4.2 Comparison of Forecast Performance for other Parsimonious Econometric models

The results of the double differenced model of Hendry (2006) are presented in Row (7) of Table 4.1. As is evident, this model performs particularly poorly, especially at long horizons where forecast errors are many times larger in magnitude than the benchmark model. This result is unsurprising given that large monthly movements in wheat prices were witnessed during the evaluation period and this model effectively extrapolates these changes over many periods into the future. The double differenced specification does not offer any improvements in directional forecasting either.

Row (8) presents results for the local drift model based upon extrapolating recent performance in spot prices using rolling regressions. The forecasts from this model are not as erratic as those of Hendry (2006). However, MSPEs are still large, in some cases twice the magnitude of the driftless random walk forecast. Once again, no significant performance in identifying the sign of the change in wheat prices is found.

Finally, we evaluate the performance of forecasts using changes in exchange rates of major exporting nations building on from the results of Chen et al. (2010). Rows (9) and (10) present results for Australia and Canada, respectively. At no horizon does inclusion of the bilateral exchange rate improve forecast performance. Similarly, no additional performance benefit at forecasting future changes of direction in wheat prices is observed.

To summarize, no-change forecasts are as accurate as those based upon simple econometric models or those which include bilateral exchange rates. These results are consistent with other commodities which are equally difficult to find simple models which significantly outperform no-change forecasts. We note two general observations from the results of forecasting nominal wheat prices. First, MSPEs of random walk forecasts are very large in comparison to previous studies of forecasting commodity prices, for instance oil (Alquist & Kilian 2010). This finding results from large fluctuations wheat prices have undergone during the evaluation period under consideration. Second, in general parsimonious econometric models perform poorly at outperforming no change forecasts of nominal monthly wheat prices.

4.4.3 Forecasting Performance of unrestricted AR, ARMA and VAR models for the Real Price of Wheat

We now turn to forecasting performance of real wheat prices, investigating the performance of AR and ARMA models for prices in both log levels and log differences. In both cases the evaluation is performed recursively with an evaluation period of 1990M1:2016M12, with forecast performance evaluated against no-change benchmark. The information set available for forecasting is effectively limited to past values of the deflated price of wheat.

We present a variety of model specifications, some using fixed lag specifications corresponding to one or two years of past observations, 12 and 24 lags, respectively. Other models utilize information criterion measures for appropriate lag selection. Both Schwarz-Bayesian information criterion (BIC) and Akaike information criterion (AIC) are also implemented for appropriate lag order selection (Marcellino et al. 2006). In both cases $AR(p)$ coefficients are evaluated for $p \in \{0, \ldots, 12\}$ using step-wise selection.

The top half of Table 4.2 presents results on the log levels, the bottom half the log differenced results. We take the forecast for h periods ahead from each model and and evaluate it as if it was a direct forecast, utilizing the Clark & West (2007) test statistics to generate p-values. (It should be noted that there are no theoretical results in the forecasting literature of how to evaluate $ARMA$ based models to a no-change forecast so these are computed for direct comparability with the previous results and results yet to be presented.)

The upper panel of Table 4.2 suggests that forecasts based on $AR(p)$ models with the relevant lag order being determined recursively improve forecast performance compared to the no-change forecast for real wheat prices. Contrary to what might be expected, forecast performance improves with longer horizon forecasts. Performance on differenced series is practically indistinguishable from no-change benchmark forecasts, with no improvements found. Interestingly, in the case where lag selection is determined by the Schwartz-Bayesian Information Criterion the model selects the random walk forecast (i.e., $ARMA(0,1,0)$ is chosen).

The unrestricted VAR model results are presented in the final two rows of Table 4.2. The variables included are the log-differenced real price of wheat, the log-differenced price of oil,

Table 4.2: Relative MSPE Ratios and Forecast Error Diagnostics for Real Wheat Prices from Autoregressive Models

		$h=3$		$h=6$		$h=9$		$h=12$	
		MSPE	HR	MSPE	HR	MSPE	HR	MSPE	HR
s_t	$AR(12)$	1.04	0.51	1.05	0.48	1.04	0.49	1.02	**0.55**
		(0.07)	(0.29)	(0.54)	(0.74)	(0.37)	(0.68)	(0.16)	(0.04)
	$AR(BIC)$	**0.99**	0.48	**0.98**	0.48	**0.96**	0.50	**0.95**	0.54
		(0.01)	(0.77)	(0.00)	(0.79)	(0.00)	(0.48)	(0.00)	(0.09)
	$AR(AIC)$	**0.99**	0.47	**0.98**	0.48	**0.97**	0.50	**0.96**	0.54
		(0.03)	(0.83)	(0.00)	(0.79)	(0.00)	(0.48)	(0.00)	(0.09)
	$ARMA(1,1)$	0.99	0.47	0.98	0.48	0.97	0.50	0.96	0.54
		(0.02)	(0.84)	(0.01)	(0.79)	(0.00)	(0.49)	(0.00)	(0.09)
Δs_t	$AR(11)$	1.05	0.51	1.07	0.48	1.06	0.49	1.05	**0.55**
		(0.15)	(0.28)	(0.75)	(0.74)	(0.65)	(0.59)	(0.42)	(0.04)
	$AR(BIC)$	-	-	-	-	-	-	-	-
		(-)	(-)	(-)	(-)	(-)	(-)	(-)	(-)
	$AR(AIC)$	1.01	0.48	1.02	0.47	1.03	0.50	1.03	**0.55**
		(0.95)	(0.77)	(1.00)	(0.87)	(0.99)	(0.43)	(0.98)	(0.05)
	$ARMA(0,1)$	1.00	0.48	1.00	0.48	1.00	0.50	1.00	0.54
		(0.16)	(0.77)	(0.22)	(0.79)	(0.15)	(0.49)	(0.35)	(0.09)
	$VAR(12)$	1.17	0.54	1.14	0.52	1.06	0.52	1.03	0.53
		(0.06)	(0.11)	(0.20)	(0.21)	(0.08)	(0.23)	(0.04)	(0.17)

Note: MSPE results are presented as ratios relative to the MSPE of the no-change forecast model. The evaluation period under consideration is 1990M1:2016M12. The training set for initial estimation is 1980M1:1989M12. AIC and BIC are implemented with an upper limit of 12 lags. p-values are presented in parentheses with a null of equal predictive accuracy as the benchmark forecast from Clark & West (2007) tests. The Hit Ratio, HR, is defined as the percentage of forecasts which correctly identify the direction of movement in the price of wheat. p-values from Pesaran & Timmermann (1992) are provided beneath.

the measure of global real economic activity and the spread. We utilize the same recursive framework as has been previously utilized for AR and $ARMA$ models. These results show that the VAR does not perform well at forecasting out-of-sample, and increasing the lag length reduces performance. Some caution must be taken to interpret p-values for Table 4.2, as the p-values may overstate the significance of short-horizon MSPE reductions (Alquist et al. 2013). This can be seen that one of the results in Table 4.2, appears to be significant with a relative MSPE greater than 1. (Alquist et al. 2013) suggest interpreting marginally significant values with caution.

Furthermore with few exceptions, there is no systematic improvement in performance for identifying the future change in the real price of wheat across all models and all specifications. However there are two interesting findings, (i) forecasting in log levels appears to provide superior performance as compared to log-differences, and (ii) models with optimal lag-lengths chosen by information criterion outperform those of fixed lag specification.

4.4.4 Forecasting Real Wheat Price Changes using Dynamic Model Averaging and Dynamic Model Selection

Our final results in Table 4.3 examine the performance of Dynamic Models to forecast log differenced real wheat prices. Forecasts are evaluated under the same conditions as AR and $ARMA$ specifications for easy comparison. Relative MSPEs are computed over an evaluation period from 1990M1:2016M12, with p-values based upon Clark & West (2007) presented in parentheses.[2] The Hit Ratio as previously described is also presented.

The results for DMA with both forgetting factors set to 0.95 is presented in Row (3). As can been seen from the Table, this is the best performing model specification, with significant predictive gains at all horizons under evaluation. This specification places relatively more weight on recent observations and allows for more rapid changes in both model specification and parameter estimates over time. Notably, $DMA^{\alpha=\lambda=0.95}$ is the only method evaluated which also presents improvements over the benchmark no change for direction testing, with significant gains above 10% for all but the three quarter ahead forecast. The performance improvement in the MSPE is around 15% for horizons of 3, 6 and 9 months and 5% for 12 months.

To illustrate the improvements that $DMA^{\alpha=\lambda=0.95}$ provides over the no-change forecast, Figure 4.4.4 presents the realized log returns of wheat, compared to the forecast from the benchmark no-change and $DMA^{\alpha=\lambda=0.95}$ predicted values. We can see that $DMA^{\alpha=\lambda=0.95}$ captures some of the complex dynamics that wheat returns have undergone over the past two decades, particularly during the crisis period c.2008.

In comparison to DMA which averages across potential forecast models, DMS which selects the most promising candidate from all combinations at each point of time, performs

[2] Clark & West (2007) test statistics are correct here as DMA produces direct forecasts.

Table 4.3: Relative MSPE Ratios and Forecast Error Diagnostics for Wheat Returns from Dynamic Models

	$h=3$		$h=6$		$h=9$		$h=12$	
	MSPE	HR	MSPE	HR	MSPE	HR	MSPE	HR
$DMA^{\alpha=\lambda=0.99}$	0.99	**0.54**	**0.98**	**0.55**	**0.97**	**0.58**	1.00	0.53
	(0.11)	(0.06)	(0.05)	(0.05)	(0.01)	(0.00)	(0.22)	(0.11)
$DMS^{\alpha=\lambda=0.99}$	1.06	0.44	1.03	0.49	1.02	0.5	1.04	0.46
	(0.85)	(0.98)	(0.96)	(0.70)	(0.85)	(0.57)	(0.74)	(0.89)
$DMA^{\alpha=\lambda=0.95}$	**0.84**	**0.66**	**0.85**	**0.63**	**0.86**	**0.59**	**0.95**	**0.61**
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
$DMS^{\alpha=\lambda=0.95}$	1.19	0.49	1.15	0.47	1.19	0.46	1.16	0.50
	(0.59)	(0.58)	(0.87)	(0.87)	(0.97)	(0.90)	(0.40)	(0.57)
$DMA^{\alpha=0.95,\lambda=1}$	1.01	**0.55**	**0.98**	0.52	**0.96**	**0.56**	1.00	0.52
	(0.15)	(0.05)	(0.01)	(0.14)	(0.00)	(0.00)	(0.20)	(0.16)
BMA	1.01	0.51	1.00	0.49	1.00	0.47	1.00	0.47
	(0.90)	(0.40)	(0.13)	(0.71)	(0.22)	(0.84)	(0.35)	(0.88)
BMS	1.00	0.50	1.00	0.52	1.00	0.51	1.01	0.47
	(0.47)	(0.49)	(0.34)	(0.27)	(0.81)	(0.31)	(0.82)	(0.84)

Note: MSPE results are presented as a ratio relative to the MSPE of the no-change forecast model. The evaluation period under consideration is 1990M1:2016M12. The training set for initial estimation is 1980M1:1989M12. p-values are presented in parentheses with a null of equal predictive accuracy as the benchmark forecast from Clark & West (2007) tests. The Hit Ratio, HR, is defined as the percentage of forecasts which correctly identify the direction of movement in the price of wheat. p-values from Pesaran & Timmermann (1992) are provided beneath.

poorly. Regardless of the specification of forgetting factors, DMS performs worse than its DMA counterpart, and never outperforms the no-change forecast.

Comparison for Dynamic Model Averaging with alternative specifications shows that increasing λ to 1 (i.e., we allow for no variation in parameter estimates) reduces performance, as does increasing both α and λ (i.e., we allow reduced variation in both parameter estimates and model selection). This finding suggests that as wheat prices have exhibited high volatility over the past two decades, models which place relatively more weight on the most recent observations can significantly outperform models which have more stability in both model specification and parameter estimates.

Leading on from the previous results, variants of this methodology which do not allow for parameters or model specifications to vary over time perform poorly. BMA and BMS (equivalent to $DMA^{\alpha=\lambda=1}$ and $DMS^{\alpha=\lambda=1}$, respectively), do not allow the model to adjust

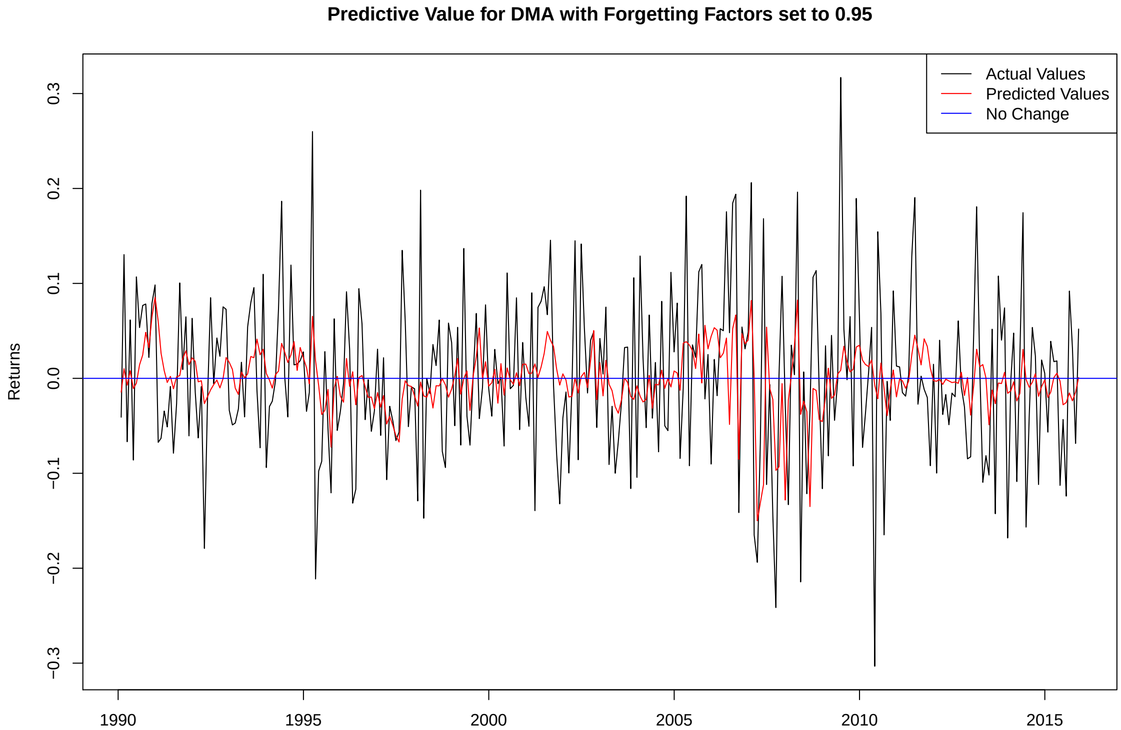

Figure 4.1: Predicted values from DMA with both forgetting factors set to 0.95 compared to the actual value and the no-change forecast.

to the volatile nature of wheat prices. The results for these are presented in Rows (6) and (7) of Table 4.3. In all cases these specifications underperformed compared to their equivalents with lower forgetting factors and the benchmark model.

4.4.5 Evaluation of variables included within Dynamic Model Averaging

DMA, with appropriate forgetting factors, appears to perform particularly well at forecasting wheat returns, given that it has time-varying parameters and the ability to modify the variables included in a forecasting model at each time period. In principle these benefits come from its ability to select parsimonious models which include fewer predictors. Here we follow Koop & Korobilis (2012) to determine how many and which of the variables are included in the DMA procedure. This allows greater insight into which variables are most important for forecasting the future path of wheat prices.

Let $Size_t^{(i)}$ be the number of chosen predictive variables, other than constant and other variables which are common to all models, in model i at time t. If the posterior probability that this model i should be included for forecasting at time t, conditional on the available information $\mathcal{F}_t$, is $p(M_i \mid \mathcal{F}_t)$, then the expected number of predictive variables at each point in time is given by:

$$E\left[Size_t\right] = \sum_{i=1}^{k} Size_t^{(i)} p(M_i \mid \mathcal{F}_t). \tag{4.45}$$

Figure 4.2 plots the expected size of $DMA^{\alpha=\lambda=0.95}$ for horizons of 3, 6, 9 and 12 months ahead. These plots give an indication of the shrinkage which DMA undergoes and show that; (i) the average model size is relatively stable across the evaluation period with one notable exception and (ii), DMA selects relatively parsimonious models with a strong preference for models with between 4 and 6 of the potential explanatory variables at all time periods. Notably during the crisis period around 2008 there is a rise in the number of included predictive variables which is most pronounced in the one-quarter ahead forecasts, but also appears in the longest horizon forecasts of one year.

However these plots do not reveal which variables are the most important, nor how they vary over time. The posterior inclusion probabilities for each variable can be computed to allow further insight into the variables which are most important for forecasting wheat prices, and how they vary across time and forecast horizon. As described in Koop & Korobilis (2012) and Catania & Nonejad (2018) for each predictor the posterior inclusion probability can be calculated as the weight that DMA assigns to models containing a given predictor, i.e., for each time t, the posterior probability of inclusion is given by:

$$\sum_{i=1}^{k} 1_{(i \subset m)} p(M_i \mid \mathcal{F}_t), \tag{4.46}$$

where $1_{(i \subset m)}$ is an indicator function taking the value of either 0 or 1 and m, $m \in \{1, \ldots, n\}$ is the mth predictor. To look at the variables included in both short and long-term forecasts, Figures (4.3) and (4.4) present the posterior inclusion probabilities for all variables for 3 and 12 months ahead, respectively. These figures demonstrate why dynamic models which vary model specification outperform those models with restrictive fixed model specifications. Over time the included model variables varies widely and the posterior inclusion probability

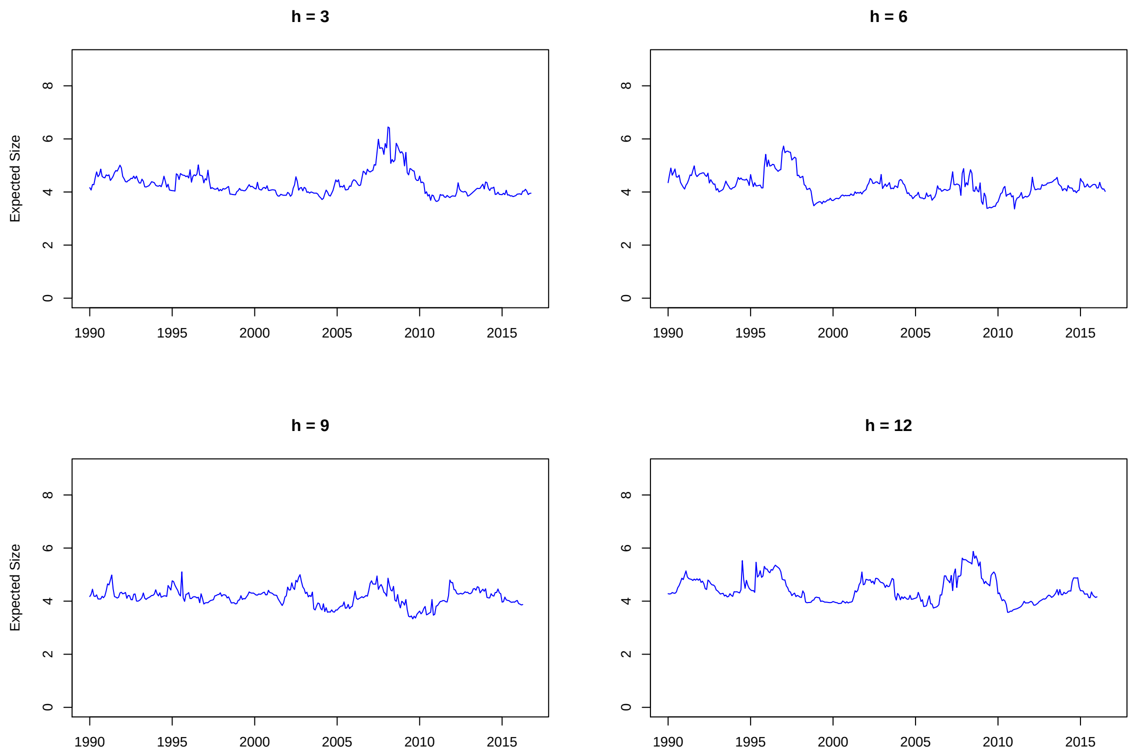

Figure 4.2: Expected size of model at each time period.

changes. A few key conclusions can be drawn from these graphs.

First examining Figure 4.3, we see that most variables are relatively stable with their inclusion probabilities for the first half of the out of sample evaluation period. The period around 2008 has a notable increase in the probability of inclusion for a number of variables. The other included agricultural commodities suddenly become much more likely to be included. First corn, then oats and finally soybeans become increasingly important for predicting wheat returns during this time horizon during the period when prices suddenly increased rapidly. That agricultural commodities comove has long been known as mentioned in previous chapters, however this sudden increase in predictive power captures the hugely increased comovement during this crisis time period. Another interesting result is the rapid increase in likelihood that one of the monetary measures, t-bills, is included in the forecasting models. Post crisis we see that the inclusion probabilities of all variables returns closer to a baseline level.

Turning to the longer forecast horizon results presented in Figure (4.4), a slightly different picture emerges. Other commodity prices are driving long horizon wheat prices throughout the mid-2000s. This result may be indicative that comovement played an increasing roll in price determination during this early period, potentially as a result of increased 'financialization'. Once again during the crisis period there is a sudden, and sustained, increase in the posterior probability of inclusion of two of the three agricultural commodity returns, corn and soybeans. From these results it appears that particularly during crisis periods there is a herding effect and that other commodities played a large roll in determining the price of wheat during this period.

Also of note from these diagrams is that at no point do the real effective exchange rates of either of the major exporting countries included in our sample become important. Similarly, Real economic activity never increases during the crisis period. This later result particularly is additional evidence that increased global real economic activity, a proxy trying to capture increased demand from developing countries such as India and China, did not have a large impact on wheat prices during the period 2007-2008.

4.5 Conclusion

This chapter has provided an extensive evaluation of how forecastable wheat prices have been over the past two decades. This has been accomplished by implementing a wide variety of parsimonious and dynamic econometric techniques to global wheat price series. As the evaluation period considered covers January 1990 until December 2016, it also covers the recent crisis period where prices underwent increased volatility and rapidly rose and fell. Forecasting future wheat prices during these periods will be of particular interest to producers, consumers and policymakers due to the potential negative impacts unexpected movements of food prices may have.

We present a few key results. First, in common with other commodity prices, simple approaches perform particularly poorly in comparison to random walk forecasts which are used as a benchmark. These models struggle to capture volatile dynamics, including periods of explosive behavior, which have characterized agricultural commodity markets in the recent past. However, dynamic models which allow for variation in both the variables included in

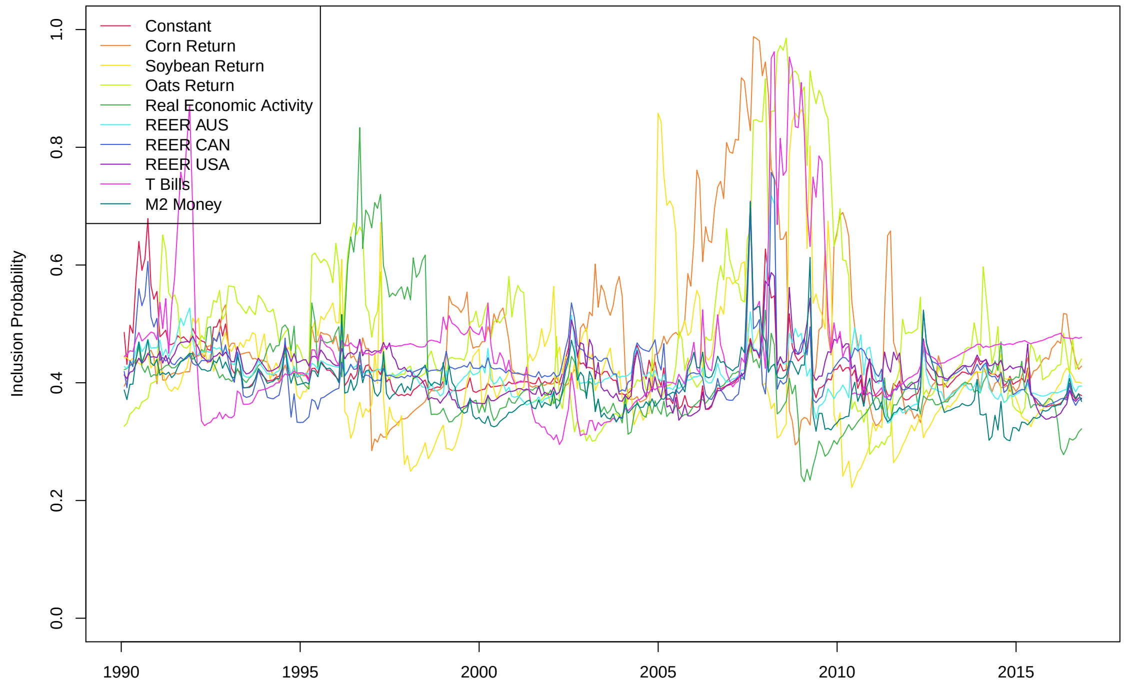

Figure 4.3: Inclusion probabilities from DMA with both forgetting factors set to 0.95 for a 3-month horizon.

the underlying linear specification, and time-varying parameters for these variables, provide superior forecast performance in comparison to other models. The DMA model with values on both forgetting parameters which are relatively low performs particularly well, as it significantly outperforms a no-change forecast at all tested horizons. The low values of both forgetting factors implies that the structure of recent market dynamics has been turbulent and so both parameter estimates and model structure must respond accordingly to improve forecast performance. This superior performance continues when examining the performance of models for predicting the future direction of wheat price changes.

The poor performance of parsimonious forecasts is not unexpected, particularly given the rapid change in prices that agricultural commodity prices have undergone in the sample period under investigation. We add to the literature which has found mixed performance of futures markets predictive content. Within this time frame futures markets were not able to consistently outperform a naive benchmark. Interestingly, all futures market models were

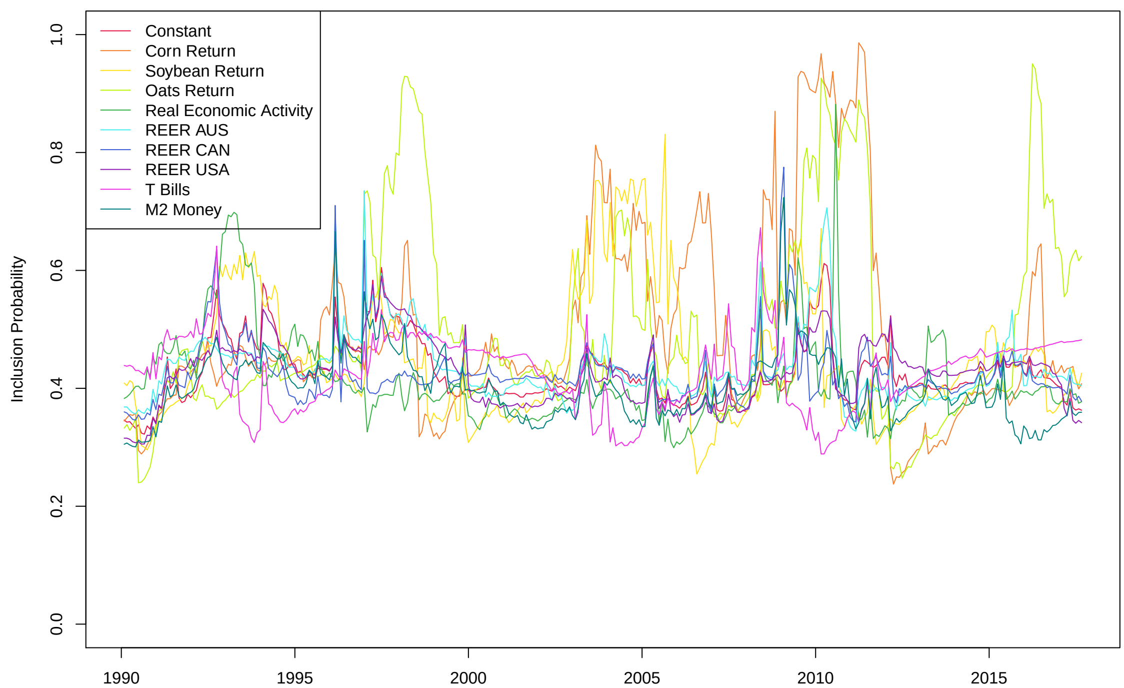

Figure 4.4: Inclusion Probabilities from DMA with both forgetting factors set to 0.95 for a 12-month horizon.

no better than random chance at predicting the direction of movement of wheat prices. Dynamic model averaging has presented excellent forecast performance in other applications and its ability to capture the underlying market dynamics lead to improvements over benchmark models.

The DMA results also allow insight into factors which have driven global wheat prices over the past two decades. During the sudden increase around 2008, changes in other grain prices become increasingly likely to be included in forecasting models, particularly for short horizons. This suggests that comovement among the asset class increased, which is in line with the view that financialization has had an impact on agricultural markets in general. Before the crisis these factors were no more likely than others to increase forecast performance. Our results also show that other factors which have been suggested as driving recent high commodity prices, such as global real economic activity, appear to have little predictive power. Overall, there is no one predictive variable which has a consistently high probability of inclusion within the dynamic model for it's out of sample forecast performance. Additionally we can analyze the shrinkage characteristics of the model over time. Relatively parsimonious models are preferred, with the optimal number of parameters remaining very stable, though it does rise during the period when prices peaked.

The best models for predicting global wheat prices are those which can adjust both the predictors which they include within the models, and the predicted estimates. These findings are important as this class of models appears to be the only one which can consistently produce out-of-sample performance which exceeds that of the random walk, in terms of both error size and directionality.

Chapter 5

Concluding Remarks

This thesis investigates the features and time series properties of global agricultural commodity prices. To do so the principle research areas examined include: (i) examining the explosive properties of global agricultural commodity prices and investigating for the presence of rational speculative bubbles; (ii) construction of a Bayesian structural vector autoregressive model and applying it to agricultural commodity prices; (iii) examining the relative importance of economic factors to wheat price dynamics and (iv) forecasting wheat price movements using a variety of econometric methods.

The first topic is examined in Chapter 2. This chapter analyses the behavior of four global agricultural commodity markets, utilizing the recently developed recursive unit root testing techniques of Phillips et al. (2011) and Phillips et al. (2015*a*). To overcome a joint hypothesis problem which arises due to the unobserved underlying fundamentals, a new time series is constructed following Pavlidis et al. (2017). Two key conclusions arise from this analysis, (i) global agricultural commodities have experienced explosive dynamics during the crisis period and (ii) contrary to previous findings in the literature, there is little evidence of rational speculative bubbles, when testing within a framework which controls for the underlying fundamental. This finding is suggestive that the rapid price changes witnessed in these markets were driven by some (unobserved) fundamental factor(s), and not by speculative bubbles. This finding further motivates Chapters 3 and 4 which investigate causal factors driving wheat price movements.

The second and third points are addressed by Chapter 3, where a Bayesian Structural VAR as proposed by Baumeister & Hamilton (2015*a*) is constructed for global wheat prices. The principle finding is that the majority of observed wheat price movements are not explained by either external commodity comovement (oil prices), or from increased global economic activity. These are two key contributory explanatory factors which have been previously proposed by the literature. This finding is suggestive that global wheat prices were not driven upwards by either index investors increasing demand for commodities, or that increased demand from India and China drove prices higher. The majority of wheat price movement is driven by residual wheat market specific supply and demand factors. However, during 2008 some of the rapid movements can be explained by precautionary demand.

Finally, Chapter 4 investigates the ability of different econometric models to forecast wheat prices over the past two decades and provides further insight into particular variables which may have contributed to wheat price dynamics. This chapter makes the following contributions; first, modern dynamic methodologies which allow for both time-varying parameters and model specifications significantly outperform almost all parsimonious modeling techniques, including the random walk. Second, when analyzing which variables are included over time, there is no single factor which has a particularly large probability of inclusion and therefore stands out as a key contributing factor to wheat price movements during the period under consideration.

Bibliography

Adjemian, M., Garcia, P., Irwin, S. & Smith, A. (2013), 'Solving the commodity marketsâĂŹ non-convergence puzzle', *Amber Waves:The Economics of Food, Farming, Natural Resources, and Rural America* (07).

Ahumada, H. & Cornejo, M. (2015), 'Explaining commodity prices by a cointegrated time series-cross section model', *Empirical Economics* **48**(4), pp. 1667–1690.

Ahumada, H. & Cornejo, M. (2016*a*), 'Forecasting Food Prices: The Case of Corn, Soybeans and Wheat', *International Journal of Forecasting* **32**, pp. 838–848.

Ahumada, H. & Cornejo, M. (2016*b*), 'Out-of-sample testing price discovery in commodity markets: the case of soybeans', *Agricultural Economics* **47**(6), pp. 709–718.

Ai, C., Chatrath, A. & Song, F. (2006), 'On the Comovement of Commodity Prices', *American Journal of Agricultural Economics* **88**(3), pp. 574–588.

Ajanovic, A. (2011), 'Biofuels versus Food Production: Does Biofuels Production Increase Food Prices?', *Energy* **36**(4), pp. 2070–2076.

Alexander, C. (2012), *Hedging the Risk of an Energy Futures Portfolio*, John Wiley and Sons, Ltd, chapter 9, pp. 117–127.

Alquist, R. & Kilian, L. (2010), 'What do we learn from the price of crude oil futures?', *Journal of Applied Econometrics* **25**(4), pp. 539–573.

Alquist, R., Kilian, L. & Vigfusson, R. (2013), Forecasting the price of oil, *in* G. Elliott & A. Timmermann, eds, 'Handbook of Economic Forecasting', Vol. 2, Elsevier, pp. 427–507.

Australian Bureau of Statistics (2008), 'Agricultural commodities, Australia, 2006-07'.
URL: *goo.gl/78uhnC*

Aye, G., Gupta, R., Hammoudeh, S. & Kim, W. J. (2015), ‘Forecasting the price of gold using dynamic model averaging’, *International Review of Financial Analysis* **41**, pp. 257–266.

Baffes, J. & Dennis, A. (2013), Long-term drivers of food prices, Policy Research Working Paper 6455, World Bank.

Barberis, N., Shleifer, A. & Wurgler, J. (2005), ‘Comovement’, *Journal of Financial Economics* **75**(2), pp. 283–317.

Baumeister, C. & Hamilton, J. (2015*a*), ‘Sign Restrictions, Structural Vector Autoregressions, and Useful Prior Information’, *Econometrica* **83**(5), pp. 1963–1999.

Baumeister, C. & Hamilton, J. (2015*b*), Optimal Inference about Impulse-Response Functions and Historical Decompositions in Incompletely Identified Structural Vector Autoregressions, Working Paper No. 24167, National Burea of Economic Research.

Baumeister, C. & Hamilton, J. (2018), ‘Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks’, *American Economic Review* . (Forthcoming).

Baumeister, C. & Kilian, L. (2012), ‘Real-time forecasts of the real price of oil’, *Journal of Business and Economic Statistics* **30**(2), pp. 326–336.

Baumeister, C. & Peersman, G. (2010), Sources of the volatility puzzle in the crude oil market, Working Paper 2010/634, University of Ghent.

Bellemare, M. (2015), ‘Rising Food Prices, Food Price Volatility, and Social Unrest’, *American Journal of Agricultural Economics* **97**(1), pp. 1–21.

Berazneva, J. & Lee, D. R. (2013), ‘Explaining the african food riots of 2007âĂŞ2008: An empirical analysis’, *Food Policy* **39**, pp. 28–39.

Blanchard, O. (1979), ‘Speculative Bubbles, Crashes and Rational Expectations’, *Economics Letters* **3**(4), pp. 387–389.

Blanchard, O. J. & Watson, M. W. (1982), 'Bubbles, rational expectations and financial markets'.

Bobenrieth, E., Bobenrieth, J. & Wright, B. (2014), Bubble troubles? rational storage, mean reversion, and runs in commodity prices, *in* J. Chavas, D. Hummels & B. Wright, eds, 'The Economics of Food Price Volatility', University of Chicago Press, pp. 193–208.

Buncic, D. & Moretto, C. (2015), 'Forecasting copper prices with dynamic averaging and selection models', *The North American Journal of Economics and Finance* **33**, pp. 1–38.

Carter, C. A., Rausser, G. C. & Smith, A. (2016), 'Commodity storage and the market effects of biofuel policies', *American Journal of Agricultural Economics* **99**(4), pp. 1027–1055.

Catania, L. & Nonejad, N. (2017), *eDMA: Dynamic Model Averaging with Grid Search*. R package version 1.4-0.
URL: *https://CRAN.R-project.org/package=eDMA*

Catania, L. & Nonejad, N. (2018), 'Dynamic Model Averaging for Practitioners in Economics and Finance: The eDMA Package', *Journal of Statistical Software, Articles* **84**(11), pp. 1–39.

Cheah, E.-T. & Fry, J. (2015), 'Speculative bubbles in bitcoin markets? an empirical investigation into the fundamental value of bitcoin', *Economics Letters* **130**, pp. 32–36.

Chen, Y., Rogoff, K. & Rossi, B. (2010), 'Can exchange rates forecast commodity prices?', *The Quarterly Journal of Economics* **125**(3), pp. 1145–1194.

Cheng, I. & Xiong, W. (2014), 'Financialization of commodity markets', *Annual Review of Financial Economics* **6**, pp. 419–441.

Clark, T. & West, K. (2007), 'Approximately normal tests for equal predictive accuracy in nested models', *Journal of Econometrics* **138**(1), pp. 291–311.

Dale, R., Johnson, J. & Tang, L. (2005), 'Financial markets can go mad: evidence of irrational behaviour during the south sea bubble1', *The Economic History Review* **58**(2), pp. 233–271.

Daviron, B., Nango Dembele, N., Murphy, S. & Rashid, S. (2011), Price volatility and food security. A report by the high level panel of experts on food security and nutrition of the committee on world food security, Technical report.

De Hoyos, R. E. & Medvedev, D. (2009), *Poverty effects of higher food prices: a global perspective*, The World Bank.

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

Dedola, L. & Neri, S. (2007), 'What Does a Technology Shock Do? A VAR Analysis with Model-Based Sign Restrictions', *Journal of Monetary Economics* **54**, pp. 512–549.

Delle Chiaie, S., Ferrara, L. & Giannone, D. (2017), Common factors of commodity prices, Banque de France Working Paper No. 645, Banque de France.

Diba, B. & Grossman, H. (1988), 'Explosive Rational Stock Bubbles in Stock Prices?', *American Economic Review* **78**(3), pp. 520–530.

Dickey, D. & Fuller, W. (1979), 'Distribution of the Estimators for Autoregressive Time Series With a Unit Root', *Journal of American Statistical Association* **74**(366), pp. 427–431.

Doan, T., Litterman, B. & Sims, C. (1984), 'Forecasting and Conditional Projection Using Realistic Prior Distributions', *Econometric Reviews* **3**(1), pp. 1–100.

Drachal, K. (2016), 'Forecasting spot oil price in a dynamic model averaging framework - have the determinants changed over time?', *Energy Economics* **60**, pp. 35–46.

Elliott, G. & Timmermann, A. (2008), 'Economic Forecasting', *Journal of Economic Literature* **46**(1), pp. 3–56.

Engsted, T. & Nielsen, B. (2012), 'Testing for rational bubbles in a coexplosive vector autoregression', *The Econometrics Journal* **15**(2), pp. 226–254.

Etienne, X., Irwin, S. & Garcia, P. (2014), 'Bubbles in food commodity markets: Four decades of evidence', *Journal of International Money and Finance* **42**, pp. 129–155.

Etienne, X., Irwin, S. & Garcia, P. (2015), 'Price Explosiveness, Speculation, and Grain Futures Prices', *American Journal of Agricultural Economics* **97**(1), pp. 65–87.

Etienne, X., Irwin, S. & Garcia, P. (2017), 'New Evidence that Index Traders Did Not Drive Bubbles in Grain Futures Markets', *Journal of Agricultural and Resource Economics* **42**(1), pp. 45–67.

Evans, G. (1991), 'Pitfalls in Testing for Explosive Bubbles in Asset Prices', *American Economic Review* **81**(4), pp. 992–930.

Fama, E. & French, K. (1987), 'Commodity Futures Prices: Some Evidence on Forecast Power, Premiums and the Theory of Storage', *The Journal of Business* **60**(1), pp. 55–73.

Fernandez-Perez, A., Frijns, B. & Tourani-Rad, A. (2016), 'Contemporaneous interactions among fuel, biofuel and agricultural commodities', *Energy Economics* **58**, pp. 1–10.

Frankel, J. & Rose, A. (2010), Determinants of agricultural and mineral commodity prices, HKS Faculty Research Working Paper Series RWP10-038, John F. Kennedy School of Government, Harvard University.

Garcia, P., Irwin, S. & Smith, A. (2015), 'Futures market failure?', *American Journal of Agricultural Economics* **97**(1), pp. 40–64.

Gerlow, M. E., Irwin, S. H. & Liu, T.-R. (1993), 'Economic Evaluation of Commodity Price Forecasting Models', *International Journal of Forecasting* **9**(3), pp. 387–397.

Gilbert, C. (2010), 'Speculative Influences on Commodity Futures Prices 2006-2008', United Nations Conference on Trade and Development Discussion Paper Number 197, United Nations.

Gutierrez, L. (2013), 'Speculative Bubbles in Agricultural Commodity Markets', *European Review of Agricultural Economics* **40**(2), pp. 217–238.

Hafner, C. M. (2018), 'Testing for Bubbles in Cryptocurrencies with Time-Varying Volatility', *Journal of Financial Econometrics* .

Hamilton, J. D. (2013), Oil prices, exhaustible resources, and economic growth, *in* R. Fouqet, ed., 'Handbook on Energy and Climate Change', Edward Elgar Publishing, Cheltenham, United Kingdom, pp. 29–63.

Hamilton, J. D. & Wu, J. C. (2015), 'Effects of index-fund investing on commodity futures prices', *International Economic Review* **56**(1), pp. 187–205.

Harvey, D. I., Leybourne, S. J., Sollis, R. & Taylor, A. R. (2016), 'Tests for explosive financial bubbles in the presence of non-stationary volatility', *Journal of Empirical Finance* **38**(Part B), pp. 548–574.

Headey, D. (2011), 'Rethinking the Global Food Crisis: The Role of Trade Shocks', *Food Policy* **36**(2), pp. 136–146.

Headey, D. & Fan, S. (2008), 'Anatomy of a Crisis: The Causes and Consequences of Surging Food Prices', *Agricultural Economics* **39**(s1), pp. 375–391.

Headey, D. & Fan, S. (2010), Reflections on the global food crisis, IFPRI Research Monograph No. 165, International Food Policy Research Institute.

Hendry, D. (2006), 'Robustifying forecasts from equilibrium-correction systems', *Journal of Econometrics* **135**(1), pp. 399–426.

Hochman, Z., Gobbett, D. & Horan, H. (2017), 'Climate trends account for stalled wheat yields in Australia since 1990', *Global Change Biology* **23**(5), pp. 2071–2081.

Irwin, S. (2013), 'Commodity Index Investment and Food Prices: Does the "Masters Hypothesis" Explain Recent Price Spikes?', *Agricultural Economics* **44**(s1), pp. 29–41.

Irwin, S. & Sanders, D. (2011), 'Testing the Masters Hypothesis in commodity futures markets', *Energy Economics* **34**, pp. 256–269.

Ivanic, M. & Martin, W. (2008), 'Implications of Higher Global Food Prices for Poverty in Low-Income Countries', *Agricultural Economics* **39**(s1), pp. 405–416.

Ivanic, M., Martin, W. & Zaman, H. (2012), 'Estimating the Short-Run Poverty Impacts of the 2010-11 Surge in Food Prices', *World Development* **40**(11), pp. 2302–2317.

Janzen, J. P., Carter, C. A., Smith, A. D. & Adjemian, M. K. (2014), Deconstructing Wheat Price Spikes: A Model of Supply and Demand, Financial Speculation, and Commodity Price Comovement, USDA-ERS Economic Research Report Number ERR-165, U.S. Department of Agriculture, Economic Research Service.

Joseph, K., Irwin, S. & Garcia, P. (2016), 'Commodity Storage under Backwardation: Does the Working Curve Still Work?', *Applied Economic Perspective and Policy* **38**(1), pp. 152–173.

Just, R. E. & Rausser, G. C. (1981), 'Commodity price forecasting with large-scale econometric models and the futures market', *American Journal of Agricultural Economics* **63**(2), pp. 197–208.

Kaldor, N. (1939), 'Speculation and Economic Stability', *The Review of Economic Studies* **7**(1), pp. 1–27.

Kaldor, N. (1940), 'A Note on the Theory of the Forward Market', *The Review of Economic Studies* **7**(3), pp. 196–201.

Kilian, L. (2009), 'Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market', *American Economic Review* **99**(3), pp. 1053–1069.

Kilian, L. (2015), Structural vector autoregressions, *in* N. Hashimzade & M. Thornton, eds, 'Handbook of Research Methods and Applications in Empirical Macroeconomics', Edward Elgar Publishing, Cheltenham, United Kingdom, pp. 515–554.

Kilian, L. & Lütkepohl, H. (2017), *Structural vector autoregressive analysis*, Cambridge University Press.

Kilian, L. & Murphy, D. (2012), 'Why agnostic sign restrictions are not enough: understanding the dynamics of oil market VAR models', *Journal of the European Economic Association* **10**(5), pp. 1166–1188.

Kilian, L. & Murphy, D. (2014), 'The role of inventories and speculative trading in the global market for crude oil', *Journal of Applied Econometrics* **29**(3), pp. 454–478.

Kilian, L. & Zhou, X. (2018), 'Modeling fluctuations in the global demand for commodities', *Journal of International Money and Finance* **88**, pp. 54–78.

Kivedal, B. K. (2013), 'Testing for rational bubbles in the us housing market', *Journal of Macroeconomics* **38**, pp. 369–381.

Koop, G. & Korobilis, D. (2012), 'Forecasting Inflation using Dynamic Model Averaging', *International Economic Review* **53**(3), pp. 867–886.

Leitch, G. & Tanner, J. E. (1991), 'Economic forecast evaluation: profits versus the conventional error measures', *American Economic Review* **81**(3), pp. 580–590.

Marcellino, M., Stock, J. & Watson, M. (2006), 'A comparison of direct and iterated multistep AR methods for forecasting macroeconomic time series', *Journal of Econometrics* **135**(1-2), pp. 499–526.

Masters, M. (2008*a*), 'Testimony before the committee on homeland security and governmental affairs'. U.S. Senate, Washington, 20 May, 2008.

Masters, M. W. (2008*b*), 'Testimony before the committee on homeland security and governmental affairs', *US Senate, Washington, May* **20**.

McPhail, L., Du, X. & Muhammad, A. (2012), 'Disentangling corn price volatility: The role of global demand, speculation and energy', *Journal of Agricultural and Applied Economics* **44**(3), pp. 401–410.

Naser, H. (2016), 'Estimating and forecasting the real prices of crude oil: A data rich model using a dynamic model averaging (dma) approach', *Energy Economics* **56**, pp. 75–87.

Neri, S., Nobili, A., Conti, A. M. et al. (2017), Low inflation and monetary policy in the Euro area, ECB Working Paper No. 2005, European Central Bank.

Pavlidis, E., Paya, I. & Peel, D. (2017), 'Testing for Speculative Bubbles using Spot and Forward Prices', *International Economic Review* **58**(4), pp. 1191–1226.

Pesaran, M. & Timmermann, A. (1992), 'A simple nonparametric test of predictive performance', *Journal of Business & Economic Statistics* **10**(4), pp. 461–465.

Phillips, P., Shi, S. & Yu, J. (2015*a*), 'Testing for Multiple Bubbles: Historical Episodes of Exuberance and Collapse in the S&P 500', *International Economic Review* **56**(4), pp. 1043–1078.

Phillips, P., Shi, S. & Yu, J. (2015*b*), 'Supplement to Two Papers on Multiple Bubbles'.
URL: *goo.gl/YEGdgD*

Phillips, P., Wu, Y. & Yu, J. (2011), 'Explosive Behaviour in the 1990s NASDAQ: When did Exuberance Escalate Asset Values?', *International Economic Review* **52**(1), pp. 201–226.

Phillips, P. & Yu, J. (2011), 'Dating the timeline of financial bubbles during the subprime crisis.', *Quantitative Economics* **2**(3), pp. 455–491.

Pindyck, R. (1993), 'The Present Value Model of Rational Commodity Pricing', *The Economic Journal* **103**(418), pp. 511–530.

Pindyck, R. & Rotemberg, J. (1990), 'The Excess Co-Movement of Commodity Prices', *Economic Journal* **100**(403), pp. 1173–1189.

Raftery, A., Kárnỳ, M. & Ettler, P. (2010), 'Online prediction under model uncertainty via dynamic model averaging: Application to a cold rolling mill', *Technometrics* **52**(1), pp. 52–66.

Reichsfeld, D. & Roache, S. (2011), Do Commodity Futures Help Forecast Spot Prices?, IMF Working Paper WP/11/254, International Monetary Fund.

Rosegrant, M. (2008), *Biofuels and grain prices: impacts and policy responses*, International Food Policy Research Institute Washington, DC.

Rubio-Ramírez, J., Waggoner, D. & Arias, J. (2014), Inference based on SVAR identified with sign and zero restrictions: Theory and applications, CEPR Discussion Paper No. DP9796, Centre for Economic Policy Research.

Rubio-Ramírez, J., Waggoner, D. & Zha, T. (2010), 'Structural Vector Autoregressions: Theory of Identificaton and Algorithms for Inference', *The Review of Economic Studies* **77**(2), pp. 665–696.

Sanders, D. & Irwin, S. (2010), 'A speculative bubble in commodity futures prices? Cross-sectional evidence', *Agricultural Economics* **41**, pp. 25–32.

Sanders, D. & Irwin, S. (2017), 'Bubbles, Froth and Facts: Another Look at the Masters Hypothesis in Commodity Futures Markets', *Journal of Agricultural Economics* **68**(2), pp. 345–365.

Shewry, P. R. & Hey, S. J. (2015), 'The contribution of wheat to human diet and health', *Food and energy security* **4**(3), pp. 178–202.

Sims, C. (1980), 'Macroeconomics and Reality', *Econometrica* **48**(1), pp. 1–48.

Sims, C. & Zha, T. (1998), 'Bayesian methods for dynamic multivariate models', *International Economic Review* **39**(4), pp. 949–968.

Stock, J. & Watson, M. (2003), 'Forecasting output and inflation: The role of asset prices', *Journal of Economic Literature* **41**(3), pp. 788–829.

Tang, K. & Xiong, W. (2012), 'Index Investment and the Financialization of Commodities', *Financial Analysts Journal* **68**(6), pp. 54–74.

Timmer, P. (2008), Causes of high food prices, ADB Economics Working Paper Series No. 128, Asian Development Bank.

Tomek, W. G. (1997), 'Commodity futures prices as forecasts', *Review of Agricultural Economics* **19**(1), pp. 23–44.

Uhlig, H. (2005), 'What are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure', *Journal of Monetary Economics* **52**(2), pp. 381–419.

Van Robays, I. (2012), Macroeconomic uncertainty and the impact of oil shocks, ECB Working Paper No. 1479, European Central Bank.

Vargas-Silva, C. (2008), 'Monetary policy and the US housing market: A VAR analysis imposing sign restrictions', *Journal of Macroeconomics* **30**(3), pp. 977–990.

Vocke, G. (2012), Wheat Year in Review (Domestic): Higher Domestic Use and Exports Lower 2010/11 Ending Stocks, ERS Report WHS-2011, United States Department of Agriculture.

Von Braun, J. (2008), 'The food crisis isn't over', *Nature* **456**(7223), 701.

Vranken, L., Avermaete, T., Petalios, D. & Mathijs, E. (2014), 'Curbing global meat consumption: Emerging evidence of a second nutrition transition', *Environmental Science & Policy* **39**, pp. 95–106.

Working, H. (1927), 'Forecasting the price of wheat', *Journal of Farm Economics* **9**(3), pp. 273–287.

Working, H. (1933), 'Price Relations between July and September Wheat Futures at Chicago since 1885', *Wheat Studies of the Food Research Institute* **9**(6), pp. 187–238.

Working, H. (1949), 'The Theory of Price of Storage', *American Economic Review* **39**(6), pp. 1254–1262.

The World Bank (2018), 'World development indicators: Structure of output'. data retrieved from World Development Indicators, `http://wdi.worldbank.org/table/4.2#`.